

Release Notes: Junos[®] OS Release 19.4R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, JRR Series, and Junos Fusion

21 April 2023

Contents	Introduction 10
	New Features in 19.4R2 10
	Junos OS Release Notes for ACX Series 12
	What's New 12
	What's New in Release 19.4R2-S2 13
	What's New in Release 19.4R2 13
	What's New in Release 19.4R1 16
	What's Changed 20
	What's Changed in 19.4R2-S4 20
	What's Changed in 19.4R2-S2 20
	What's Changed in 19.4R2 20
	What's Changed in 19.4R1 21
	Known Limitations 22
	General Routing 23
	Open Issues 24
	General Routing 24
	Layer 2 Ethernet Services 25
	Multicast 26
	Routing Protocols 26

Resolved Issues | 26**Resolved Issues: 19.4R2 | 27****Resolved Issues: 19.4R1 | 28****Documentation Updates | 30****Feature Guides Are Renamed As User Guides | 31****Migration, Upgrade, and Downgrade Instructions | 31****Upgrade and Downgrade Support Policy for Junos OS Releases | 31****Junos OS Release Notes for EX Series Switches | 32****What's New | 33****What's New in 19.4R2 | 33****What's New in 19.4R1 | 34****What's Changed | 43****What's Changed in 19.4R2-S4 | 44****What's Changed in 19.4R2 | 44****What's Changed in 19.4R1 | 45****Known Limitations | 47****Open Issues | 47****Authentication and Access Control | 48****EVPN | 48****General Routing | 48****Infrastructure | 49****Interfaces and Chassis | 49****Layer 2 Features | 50****Platform and Infrastructure | 50****Routing Protocols | 50****Resolved Issues | 51****Class of Service (CoS) | 51****EVPN | 52****Forwarding and Sampling | 52****General Routing | 52****Infrastructure | 53****Interfaces and Chassis | 53****Junos Fusion for Enterprise | 53****Junos Fusion Satellite Software | 54**

Layer 2 Features	54
Layer 2 Ethernet Services	54
MPLS	54
Platform and Infrastructure	54
Routing Protocols	55
User Interface and Configuration	55
Virtual Chassis	55
Documentation Updates	55
Feature Guides Are Renamed As User Guides	56
Migration, Upgrade, and Downgrade Instructions	56
Upgrade and Downgrade Support Policy for Junos OS Releases	56
Junos OS Release Notes for JRR Series	57
What's New	58
What's New in 19.4R2 Release	58
What's New in 19.4R1 Release	58
What's Changed	59
Known Limitations	59
Open Issues	60
Resolved Issues	60
Resolved Issues: 19.4R2	61
Resolved Issues: 19.4R1	61
Documentation Updates	61
Feature Guides Are Renamed As User Guides	62
Migration, Upgrade, and Downgrade Instructions	62
Upgrade and Downgrade Support Policy for Junos OS Releases	62
Junos OS Release Notes for Junos Fusion Enterprise	63
What's New	64
What's Changed	64
Known Limitations	65
Open Issues	65
Junos Fusion for Enterprise	66
Resolved Issues	66
Resolved Issues: Release 19.4R2	67

Documentation Updates | 67

Feature Guides Are Renamed As User Guides | 67

Migration, Upgrade, and Downgrade Instructions | 68

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 68

Upgrading an Aggregation Device with Redundant Routing Engines | 70

Preparing the Switch for Satellite Device Conversion | 71

Converting a Satellite Device to a Standalone Switch | 72

Upgrade and Downgrade Support Policy for Junos OS Releases | 72

Downgrading from Junos OS | 73

Junos OS Release Notes for Junos Fusion Provider Edge | 74

What's New | 74

What's Changed | 75

Known Limitations | 75

Open Issues | 76

Resolved Issues | 76

Resolved Issues: 19.4R2 Release | 76

Resolved Issues: 19.4R1 Release | 77

Documentation Updates | 77

Feature Guides Are Renamed As User Guides | 77

Migration, Upgrade, and Downgrade Instructions | 78

Basic Procedure for Upgrading an Aggregation Device | 78

Upgrading an Aggregation Device with Redundant Routing Engines | 81

Preparing the Switch for Satellite Device Conversion | 81

Converting a Satellite Device to a Standalone Device | 83

Upgrading an Aggregation Device | 85

Upgrade and Downgrade Support Policy for Junos OS Releases | 86

Downgrading from Junos OS Release 19.4 | 86

Junos OS Release Notes for MX Series 5G Universal Routing Platform | 87

What's New | 87

What's New in 19.4R2 Release | 88

What's New in 19.4R1 Release | 89

What's Changed | 110

What's Changed in 19.4R2-S4 | 111

What's Changed in 19.4R2-S2 Release | 111

What's Changed in 19.4R2-S1 Release	111
What's Changed in 19.4R2 Release	111
What's Changed in 19.4R1 Release	112
Known Limitations	116
General Routing	116
Interfaces and Chassis	118
Platform and Infrastructure	119
Routing Protocols	119
Open Issues	119
Class of Service (CoS)	120
EVPN	120
Forwarding and Sampling	121
General Routing	121
Interfaces and Chassis	128
Layer 2 Ethernet Services	130
Layer 2 Features	130
MPLS	130
Network Management and Monitoring	131
Platform and Infrastructure	131
Routing Protocols	133
Services Applications	134
VPNs	134
Resolved Issues	135
Resolved Issues: 19.4R2	135
Resolved Issues: 19.4R1	147
Documentation Updates	168
Advanced Subscriber Management Provider	169
Feature Guides Are Renamed As User Guides	169
Migration, Upgrade, and Downgrade Instructions	169
Basic Procedure for Upgrading to Release 19.4	170
Procedure to Upgrade to FreeBSD 11.x based Junos OS	170
Procedure to Upgrade to FreeBSD 6.x based Junos OS	173
Upgrade and Downgrade Support Policy for Junos OS Releases	175
Upgrading a Router with Redundant Routing Engines	175

Downgrading from Release 19.4	176
Junos OS Release Notes for NFX Series	176
What's New	177
General routing	178
Hardware	178
Architecture	178
What's Changed	179
What's Changed in 19.4R2 Release	179
What's Changed in 19.4R1 Release	180
Known Limitations	180
Interfaces	181
Platform and Infrastructure	181
Open Issues	181
Interfaces	182
Platform and Infrastructure	182
Virtual Network Functions (VNFs)	183
Resolved Issues	183
Resolved Issues: 19.4R2	184
Resolved Issues: 19.4R1	185
Documentation Updates	188
Feature Guides Are Renamed As User Guides	188
Migration, Upgrade, and Downgrade Instructions	188
Upgrade and Downgrade Support Policy for Junos OS Releases	189
Basic Procedure for Upgrading to Release 19.4	189
Junos OS Release Notes for PTX Series Packet Transport Routers	191
What's New	191
What's New in 19.4R2	192
What's New in 19.4R1	192
What's Changed	199
What's Changed in 19.4R2-S4	199
What's Changed in 19.4R2	199
What's Changed in 19.4R1	200
Known Limitations	202
General Routing	202

Open Issues | 203**General Routing | 204****Multiprotocol Label Switching (MPLS) | 205****Routing Protocols | 205****Resolved Issues | 206****General Routing | 206****Infrastructure | 207****Layer 2 Ethernet Services | 207****Multiprotocol Label Switching (MPLS) | 207****Routing Protocols | 207****Documentation Updates | 208****Feature Guides Are Renamed as User Guides | 208****Migration, Upgrade, and Downgrade Instructions | 209****Basic Procedure for Upgrading to Release 19.4 | 209****Upgrade and Downgrade Support Policy for Junos OS Releases | 212****Upgrading a Router with Redundant Routing Engines | 212****Junos OS Release Notes for the QFX Series | 213****What's New | 214****What's New in Release 19.4R2 | 214****What's New in Release 19.4R1 | 214****What's Changed | 226****What's Changed in 19.4R2-S4 | 226****What's Changed in 19.4R2 | 226****What's Changed in 19.4R1 | 228****Known Limitations | 230****Layer 2 Ethernet Services | 230****Layer 2 Features | 230****Network Management and Monitoring | 230****Platform and Infrastructure | 231****Routing Protocols | 231****Open Issues | 232****Class of Service (CoS) | 232****EVPN | 232****Interfaces and Chassis | 232**

Junos Fusion for Provider Edge	233
Layer 2 Features	233
Platform and Infrastructure	233
Routing Protocols	235
User Interface and Configuration	236
Virtual Chassis	236
Resolved Issues	236
Resolved Issues: 19.4R2	237
Resolved Issues: 19.4R1	242
Documentation Updates	249
Feature Guides Are Renamed As User Guides	249
Migration, Upgrade, and Downgrade Instructions	250
Upgrading Software on QFX Series Switches	250
Installing the Software on QFX10002-60C Switches	253
Installing the Software on QFX10002 Switches	253
Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches	254
Installing the Software on QFX10008 and QFX10016 Switches	256
Performing a Unified ISSU	260
Preparing the Switch for Software Installation	261
Upgrading the Software Using Unified ISSU	261
Upgrade and Downgrade Support Policy for Junos OS Releases	263
Junos OS Release Notes for SRX Series	264
What's New	265
Release 19.4R2 New and Changed Features	265
Release 19.4R1 New and Changed Features	265
What's Changed	273
What's Changed in Release 19.4R2-S4	274
What's Changed in Release 19.4R2	274
What's Changed in Release 19.4R1	275
Known Limitations	279
Application Layer Gateways (ALGs)	279
DHCP	279
Flow-Based and Packet-Based Processing	279

J-Web	280
Platform and Infrastructure	281
Routing Policy and Firewall Filters	281
Switching	281
VPNs	281
Open Issues	282
Intrusion Detection and Prevention (IDP)	283
J-Web	283
Routing Policy and Firewall Filters	283
VPNs	283
Resolved Issues	284
Resolved Issues: 19.4R2	284
Resolved Issues: 19.4R1	288
Documentation Updates	295
Feature Guides Are Renamed As User Guides	296
Migration, Upgrade, and Downgrade Instructions	296
Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases	296
Upgrading Using ISSU	297
Licensing	298
Compliance Advisor	298
Finding More Information	298
Documentation Feedback	299
Requesting Technical Support	300
Self-Help Online Tools and Resources	300
Creating a Service Request with JTAC	301
Revision History	301

Introduction

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

These release notes accompany Junos OS Release 19.4R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, JRR Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

- **Feature Guides Are Renamed As User Guides**—Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the BGP Feature Guide is now the BGP User Guide. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this link.
- **New In Focus Guide**—Starting on Junos Release 19.4R2, we are introducing a new document called In Focus that provides details on the most important features for the release in one place. We hope this document will quickly get you to the latest information about Junos features. Let us know if you find this information useful by sending email to techpubs-comments@juniper.net.
- **Important Information:**
 - [Upgrading Using ISSU on page 297](#)
 - [Licensing on page 298](#)
 - [Compliance Advisor on page 298](#)
 - [Finding More Information on page 298](#)
 - [Documentation Feedback on page 299](#)
 - [Requesting Technical Support on page 300](#)

New Features in 19.4R2

Feature	Release Note Section
Support for segment routing (ACX6360)	"What's New" on page 12

Feature	Release Note Section
Support for DHCPv6-PD on EVPN IRB synchronization between multiple PE routers (ACX 5448, ACX5448-M, and ACX5448-D)	“What's New” on page 12
Support for DHCP Option 82 over EVPN (ACX5448, ACX5448-M, and ACX5448-D)	“What's New” on page 12
Support for persistent storage of DHCPv4 and DHCPv6 bindings over EVPN-IRB (ACX5448, ACX5448-M, and ACX5448-D)	“What's New” on page 12
Multicast with IGMP or MLD snooping across VLANs for EVPN-MPLS (ACX5448)	“What's New” on page 12
Multicast with IGMP or MLD snooping within VLANs for EVPN-MPLS (ACX5448, ACX5448-M, and ACX5448-D)	“What's New” on page 12
Support for IRB and IRBv6 on EVPN-MPLS networks (ACX5448)	“What's New” on page 12
Support for EVPN E-Tree service (ACX5448)	“What's New” on page 12
L2 CoS classifiers and rewrite feature on MX Series routers (MX240, MX480, MX960) with the MPC10 line card	“What's New” on page 87
EVPN on MPLS-over-UDP tunnels (MX series and vMX)	“What's New” on page 87
Transparent forwarding of CFM packets over VPLS (MX Series)	“What's New” on page 87
Delegate segment routing LSPs to a PCE (MX Series)	“What's New” on page 87
RADIUS-sourced connection status updates to CPE devices (MX Series)	“What's New” on page 87
VMware NSX Data Center for vSphere 6.4.5 and 6.4.6 certification (QFX5100 Virtual Chassis)	“What's New” on page 214

Junos OS Release Notes for ACX Series

IN THIS SECTION

- What's New | 12
- What's Changed | 20
- Known Limitations | 22
- Open Issues | 24
- Resolved Issues | 26
- Documentation Updates | 30
- Migration, Upgrade, and Downgrade Instructions | 31

These release notes accompany Junos OS Release 19.4R2 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- What's New in Release 19.4R2-S2 | 13
- What's New in Release 19.4R2 | 13
- What's New in Release 19.4R1 | 16

Learn about new features introduced in the Junos OS main and maintenance releases for ACX Series routers.

What's New in Release 19.4R2-S2

Routing Protocols

- **Support for segment routing (ACX6360)**—Starting in Junos OS Release 19.4R2-S2, ACX6360 routers support:
 - IPv4 OSPF segment routing enabled through MPLS labels.
 - Configurable segment routing global block (SRGB) for source packet routing in networking (SPRING) protocols (OSPF).
 - Support for SPRING anycast and prefix segments and inter area (OSPF).
 - SPRING topology independent loop-free alternate (TI-LFA) procedures for link and node protection.
 - Unnumbered interface support for OSPFv2 with topology-independent loop-free alternate.

[See [Understanding Source Packet Routing in Networking \(SPRING\)](#), [Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING](#), [Topology-Independent Loop-Free Alternate with Segment Routing for OSPF](#).]

What's New in Release 19.4R2

Dynamic Host Configuration Protocol (DHCP)

- **Support for DHCPv6-PD on EVPN IRB synchronization between multiple PE routers (ACX 5448, ACX5448-M and ACX5448-D)**—Starting in Junos OS Release 19.4R2, you can use DHCPv6 prefix delegation (DHCPv6-PD) to automate the delegation of IPv6 prefixes to a requesting router on EVPN IRB. DHCPv6 prefix delegation is configured on EVPN IRB and provides IPv6 prefixes to the requesting clients, instead of the unique address. The DHCPv6-PD server acts as a provider edge (PE) router that provides the delegates through the relay (PE router) operating in the EVPN instance.

If one PE router fails, appropriate DHCPv6-PD state is made available for the remaining PE routers participating in the DHCP-PD process for the given VLAN. This is done using automatic synchronization of DHCPv6-PD states between multiple PE routers connected to the same Ethernet Segment Identifier (ESI) through EVPN BGP messages.

- **Support for DHCP Option 82 over EVPN (ACX5448, ACX5448-M and ACX5448-D)**—Starting in Junos OS Release 19.4R2, we have inserted Option 82 flags in the DHCP packets for enhanced security when the packet is sent to the server. The provider edge router that is part of the EVPN instance acts as the relay agent and inserts these flags to the DHCP packets.

DHCPv4 and DHCPv6 relay of packets uses this process. With the introduction of EVPN IRB, the relay agent uses the IRB interface with EVPN for forwarding the requests and replies to and from the client or the server instead of using the default routing.

- **Support for persistent storage of DHCPv4 and DHCPv6 bindings over EVPN IRB (ACX5448, ACX5448-M and ACX5448-D)**—Starting in Junos OS Release 19.4R2, ACX5448, ACX5448-M and ACX5448-D routers, when configured to function as a DHCP relay agent, can also be configured to preserve the

DHCPv4 and DHCPv6 subscriber bindings across reboots. Existing bindings are written to a local file in `/var/preserve`. After reboot, the binding table is populated with the contents of the file and the router identifies each subscriber that was on the deleted interface, and resumes normal packet processing for subscribers when the interface is restored. To preserve the subscriber binding information, enable the **persistent-storage** statement under the `[edit system services dhcp-local-server]` hierarchy.

[See [Preserving Subscriber Binding Information](#) and [DHCPv6 Relay Agent Overview](#).]

EVPN

- **Multicast with IGMP or MLD snooping across VLANs for EVPN-MPLS (ACX5448)**—Starting in Junos OS Release 19.4R2, ACX5448 routers support inter-VLAN multicast forwarding with IGMP or MLD snooping in an EVPN over MPLS network as follows:
 - Multicast sources are external to the EVPN instance in a Layer 3 Protocol Independent Multicast (PIM) domain.
 - All provider edge (PE) devices in the EVPN instance connect to a PIM gateway router in the PIM domain using Layer 3 interfaces, and send PIM join messages toward the PIM rendezvous point (RP) for their receivers that want to join a multicast group.
 - The PIM gateway router forwards traffic for the multicast group from the external sources to all PE devices in the EVPN instance.
 - On each PE device, you configure IRB interfaces with PIM in distributed designated router (DDR) mode. The PE device uses these IRB interfaces only to forward or route the multicast traffic on one or more VLANs to interested receivers locally (on its own device), rather than to send traffic out to other PE devices in the EVPN-MPLS network.
 - You can configure IGMP or MLD snooping in one or more routing instances of type **evpn**.
 - You must configure all VLANs on all PE devices.
 - Receivers in an EVPN instance can be single-homed to one PE device or multihomed to multiple PE devices in all-active mode. Peer PE devices for an EVPN segment synchronize IGMP or MLD state information by using BGP EVPN Type 7 and Type 8 (Join Sync and Leave Sync Route) messages.
 - The PE devices support IGMPv2 and IGMP snooping with any-source multicast (ASM) `[(*,G)]`.
 - The PE devices support MLDv1, MLDv2, and MLD snooping as follows:
 - By default when you enable MLD snooping, the PE devices process MLDv1 and MLDv2 reports only in ASM `(*,G)` mode. They discard MLDv2 source-specific multicast (SSM) `(S,G)` reports.
 - If you enable MLD snooping with the SSM-only processing option (**set evpn-ssm-reports-only** at the `[edit routing-instance instance-name protocols mld-snooping]` hierarchy level), the PE devices process MLDv2 reports as SSM `(S,G)` only. They discard MLDv1 or MLDv2 ASM `(*,G)` reports. We also recommend that you configure MLDv2 on all IRB interfaces used for inter-VLAN routing in this case.

PE devices can't process both ASM and SSM reports together.

[See [Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment](#).]

- **Multicast with IGMP or MLD snooping within VLANs for EVPN-MPLS (ACX5448, ACX5448-M, and ACX5448-D)**—Starting in Junos OS Release 19.4R2, ACX5448 routers support intra-VLAN multicast forwarding with IGMP or MLD snooping in an EVPN over MPLS network as follows:
 - The multicast sources and receivers must be within an EVPN instance. You can configure IGMP or MLD snooping for one or more routing instances of type **evpn**.
 - Receivers can be single-homed to one provider edge (PE) device or multihomed to multiple PE devices (in all-active mode only).
 - The PE devices support IGMPv2 and IGMP snooping with any-source multicast (ASM).
 - The PE devices support MLDv1, MLDv2, and MLD snooping in intra-VLAN configurations as follows:
 - By default with MLD snooping enabled, the PE devices can process MLDv1 ASM (*,G) reports but discard MLDv2 reports.
 - With MLD snooping enabled, if you configure all interfaces that receive multicast traffic with MLDv2, the PE devices can process both MLDv1 and MLDv2 reports in ASM (*,G) mode, but they discard MLDv2 source-specific multicast (SSM) (S,G) reports.
 - You must configure each VLAN on all PE devices in an EVPN instance.
 - The PE devices flood multicast source traffic on a VLAN into the EVPN instance to reach all other PE devices. When they receive the traffic from the EVPN network, the PE devices use snooping information to forward the traffic only to the interested listeners in the VLAN.
 - Multihoming peer PE devices for an Ethernet segment (ES) exchange BGP EVPN Type 7 and Type 8 (Join Sync and Leave Sync Route) messages to synchronize the IGMP or MLD state information. When they receive multicast traffic from the EVPN core, the designated forwarder (DF) PE device forwards the traffic only to interested receivers based on IGMP snooping reports and Type 7 routes.

[See [Overview of Multicast Forwarding with IGMP or MLD Snooping in an EVPN-MPLS Environment](#).]

- **Support for IRB and IRBv6 on EVPN-MPLS networks (ACX5448)**—Starting with Junos OS Release 19.4R2, you can configure EVPN with IRB or IRBv6 over an EVPN-MPLS network to extend Layer 2 connectivity across data centers. Both EVPN type 2 (MAC/IP Advertisement) and EVPN type 5 (IP Prefix Advertisement) routes are supported.

[See [EVPN with IRB Solution Overview](#).]

- **Support for EVPN E-Tree service (ACX5448)**—Starting in Junos OS 19.4R2, you can configure an Ethernet VPN Ethernet Tree (E-Tree) service on ACX5448 routers.

[See [EVPN-ETREE Overview](#).]

What's New in Release 19.4R1

EVPN

- **SPRING support for EVPN (ACX5448)**—Starting in Junos OS Release 19.4R1, you can use Source Packet Routing in Networking (SPRING) as the underlay transport in EVPN on ACX5448 routers. SPRING tunnels enable routers to steer a packet through a specific set of nodes and links in the network.

To configure SPRING, use the **source-packet-routing** statement at the **[edit protocols isis]** hierarchy level.

[See [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

- **Support for EVPN features on new hardware (ACX5448)**—Starting with Release 19.4R1, Junos OS supports the following EVPN features:
 - ARP/NDP proxy and suppression with proxy MAC responses. [See [EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression](#).]
 - EVPN with segment routing (SPRING). [See [Understanding Source Packet Routing in Networking \(SPRING\)](#).]
 - EVPN E-LAN services over MPLS, including support for VLAN-based and VLAN-bundles services. [See [EVPN Overview](#) and [Overview of VLAN Services for EVPN](#).]
 - EVPN multihoming active/active. [See [EVPN Multihoming Overview](#).]
- **Support for EVPN routing policies (ACX5448, EX4600, EX4650, EX9200, MX Series, QFX Series, and vMX)**—Starting in Junos OS Release 19.4R1, Junos OS has expanded routing policy support to include the creation and application of policy filters specific to EVPN routes. You can create policies and apply policy filters to import and export EVPN routes at the routing-instance level or at the BGP level. Junos OS supports the following matching criteria for EVPN routes:
 - Route distinguisher ID
 - NLRI route type
 - EVPN Ethernet tag
 - BGP path attributes
 - Ethernet Segment Identifier
 - MAC Address on EVPN route type 2 routes
 - IP address on EVPN route type 2 and EVPN route type 5 routes
 - Extended community

[See [Routing policies for EVPN](#).]

Junos OS XML API and Scripting

- **Automation script library upgrades (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, devices running Junos OS that support the Python extensions package include upgraded Python modules. Python scripts can leverage the upgraded versions of the following modules:
 - **idna** (2.8)
 - **jinja2** (2.10.1)
 - **jnpr.junos** (Junos PyEZ) (2.2.0)
 - **lxml** (4.3.3)
 - **markupsafe** (1.1.1)
 - **ncclient** (0.6.4)
 - **packaging** (19.0)
 - **paho.mqtt** (1.4.0)
 - **pyasn1** (0.4.5)
 - **yaml** (PyYAML package) (5.1)

[See [Overview of Python Modules Available on Devices Running Junos OS.](#)]

- **Python 3 support for commit, event, op, and SNMP scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, you can use Python 3 to execute commit, event, op, and SNMP scripts on devices running Junos OS. To use Python 3, configure the **language python3** statement at the **[edit system scripts]** hierarchy level. When you configure the **language python3** statement, the device uses Python 3 to execute scripts that support this Python version and uses Python 2.7 to execute scripts that do not support Python 3 in the given release.

The Python 2.7 end-of-support date is January 1, 2020, and Python 2.7 will be EOL in 2020. The official upgrade path for Python 2.7 is to Python 3. As support for Python 3 is added to devices running Junos OS for the different types of onbox scripts, we recommend that you migrate supported script types from Python 2 to Python 3, because support for Python 2.7 might be removed from devices running Junos OS in the future.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

MPLS

- **Support for Topology Independent Loop-Free Alternate, advertising MPLS labels, and configuring SRGB for SPRING for ISIS and OSPF (ACX5448-D and ACX5448-M)**—Starting with Junos OS Release 19.4R1, ACX5448-D and ACX5448-M router supports topology independent (TI)-loop-free alternate (LFA), advertise MPLS labels (ISIS, OSPF), and segment routing global block (SRGB) for SPRING (ISIS, OSPF).

[See [Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS](#), [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

OAM

- **Support for Ethernet OAM and Metro Ethernet services over segment routing (ACX5448-D, ACX5448-M, MX Series)**—Starting with Junos OS Release 19.4R1, ACX5448-D, ACX5448-M and MX Series routers support Ethernet OAM and Metro Ethernet services over segment routing.

[See [Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING for IS-IS Protocol](#), [Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS](#), [Ethernet OAM Connectivity Fault Management](#) .]

Routing Protocols

- **Support for configurable SRGB used by SPRING in OSPF protocols (ACX5448)**— Starting in Junos OS Release 19.4R1, you can configure the segment routing global block (SRGB) range label used by segment routing. Labels from this range are used for segment routing functionality in OSPF domain.

The SRGB is a range of the label values used in the segment routing. Prior to Junos OS Release 19.4R1, you could not configure the range for the SRGB block.

Locally you can configure `srgb start-label <label-range> index-range <index-range>` command under `[edit protocols ospf source-packet-routing]` hierarchy or globally under `[edit protocols mpls label-range]` hierarchy.

Following are the SRGB precedences for OSPF protocol:

- Local SRGB
- Global SRGB
- Node-segment implementation of 256 label block
- **Unnumbered interface support for IS-IS and OSPFv2 with topology-independent loop-free alternate (ACX Series, MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, you can enable IPv4 processing on a point-to-point interface without assigning it an explicit IPv4 address. The router borrows the IPv4 address of another Ethernet or loopback interface already configured on the router and assigns it to the unnumbered interface to conserve IPv4 addresses.

To enable IPv4 processing for unnumbered interfaces include `unnumbered-address source` at the `[edit interfaces [name] unit [name] family inet]` hierarchy level.

[See [Configuring an Unnumbered Interface](#).]

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:
 - Additional information is now logged for MGD-FFP intermodule communication.
 - Commit errors that previously were only shown onscreen are now logged.

We provide a new operational command, **request debug information**, to speed up the initial information-gathering phase of debugging.

[See [request debug information](#).]

Software Defined Networking (SDN)

- **Tunnel templates for PCE-initiated segment routing LSPs (ACX Series)**—Starting in Junos OS Release 19.4R1, you can configure a tunnel template for Path Computation Element (PCE)-initiated segment routing LSPs and apply it through policy configuration. These templates enable dynamic creation of segment routing tunnels with two additional parameters – Bidirectional forwarding detection (BFD) and LDP tunneling.

With the support for tunnel configuration, the LSPs that you would configure statically can now be automatically created from the PCE, thereby providing the benefit of reduced configuration on the device.

[See [Understanding Static Segment Routing LSP in MPLS Networks](#).]

SEE ALSO

[What's Changed | 20](#)

[Known Limitations | 22](#)

[Open Issues | 24](#)

[Resolved Issues | 26](#)

[Documentation Updates | 30](#)

[Migration, Upgrade, and Downgrade Instructions | 31](#)

What's Changed

IN THIS SECTION

- [What's Changed in 19.4R2-S4 | 20](#)
- [What's Changed in 19.4R2-S2 | 20](#)
- [What's Changed in 19.4R2 | 20](#)
- [What's Changed in 19.4R1 | 21](#)

Learn about what changed in the Junos OS main and maintenance releases for ACX Series routers.

What's Changed in 19.4R2-S4

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the `cspf-backoff-time` statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

See [cspf-backoff-time](#).

What's Changed in 19.4R2-S2

General Routing

- **Support for `gigether-options` statement (ACX5048, ACX5096)**—Junos OS supports the `gigether-options` statement at the `edit interfaces interface-name` hierarchy on the ACX5048 and ACX5096 routers. Previously, support for the `gigether`-statement was deprecated. See [gigether-options](#) and

What's Changed in 19.4R2

General Routing

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:
 - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.

- ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Software Installation and Upgrade

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

What's Changed in 19.4R1

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, MX Series, PTX Series, and SRX Series)**—Starting with Junos OS Release 19.4R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

[See [commit \(System\)](#).]

- **Support for gigheter-options statement (ACX5048, ACX5096)**—Junos OS supports the gigheter-options statement at the **edit interfaces interface-name** hierarchy on the ACX5048 and ACX5096 routers. Previously, support for the gigheter-statement was deprecated. See [gigheter-options and](#)

Routing Protocols

- XML RPC equivalent included for the `show bgp output-scheduler | display xml rpc` CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—Starting in Junos OS Release 19.4R1, we have included an XML RPC equivalent for the `show bgp output-scheduler | display xml rpc` CLI command. In Junos OS releases before Release 19.4R1, the `show bgp output-scheduler | display xml rpc` CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler.](#)]

System Logging

- Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

What's New	12
Known Limitations	22
Open Issues	24
Resolved Issues	26
Documentation Updates	30
Migration, Upgrade, and Downgrade Instructions	31

Known Limitations

IN THIS SECTION

- [General Routing](#) | 23

Learn about known limitations in Junos OS Release 19.4R2 for ACX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- All PTP packets go to the best-effort queue instead of Network Control queue. [PR1361315](#)
- ACX6360 Junos telemetry interface or telemetry infrastructure does not support the interface-filtering capability. Therefore, after you enable a particular sensor for telemetry, it is turned on for all the interfaces. [PR1371996](#)
- When a timing configuration and the corresponding interface configuration are flapped for multiple times in iteration, PTP is stuck in "INITIALIZE" state where the ARP for the neighbor is not resolved. In issue state, BCM hardware block gets into inconsistency state, where the lookup is failing. [PR1410746](#)
- The port LEDs glowing during system/vmhost halt state is the expected behavior across all ACX Series platforms. Even the system LED glows during halt state. [PR1430129](#)
- With an asymmetric network connection, EX: 10G MACsec port connected to a 10GbE channelized port, high and asymmetric T1 and T4 time errors introduce a high two-way time error. This introduces different CF updates in forward and reverse paths. [PR1440140](#)
- With the MACsec feature enabled and introduction of traffic, the peak-to-peak value varies with the percentage of traffic introduced. Find the maximum and mean values of the Time errors with different traffic rates (two router scenario) who can have maximum value jumps as high as 1054ns with 95% traffic, 640ns with 90% traffic, and 137ns with no traffic. [PR1441388](#)
- ACX Series do not support directly associating a policer in a logical interface. The association needs to be achieved using filters only. In this case, **family any** filter can be configured. Also, the ACX5448 hardware ASIC does not support egress policing. Egress shaping (H-QoS at the logical-interface level) can be used instead. [PR1446376](#)
- ACX Series routers support only 900 joins of IGMPv3 users per second. [PR1448146](#)

SEE ALSO

[What's New | 12](#)

[What's Changed | 20](#)

[Open Issues | 24](#)

[Resolved Issues | 26](#)

[Documentation Updates | 30](#)

[Migration, Upgrade, and Downgrade Instructions | 31](#)

Open Issues

IN THIS SECTION

- [General Routing | 24](#)
- [Layer 2 Ethernet Services | 25](#)
- [Multicast | 26](#)
- [Routing Protocols | 26](#)

Learn about open issues in Junos OS Release 19.4R2 for ACX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- On ACX1000, ACX2000, ACX4000, ACX5048, and ACX5096 platforms, after a new child logical interface with VLAN and filter is added on an aggregated Ethernet physical interface or if the VLAN ID of a child logical interface is changed with filter, traffic over the physical interface might get filtered with that filter on the child logical interface. For example, ae-0/0/0 is a physical interface and ae-0/0/0.100 is a logical interface. [PR1407855](#)
- Layer 2 rewrite is happening on regular bridge domain and VLAN interfaces, though there are some service dependencies (VPLS in this case) in which the egress interface map table is not updated properly with Layer 2 rewrite map ID. As a result, that causes rewrite is not working. [PR1414414](#)
- CoS table error can sometimes cause traffic outages and SNMP timeouts if the optic is plugged out and inserted back. [PR1418696](#)
- On an ACX5000 platform, high CPU usage by the fxpc process might be seen under a rare condition if parity errors are detected in devices. This issue has no direct service/traffic impact. However, because CPU utilization is high during this issue, there are some side effects. For example, the issue could impact time-sensitive features such as BFD. [PR1419761](#)
- On an ACX5448, the **request system reboot** command triggers a reboot on the host (Linux) instead of just being limited to Junos OS. [PR1426486](#)
- The em2 interface configuration is causing the FPC to crash during initialization and the FPC does not come online. After deleting the em2 configuration and restarting, the router FPC comes online. [PR1429212](#)

- In an ACX5448 with Ethernet OAM scenario, if connectivity fault management (CFM) sessions are based on 40GbE or 100GbE ports, CCM (continuity check message) packets might be dropped seen after scaling up beyond 650 CFM sessions. [PR1432720](#)
- Protocols get forwarded when using non-existing SSM map source address in IGMPv3 instead of pruning. This is a day 1 design issue, which needs to be redesigned. [PR1435648](#)
- Timing on 1G, performance is not at par compared with 10G, compensation is done to bring the mean value under class-A but the peak-to-peak variations are high and can go beyond 100ns. It has a latency variation with peak to peak variations of around 125ns–250ns (that is, 5–10% of the mean latency introduced by the each phy which is of around 2.5us) without any traffic. [PR1437175](#)
- Memory leaks are expected in this release. [PR1438358](#)
- Recovery of Junos volume isn't possible from the OAM menu. [PR1446512](#)
- ACX Series routers support only 900 joins of IGMPv3 users per second. [PR1448146](#)
- Drop profile max threshold might not be reached to its limit when the packet size is other than 1000 bytes. This is due to current design limitation. [PR1448418](#)
- It is not possible to form 125,000 IGMP groups with the ACX5448 router receiving 125,000 IGMP v2 reports per second. This is a product limitation from BCM and CPU host path queuing model. [PR1454465](#)
- In some cases when there is a failure of the I2C daemon, this might result in riofeb crash on ACX5448. [PR1455928](#)
- The CFM remote MEP is not coming up after configuration or remains in Start state. [PR1460555](#)
- When both MSTP and ERP are enabled on the same interface, ERP does not come up properly. [PR1473610](#)
- Issue seen during unified ISSU to Junos OS Release 20.1. - The upgrade, but the forwarding plane (PFE) does not function. As a result, forwarding is affected. [PR1483959](#)
- Port mirror in ACX6360-OR will not be supported in Junos OS Release 19.4R2 release. Hence configuration firewall, then port-mirror will not work. [PR1491789](#)
- This issue is not consistently seen, for every 7 iterations it fails couple of times. The issue is seen only in Non-IRB case and specific to MLD. IGMP works fine in this case. For Junos OS Release 19.4R2, the issue shall be mitigated as this scenario is not the specific customer ask which involves North bound EVPN Neighbour. [PR1505849](#)

Layer 2 Ethernet Services

- With the addition of mandatory sub-option 12(RELAY-ID), the available number bytes for circuit Id is reduced. [PR1502926](#)
- After performing feb restart, jdhcpd core files were observed on another Provider Edge device in multi home topology. [PR1504365](#)

Multicast

- When the DF interface is flapped, the traffic is doubled. As the current non-DF also forwards the traffic in EVPN multihoming scenario. As a workaround, clear the multicast route. [PR1502411](#)
- Observed **core-ripsaw-node-aftd-expr** on ACX6360-OR. Unable to back trace the core. [PR1504717](#)
- Connectivity with PIC is affected, and the SFW session stats are not fetched. [PR1505089](#)
- mpls_lsp check is failing while verifying basic **lsp_retry_limit**. After the fix for PR-1487532, we reset the **src_address** of the lsp to 0 (if src_address is not configured) whenever it changes its state from Up to down. So when the ingress-lsp goes to down state, we are resetting it to 0. The script is failing because the script is checking for src_address to be present for the ingress-lsp session. Script change is required. [PR1505474](#)

Routing Protocols

- Group limit to restrict the groups learning is not the customer req for 19.4R2. It is a generic issue and not specific to EVPN and not specific to customer POC. [PR1501485](#)

SEE ALSO

What's New	 12
What's Changed	 20
Known Limitations	 22
Resolved Issues	 26
Documentation Updates	 30
Migration, Upgrade, and Downgrade Instructions	 31

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.4R2](#) | [27](#)
- [Resolved Issues: 19.4R1](#) | [28](#)

Learn which issues were resolved in the Junos OS main and maintenance releases for ACX Series routers.

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

Resolved Issues: 19.4R2

General Routing

- On ACX5000, **MacDrainTimeOut** and **bcm_port_update** failed: Internal error. [PR1284590](#)
- On ACX5048 and ACX5096 platforms, high CPU for fxpc processes might be observed on class-of-service configuration changes on interfaces. [PR1407098](#)
- Drift messages in ACX2200, which is a PTP hybrid (PTP + Synchronous Ethernet) device. [PR1426910](#)
- ACX5448-D interfaces support: The input bytes value in the **show interfaces extensive** is not in par with older ACX or MX. [PR1430108](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- Fans on an ACX5448-M might not be running at the correct speed. [PR1448884](#)
- ACX5048 SNMP polling is stalled after the link is flapped or the SFP transceiver is replaced, and **ACX_COS_HALP(acx_cos_gport_sched_set_strict_priority:987): Failed to detach logs** is seen. [PR1455722](#)
- ACX5448-D and ACX5448-M devices do not display airflow information and temperature sensors as expected. [PR1456593](#)
- [interface] [generic] [ACX] not able to add more than 16 links in a LAG. [PR1463253](#)
- RED drop on interface is seen even when there is no congestion. [PR1470619](#)
- CoS: Egress queue statistics are not applicable to ae interfaces on ACX5048. [PR1472467](#)
- dcpfe core files are generated when you disable and then enable MACsec through Toby scripts. [PR1479710](#)
- ACX5448 Layer2 VPN with interface **ethernet-ccc input-vlan-map/output-vlan-map** might cause traffic to be silently dropped. [PR1485444](#)
- LSP (primary/standby) does not Act/Up after routing or rpd restart. [PR1494210](#)

Interfaces and Chassis

- MC-AE interface might display unknown status if the subinterface is added as part of the VLAN on the peer MC-AE node. [PR1479012](#)

Layer 2 Ethernet Services

- Member links state might be asynchronized on a connection between PE and CE devices in an EVPN A/A scenario. [PR1463791](#)

MPLS

- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. [PR1493431](#)

Resolved Issues: 19.4R1

General Routing

- On ACX5000 **MacDrainTimeOut** and **bcm_port_update failed: Internal error** error is seen. [PR1284590](#)
- bcmDPC task is high even though Interuppt START_BY_START flag is set to 0. [PR1329656](#)
- The AE interface with LACP stays down after the router reboots if link-speed is configured. [PR1357012](#)
- On ACX Series devices, the LED on the GE interface goes down when speed 10M is added. [PR1385855](#)
- Link Fault Signaling (LFS) doesn't work on ACX5448 10-, 40-, and 100-Gigabit Ethernet interfaces. [PR1401718](#)
- The optic comes with Tx enabled by default. Because the port is administratively disabled, the port is stopped. However, because the port has not been started, it does not disable Tx. [PR1411015](#)
- The ACX5448:40G FEC on ACX5448, which is FEC enabled by default, must be aligned with the MX and QFX platforms, where FEC is NONE. [PR1414649](#)
- On the ACX5448-X:SKU and ACX5448-D, 96000 ARPs get populated. However, only 47000 NH entries are present. Around 50 percent of packet drop is observed. [PR1426734](#)
- Chassisd might crash with unsupported hcos configuration when an MX104 is used as a fusion aggregation device. [PR1430076](#)
- The l2cpd process might crash and generate a core file when interfaces flaps. [PR1431355](#)
- Deviation in traffic rate in the queue is around 8 % to 10% percent in some cases. [PR1436297](#)
- 1PPS performance metrics (class A) of G.8273.2 are not met for 1G interfaces because of the variable latency added by the Vitesse PHY. [PR1439231](#)
- The interface on ACX1100 devices remains down when using SFP-1FE-FX (740-021487). [PR1439384](#)
- Transit DHCP packets are not punted to CPU and are transparently passthrough. [PR1439518](#)

- When the interface is flapped between channelized configurations, 25-Gigabit Ethernet to 100-Gigabit the aggregated Ethernet interface configuration is not cleaned up properly. [PR1441374](#)
- In an ACX5448 platforms, when the PFE failed to allocate packet buffer, portion of packet memories may not be freed. [PR1442901](#)
- RED drops might be seen after link flaps or CoS configuration changes. [PR1443466](#)
- ACX5448/18.3R1-S4.1 is not performing proper dot1p CoS rewrite on interfaces configured with l2circuit/local-switching/family ccc. [PR1445979](#)
- On ACX Series, the auto exported route between VRFs might not respond for icmp echo requests. [PR1446043](#)
- l2circuit with a **backup-neighbor** (hot-standby) configured might stop forwarding traffic after failovers. [PR1449681](#)
- oper-state for et interface does not transition from 'init' to 'Normal'. [PR1449937](#)
- RMPC core files are found after configuration changes done on the network for PTP/Clock Synchronization. [PR1451950](#)
- After disabling 100G and 40G interface Laser output power in **show interfaces diagnostics optics** shows some values. [PR1452323](#)
- ACX5448 FPC crashed due to segmentation fault. [PR1453766](#)
- Incorrect operating state is displayed in snmp trap for fan removal. [PR1455577](#)
- Enable gigheter option to configure Ethernet FEC on client ports. [PR1456293](#)
- ACX5448-D and ACX5448-M Devices does not display airflow information and temperature sensors as expected. [PR1456593](#)
- ACX5448 Layer2 VPN with encapsulation-type ethernet stops passing traffic after a random port is added with vlan configuration. [PR1456624](#)
- The rpd crash might be seen if BGP route is resolved over same prefix protocol next-hop in inet.3 table which has both RSVP and LDP routes. [PR1458595](#)
- Route resolve resolution is not happening when the packet size is 10000. [PR1458744](#)
- The traffic might be blackholed during link recovery in an open ethernet access ring with ERPS configured. [PR1459446](#)
- ACX5000: SNMP mib walk for jnxOperatingTemp not returning anything for FPC in new versions. [PR1460391](#)
- ACX5448-M Interfaces and Optics support: on enabling local loopback 10G interface is going down. [PR1460715](#)
- ACX5448-D Interfaces and Optics support: sometimes during the bring up of AE interface there are ARP resolution issues. [PR1461485](#)

- ACX Series routers LLDP neighbor not up on lag after software upgrade to Junos OS Release 18.2R3-S1. [PR1461831](#)
- RED drop on interface, no congestion. [PR1470619](#)

Layer 2 Ethernet Services

- DHCP request might get dropped in DHCP relay scenario. [PR1435039](#)

Platform and Infrastructure

- REST API process will get non-responsive when a number of request coming with a high rate. [PR1449987](#)

Routing Protocols

- Loopback address are exported into other VRF instance might not work on EX/QFX/ACX platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even the LDP neighbor's MAC changes. [PR1451217](#)
- The rpd might crash continuously due to memory corruption in an IS-IS setup. [PR1455432](#)

SEE ALSO

What's New 12
What's Changed 20
Known Limitations 22
Open Issues 24
Documentation Updates 30
Migration, Upgrade, and Downgrade Instructions 31

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 31](#)

This section lists the errata and changes in Junos OS Release 19.4R2 for the ACX Series documentation.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

[What's New | 12](#)[What's Changed | 20](#)[Known Limitations | 22](#)[Open Issues | 24](#)[Resolved Issues | 26](#)[Migration, Upgrade, and Downgrade Instructions | 31](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 31](#)

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Router. Upgrading or downgrading Junos OS might take several minutes, 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

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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 information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

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

SEE ALSO

[What's New | 12](#)

[What's Changed | 20](#)

[Known Limitations | 22](#)

[Open Issues | 24](#)

[Resolved Issues | 26](#)

[Documentation Updates | 30](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- [What's New | 33](#)
- [What's Changed | 43](#)
- [Known Limitations | 47](#)
- [Open Issues | 47](#)
- [Resolved Issues | 51](#)
- [Documentation Updates | 55](#)
- [Migration, Upgrade, and Downgrade Instructions | 56](#)

These release notes accompany Junos OS Release 19.4R2 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- [What's New in 19.4R2 | 33](#)
- [What's New in 19.4R1 | 34](#)

Learn about new features introduced in the Junos OS main and maintenance releases for EX Series switches.

NOTE: The following EX Series switches are supported in Release 19.4R2: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in 19.4R2

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.4R2.

What's New in 19.4R1

Authentication, Authorization, and Accounting

- **Disable LLDP TLV messages (EX4300-48MP switches)**—Starting in Junos OS Release 19.4R1, you can disable nonmandatory time, length, and value (TLV) messages so they will not be advertised by the Link Layer Discovery Protocol (LLDP) or Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED).

[See [Device Discovery Using LLDP and LLDP-MED on Switches](#).]

Class of Service

- **Support for 802.1p rewrite of host outbound traffic (EX4300-MP)**—Starting in Junos OS Release 19.4R1, support is provided for 802.1p rewrite of host outbound traffic on EX4300-MP devices.

[See [Applying Egress Interface Rewrite Rules to the IEEE 802.1p Field for All Host Outbound Traffic on the Interface](#).]

EVPN

- **Support for EVPN routing policies (ACX5448, EX4600, EX4650, EX9200, MX Series, QFX Series, and vMX)**—Starting in Junos OS Release 19.4R1, Junos OS has expanded routing policy support to include the creation and application of policy filters specific to EVPN routes. You can create policies and apply policy filters to import and export EVPN routes at the routing-instance level or at the BGP level. Junos OS supports the following matching criteria for EVPN routes:

- Route distinguisher ID
- NLRI route type
- EVPN Ethernet tag
- BGP path attributes
- Ethernet segment identifier
- MAC address on EVPN Type 2 routes
- IP address on EVPN Type 2 and EVPN Type 5 routes
- Extended community

[See [Routing policies for EVPN](#).]

- **Access security support in EVPN-VXLAN overlay networks (EX4300-48MP)**—Starting in Junos OS Release 19.4R1, we support access security features on EX4300-48MP switches that function as Layer 2 VXLAN gateways in an EVPN-VXLAN centrally-routed overlay network (two-layer IP fabric). We support the following features on Layer 2 server-facing interfaces that are associated with VXLAN-mapped VLANs:

- DHCPv4 and DHCPv6 snooping. [See [DHCP Snooping](#).]
- Dynamic ARP inspection (DAI). [See [Understanding and Using Dynamic ARP Inspection \(DAI\)](#).]

- Neighbor discovery inspection (NDI). [See [IPv6 Neighbor Discovery Inspection](#).]
- IPv4 and IPv6 source guard. [See [Understanding IP Source Guard for Port Security on Switches](#).]
- Router advertisement (RA) guard. [See [Understanding IPv6 Router Advertisement Guard](#).]

The access security features function the same and you configure them in the same way in an EVPN-VXLAN environment as you do in a non-EVPN-VXLAN environment. However, keep these differences in mind:

- We do not support these features on multihomed servers.
- These features do not influence the VXLAN tunneling and encapsulation process.
- **Layer 3 VXLAN gateway support in EVPN-VXLAN overlay network (EX4300-48MP)**—Starting in Junos OS Release 19.4R1, the EX4300-48MP switch can function as a Layer 3 VXLAN gateway in an EVPN-VXLAN centrally-routed bridging overlay (two-layer IP fabric) and an edge-routed bridging overlay (collapsed IP fabric). As a Layer 3 VXLAN gateway, the switch supports these features:
 - Default gateway function through the configuration of an IRB interface. [See [Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network](#).]
 - Routing of IPv6 data traffic through an EVPN-VXLAN overlay network with an IPv4 underlay. [See [Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay](#).]
 - EVPN pure Type 5 routes. [See [Understanding EVPN Pure Type-5 Routes](#).]
- **Features supported on EX4650 and QFX5120 switches**—Starting with Junos OS Release 19.4R1, the following Junos OS features are supported on EX4650 and QFX5120 switches:
 - Automatically generated Ethernet segment identifiers (ESIs) in EVPN-VXLAN and EVPN-MPLS networks. [See [Understanding Automatically Generated and Assigned ESIs in EVPN Networks](#).]
 - Firewall filtering and policing on EVPN-VXLAN traffic. [See [Understanding VXLANs](#) and [Overview of Firewall Filters](#).]
 - Graceful restart on EVPN-VXLAN. [See [Graceful Restart in EVPN](#).]
 - IGMPv2 snooping for EVPN-VXLAN in a multihomed environment. [See [Overview of IGMP Snooping in an EVPN-VXLAN Environment](#).]
 - IPv6 data traffic support through an EVPN-VXLAN overlay network. [See [Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay](#).]
 - Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface. [See [Understanding Flexible Ethernet Services Support with EVPN-VXLAN](#).]
 - MAC limiting, storm control, and port mirroring support in EVPN-VXLAN overlay networks. [See [MAC Limiting, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment](#).]

- Multihomed proxy advertisement.
[See [EVPN Multihoming Implementation](#).]
- Selective multicast forwarding and SMET route support in EVPN-VXLAN.
[See [Overview of Selective Multicast Forwarding](#).]
- Standard class-of-service (CoS) features—classifiers, rewrite rules, and schedulers—are supported on VXLAN interfaces.
[See [Understanding CoS on OVSDB-Managed VXLAN Interfaces](#).]
- VMTO for ingress traffic.
[See [Ingress Virtual Machine Traffic Optimization](#).]
- **Exclusion list with MAC pinning in an EVPN network (EX9200 and MX Series)**—When you enable **mac-pinning** on an interface, all MAC addresses that are learned on that interface will be pinned and cannot be relearned on the other interfaces in the EVPN network. Starting in Junos OS Release 19.4R1, you can create a list of MAC addresses that would be excluded from being pinned and the MAC address can be moved and relearned on another interface within the EVPN network. While MAC pinning is configured on the interface, the exclusion list is configured for the device. To create an exclusion list, include a list of MAC addresses with the **exclusive-mac** parameter at the **[edit protocols l2-learning global-mac-move]** hierarchy level.
[See [Creating exclusion list for MAC Pinning](#).]

Junos OS XML, API, and Scripting

- **Python 3 support for commit, event, op, and SNMP scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, you can use Python 3 to execute commit, event, op, and SNMP scripts on devices running Junos OS. To use Python 3, configure the **language python3** statement at the **[edit system scripts]** hierarchy level. When you configure the **language python3** statement, the device uses Python 3 to execute scripts that support this Python version and uses Python 2.7 to execute scripts that do not support Python 3 in the given release.

The Python 2.7 end-of-support date is January 1, 2020, and Python 2.7 will be EOL in 2020. The official upgrade path for Python 2.7 is to Python 3. As support for Python 3 is added to devices running Junos OS for the different types of onbox scripts, we recommend that you migrate supported script types from Python 2 to Python 3, because support for Python 2.7 might be removed from devices running Junos OS in the future.

[See [Understanding Python Automation Scripts for Devices Running Junos OS](#).]

- **Automation script library upgrades (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, devices running Junos OS that support the Python extensions package include upgraded Python modules. Python scripts can leverage the upgraded versions of the following modules:

- **idna** (2.8)
- **jinja2** (2.10.1)
- **jnpr.junos** (Junos PyEZ) (2.2.0)
- **lxml** (4.3.3)
- **markupsafe** (1.1.1)
- **ncclient** (0.6.4)
- **packaging** (19.0)
- **paho.mqtt** (1.4.0)
- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

[See [Overview of Python Modules Available on Devices Running Junos OS.](#)]

Junos Telemetry Interface

- **JTI Packet Forwarding Engine and Routing Engine sensor support (EX4300-MP switches)**—Starting in Junos OS Release 19.4R1, you can use the Junos Telemetry Interface (JTI) and remote procedure calls (gRPC) to stream statistics from EX4300-MP switches to an outside collector.

The following Routing Engine statistics are supported:

- LACP state export
- Chassis environmentals export
- Network discovery chassis and components
- LLDP export and LLDP model
- BGP peer information (RPD)
- RPD task memory utilization export
- Network discovery ARP table state
- Network discovery NDP table state

The following Packet Forwarding Engine statistics are supported:

- Congestion and latency monitoring
- Logical interface
- Filter
- Physical interface

- NPU/LC memory
- Network discovery NDP table state

To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

[See [Configuring a Junos Telemetry Interface Sensor \(CLI Procedure\)](#), [Configure a Telemetry Sensor in Junos](#) and [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **JTI and OpenConfig support for VLAN sensors (EX4650, QFX5120)**—Junos OS Release 19.4R1 supports the export of VLAN statistics using either Junos Telemetry Interface (JTI) services or remote procedure call (gRPC) services. You can export statistics at configurable intervals to an outside collector.

This feature includes OpenConfig support for the data model **openconfig-vlan.yang** for VLAN configuration version 1.0.2.

Use the following resource paths in a gRPC or gNMI subscription:

- **/vlans/**
- **/vlans/vlan/state/name**
- **/vlans/vlan/state/vlan-id**
- **/vlans/vlan/state/status**
- **/vlans/vlan/members/**
- **/vlans/vlan/members/member/interface-ref/state/interface/**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/interface-mode**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/native-vlan**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/access-vlan**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/trunk-vlan**
- **/vlans/vlan/members/member/interface-ref/state/interface/vlan/state/vlan-id**

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Layer 2 Features

- **Redundant trunk group support (EX4650 and QFX5120)**—Starting with Junos OS Release 19.4R1, EX4650 and QFX5120 switches support redundant trunk group (RTG) links.

[See [Redundant Trunk Groups](#).]

- **Ethernet ring protection switching (ERPS)(EX4300-MP)**—Starting in Junos OS Release 19.4R1, the EX4300-MP supports Ethernet ring protection switching (ERPS) to reliably achieve carrier-class network

requirements for Ethernet topologies forming a closed loop. The ITU-T Recommendation is G.8032 version 1.

ERPS version 1 comprises the following features:

- Revertive mode of operation of the Ethernet ring
- Multiple ring instances on the same interfaces
- Multiple ring instances on different interfaces
- Interworking with Spanning Tree Protocol, Multiple Spanning Tree Protocol, and redundant trunk groups

[See [Ethernet Ring Protection Switching Overview](#).]

- **Ethernet ring protection switching (ERPS)(EX4650 and QFX5120)**—Starting in Junos OS Release 19.4R1, the EX4650 and QFX5120 support Ethernet ring protection switching (ERPS) to reliably achieve carrier-class network requirements for Ethernet topologies forming a closed loop. The ITU-T Recommendation is G.8032 version 1.

ERPS version 1 comprises the following features:

- Revertive mode of operation of the Ethernet ring
- Multiple ring instances on the same interfaces
- Multiple ring instances on different interfaces
- Interworking with Spanning Tree Protocol, Multiple Spanning Tree Protocol, and redundant trunk groups

[See [Ethernet Ring Protection Switching Overview](#).]

MPLS

- **MPLS scaling enhancements (EX4650 and QFX5120)**—Starting in Junos OS Release 19.4R1, MPLS scaling is enhanced on EX4650 and QFX5120 switches. For instance, you can increase the scale from its default 1024 to 8192 on QFX5120 switches. This enhancement optimizes and increases the ingress tunnel scale to address the current needs of data center networks either in IP-CLOS or IP over MPLS application spaces.

[See [Supported MPLS Scaling Values](#).]

Multicast

- **Multicast VLAN registration (MVR) (EX4300-48MP switches and Virtual Chassis)**—Starting in Junos OS Release 19.4R1, EX4300 multigigabit (EX4300-48MP) switches and Virtual Chassis support multicast VLAN registration (MVR). MVR efficiently distributes IPTV multicast streams across an Ethernet ring-based Layer 2 network and reduces the bandwidth needed for this traffic. MVR uses a multicast VLAN (MVLAN) as a source VLAN associated with one or more multicast group addresses, and you designate other VLANs as MVR receiver VLANs that have listeners interested in the MVLAN traffic. The device selectively forwards the traffic from source interfaces on the MVLAN to receiver interfaces that are on the MVR receiver VLANs (but not on the MVLAN).

[See [Understanding Multicast VLAN Registration](#).]

Operation, Administration, and Maintenance (OAM)

- **Ethernet CFM support (EX4300-MP switches)**—Starting with Junos OS Release 19.4R1, the EX4300-MP switch supports Ethernet connectivity fault management (CFM). You can use Ethernet CFM to:
 - Monitor faults, using the continuity check messages (CCM) protocol to discover and maintain adjacencies at the VLAN or link level.
 - Discover paths and verify faults, using the linktrace protocol to map the path taken to a destination MAC address.
 - Isolate faults, using loopback messages, and troubleshoot.

You configure Ethernet CFM using the **set protocols oam ethernet connectivity-fault-management** command, and verify the configuration using the **show oam ethernet connectivity-fault-management** command.

- **Support for Ethernet CFM (EX4650)**—Starting with Junos OS Release 19.4R1, the EX4650 switch supports Ethernet connectivity fault management (CFM). You can use Ethernet CFM to:
 - Monitor faults, using the continuity check messages (CCMs) to discover and maintain adjacencies at the VLAN or link level.
 - Discover paths and verify faults, using the Link Trace protocol to map the path taken to a destination MAC address.
 - Isolate and troubleshoot faults, using loopback messages. .

NOTE: Only down maintenance association end points (MEPs) are supported in CFM.

You configure Ethernet CFM using the **set protocols oam ethernet connectivity-fault-management** command, and verify the configuration using the **show oam ethernet connectivity-fault-management** command.

[See [Understanding Ethernet OAM Connectivity Fault Management for Switches](#).]

Support for LFM (EX4650)—Starting with Junos OS Release 19.4R1, the EX4650 switch supports OAM link fault management (LFM). You can configure OAM LFM on point-to-point Ethernet links that are connected directly or through Ethernet repeaters, and on aggregated Ethernet interfaces. The LFM status of individual links determines the LFM status of the aggregated Ethernet interface. The EX4650 supports the following OAM LFM features:

- Discovery and link monitoring
- Remote fault detection
- Remote loopback

[See [IEEE 802.3ah OAM Link-Fault Management Overview](#).]

Port Security

- **Stateless address autoconfiguration (SLAAC) snooping (EX4300-48MP)**—Starting in Junos OS Release 19.4R1, the EX4300-48MP switch supports Stateless address auto configuration (SLAAC) snooping. The switch validates IPv6 clients that use SLAAC for dynamic address assignment against the SLAAC snooping binding table before allowing the clients access to the network.

[See [IPv6 Stateless Address Auto-configuration \(SLAAC\) Snooping](#).]

- **Untrusted mode on trunk interfaces for DHCP security (EX4300-48MP)**—Starting in Junos OS Release 19.4R1, you can configure a trunk interface as untrusted for DHCP security features on EX4300-48MP switches. Trunk interfaces in untrusted mode support DHCP snooping and DHCPv6 snooping, dynamic ARP inspection (DAI), and IPv6 neighbor discovery inspection.

[See [Understanding Trusted and Untrusted Ports](#).]

- **MACsec license enforcement (EX4300-48MP)**—Starting in Junos OS Release 19.4R1, you must install a Media Access Security (MACsec) feature license if you want MACsec functionality on your EX4300-48MP switch. If the MACsec license is not installed, MACsec functionality cannot be activated. You add the MACsec license using the **request system license add** command.

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

Routing Policy and Firewall Filters

- **Firewall filter support on IPv6 egress interfaces (EX4300-48MP)**—Starting in Junos OS Release 19.4R1, you can configure a firewall filter on an IPv6 egress interface to match the specified IPv6 source or destination addresses, for example, to protect a third-party device connected to the switch.

[See [eracl-ip6-match](#) and [Configuring an Egress Filter Based on IPv6 Source or Destination IP Addresses.](#)]

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:
 - Additional information is now logged for MGD-FFP intermodule communication.
 - Commit errors that previously were only shown onscreen are now logged.

We provide a new operational command, **request debug information**, to speed up the initial information-gathering phase of debugging.

[See [request debug information.](#)]

System Management

- **Change status LED for network port to chassis beacon light (EX2300, EX2300 Virtual Chassis, EX3400, EX3400 Virtual Chassis)**—By default, when a network port and its associated link are active, the status LED for that port blinks green 8 times per second. Starting in Junos OS Release 19.4R1, you can use the **request chassis beacon** command to slow down the current blinking rate to 2 blinks per second. The slower-blinking and steadier green light acts as a beacon that leads you to an EX2300 or EX3400 switch or a particular port in a busy lab.

Using options with the **request chassis beacon** command, you can do the following for one or all network port status LEDs on a specified FPC):

- Turn on the beacon light for:
 - 5 minutes (default)
 - A specified number of minutes (1 through 120)
- Turn off the beacon light:
 - Immediately
 - After a specified number of minutes (1 through 120)

After the beacon light is turned off, the blinking rate for the network port's status LED returns to 8 blinks per second.

[See [request chassis beacon](#).]

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (EX4300-48MP, EX9251, and EX9253 switches)**—Starting in Junos OS Release 19.4R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database on EX4300-48MP, EX9251, and EX9253 switches. The ephemeral database provides a fast programmatic interface that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. The device’s active configuration is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database.

[See [Understanding the Ephemeral Configuration Database](#).]

SEE ALSO

What's Changed 43
Known Limitations 47
Open Issues 47
Resolved Issues 51
Documentation Updates 55
Migration, Upgrade, and Downgrade Instructions 56

What's Changed

IN THIS SECTION

- [What's Changed in 19.4R2-S4 | 44](#)
- [What's Changed in 19.4R2 | 44](#)
- [What's Changed in 19.4R1 | 45](#)

Learn about what changed in Junos OS main and maintenance releases for EX Series.

What's Changed in 19.4R2-S4

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the `cspf-backoff-time` statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

See [cspf-backoff-time](#).

What's Changed in 19.4R2

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (EX Series and QFX Series)**—Starting with Junos OS Release 19.4R2, the `persist-groups-inheritance` option at the `[edit system commit]` hierarchy level is enabled by default. To disable this option, use `no-persist-groups-inheritance`.

[See [commit \(System\)](#).]

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:
 - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.
 - ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Multicast

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 19.4R2, EX4600, EX4650, and the QFX5000 line of switches provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... n packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). In earlier Junos OS releases, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** do not provide valid statistics for multicast traffic at Layer 2.

[See [show multicast snooping route](#).]

Routing Protocols

- **Advertising /32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—In Junos OS Release, multiple loopback addresses export into `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier Junos OS releases, multiple secondary loopback addresses in TED were added into `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router-id.

Software Installation and Upgrade

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

What's Changed in 19.4R1

General Routing

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the **show interfaces mc-ae extensive** command. The output now displays the following two additional fields:
 - Local Partner System ID?LACP partner system ID as seen by the local node.
 - Peer Partner System ID?LACP partner system ID as seen by the MC-AE peer node.

Previously, the **show interfaces mc-ae extensive** command did not display these additional fields.

[See [show interfaces mc-ae.](#)]

Interfaces and Chassis

- **Logical Interface is created along with physical interface by default (MX Series, QFX Series, EX Series)**—Starting in Junos OS Release 19.4R1, logical interfaces are created on ge, et, and xe interfaces along with the physical interface, by default. In earlier Junos OS releases, by default, only physical interfaces are created.

For example, for ge interfaces, previously when you viewed the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

Routing Protocols

- **XML RPC equivalent included for the show bgp output-scheduler | display xml rpc CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.4R1, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler.](#)]

System Logging

- **Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

[What's New | 33](#)

[Known Limitations | 47](#)

[Open Issues | 47](#)

[Resolved Issues | 51](#)

[Documentation Updates | 55](#)

[Migration, Upgrade, and Downgrade Instructions | 56](#)

Known Limitations

There are no known limitations for the EX Series switches in Junos OS Release 19.4R2.

SEE ALSO

[What's New | 33](#)[What's Changed | 43](#)[Open Issues | 47](#)[Resolved Issues | 51](#)[Documentation Updates | 55](#)[Migration, Upgrade, and Downgrade Instructions | 56](#)

Open Issues

IN THIS SECTION

- [Authentication and Access Control | 48](#)
- [EVPN | 48](#)
- [General Routing | 48](#)
- [Infrastructure | 49](#)
- [Interfaces and Chassis | 49](#)
- [Layer 2 Features | 50](#)
- [Platform and Infrastructure | 50](#)
- [Routing Protocols | 50](#)

Learn about open issues in Junos OS Release 19.4R2 for EX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Authentication and Access Control

- In a server-fail scenario, when tagged traffic is sent for the first client, MAC learning happens for both data and voice. But for the second client on the same interface, learning happens only for voice. This is because the VLAN is already added for an interface due to first client authentication process. [PR1462479](#)
- In case of captive portal, authentication fails for captive portal user when receives the filter from the RADIUS for that user. This issue is seen only in the captive portal configuration with RADIUS filter when dot1x stanza is not present. [PR1504818](#)

EVPN

- In any EVPN scenario (for example, active-active multi-homing mode, active-standby multi-homing mode, EVPN-VXLAN, or EVPN-MPLS), the I2ald memory might slowly come up when the local CE device or core face interfaces continuously flap. If the memory of I2ald is exhausted, it causes the I2ald to crash. [PR1498023](#)

General Routing

- On an EX2300 switch, the output of the **show chassis routing-engine** command might display an incorrect value of **mac reset** for the **last reboot reason** field. [PR1331264](#)
- Scale of 150 VRRP is not tested before; there are no issues observed for 100 VRRP groups. At the higher scale, there are no drops but traffic gets flooded for groups beyond 100. [PR1371520](#)
- On an EX9208 switch, a few xe- interfaces are going down with the **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error** error message. [PR1377840](#)
- On EX Series platforms, the DHCP/PPP subscribers might fail to bind. The reason is that when you install a new software image, its shared memory (created by the previously running image) might not be cleared out. The issue persists until the previous values in the shared memory are removed and the daemons affected by the data in the shared memory continue generating core files. Thus they will not be able to function properly. [PR1396470](#)
- On EX Series platforms, unicast RPF check in strict mode might not work properly. [PR1417546](#)
- When the chassisd process receives incorrect values from LCMD for the RPM values, the fan status changes to **Failed** from **OK**, and vice versa. [PR1417839](#)
- The EX4300-48MP switch cannot learn MAC addresses through some access ports that are directly connected to a host when autonegotiation is used. [PR1430109](#)
- The time it takes to install or delete IPv4 or IPv6 routes into the FIB is slowed down. Analysis shows that rpd learning rates are not degraded, but RIB-to-FIB download rate is degraded. [PR1441737](#)
- On the EX9214 device, if the MACsec-enabled link flaps after reboot, the error **errorlib_set_error_log(): err_id(-1718026239)** is observed. [PR1448368](#)

- In overall commit time, the evaluation of mustd constraints is taking 2 seconds more than usual. This is because the **persist-group-inheritance** feature has been made as a default feature in the latest Junos OS releases. Eventually, this feature helps improve the subsequent commit times for scaled configurations significantly. The **persist-group-inheritance** feature is useful in customer scenarios where groups and nested groups are used extensively. In those scenarios, the group inheritance paths are not built every time, thus subsequent commits are faster. [PR1457939](#)
- On EX4600 platforms with MACsec configured, if traffic flows through the MACsec-enabled link, increase in framing errors or runts statistics might be seen in the **show interfaces extensive <>** command for the affected interface. Traffic loss might also happen because of this issue. [PR1469663](#)
- When both MSTP and ERP are enabled on the same interface, ERP might not come up properly. [PR1473610](#)
- On an EX9214, while verifying the last-change op-state value through XML, **rpc-reply** message is inappropriate. [PR1492449](#)
- Memory issues are seen while doing NSSU from lower version to Junos OS Release 20.1R1. Manually clean up the space and run the **request system storage cleanup** command. [PR1494963](#)
- On EX4300-48MP and EX4300 Virtual Chassis, an OSPF session on IRB over lag interface with 40G port as member, the session might stuck in exstart. This issue is Virtual Chassis specific and can be recovered by rebooting master, and does not appear again. [PR1498903](#)
- After configuring MACsec in DUT, the **show security macsec statistics** command output shows zero packets received. [PR1503404](#)
- On EX4600 platform, device is struck at "emergency:/#" prompt while loading Junos OS Release 19.4R2 using PXE image. [PR1508345](#)

Infrastructure

- When xSTP/RTG is not configured in the network and there is a traffic loop, after the network loop is broken, sometimes MAC address learning might not happen. As a workaround, restart the pfem. [PR473454](#)
- On EX Series switches, If you are configuring a large-scale number of firewall filters on some interfaces, the FPC might crash and generate core files. [PR1434927](#)
- On an EX9251 switch, **IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) failed** error message is observed continuously in AD with base configurations. [PR1485038](#)

Interfaces and Chassis

- On GRES, the VSTP port cost on aggregated Ethernet interfaces might get changed, leading to a topology change. [PR1174213](#)
- A few of the clients are not bound while verifying the DHCP binding hosts in the default routing instance though the relays are reachable. [PR1477084](#)

- When MTU is changed on a scale number of interfaces with interface range configuration on Virtual Chassis, dcd might be crashed and the configured MTU might not get applied. [PR1503496](#)

Layer 2 Features

- GARP error message **DND_ANUPAM_JUNE_10_client eswd[1203]: vlan_interface_admin_up: vif ifl flags 0xc000** is generated when there is a deletion or addition of MAC address in the FDB. [PR1192520](#)

Platform and Infrastructure

- The runt counter never incremented in the output of the **show interface x/x/x extensive** command when a runt packet with less than 64 Bytes frame size was received. The packet will be dropped as expected. [PR1419724](#)
- On the EX9208 devices, traffic loss is observed if ingress and egress ports are in different FPCs. [PR1429714](#)
- On EX9208, 33 percent degradation with MAC learning rate is observed in Junos OS Release 19.3R1 while comparing with Junos OS Release 18.4R1. [PR1450729](#)
- On EX2300 and EX3400 platforms, when doing an upgrading operation, as image size grows over a period of time and subsequently storage is insufficient to install images, the upgrade might fail with the error message **not enough space to unpack**. [PR1464808](#)

Routing Protocols

- On EX9251 platform, ECDSA256+SHA256 is not used for software integrity checking because of issue with FIPS mode and telnet to device does not work. [PR1504211](#)

SEE ALSO

[What's New | 33](#)

[What's Changed | 43](#)

[Known Limitations | 47](#)

[Resolved Issues | 51](#)

[Documentation Updates | 55](#)

[Migration, Upgrade, and Downgrade Instructions | 56](#)

Resolved Issues

IN THIS SECTION

- Class of Service (CoS) | 51
- EVPN | 52
- Forwarding and Sampling | 52
- General Routing | 52
- Infrastructure | 53
- Interfaces and Chassis | 53
- Junos Fusion for Enterprise | 53
- Junos Fusion Satellite Software | 54
- Layer 2 Features | 54
- Layer 2 Ethernet Services | 54
- MPLS | 54
- Platform and Infrastructure | 54
- Routing Protocols | 55
- User Interface and Configuration | 55
- Virtual Chassis | 55

Learn which issues were resolved in Junos OS main and maintenance releases for EX Series.

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

Class of Service (CoS)

- Shaping does not work after the reboot if **shaping-rate** is configured. [PR1432078](#)
- The traffic is placed in network-control queue on an extended port even if it comes in with a different DSCP marking. [PR1433252](#)

EVPN

- The ESI of IRB interfaces does not update after an autonomous-system number change if the interface is down. [PR1482790](#)

Forwarding and Sampling

- Type 1 ESI/AD route might not be generated locally on an EVPN PE device in all-active mode. [PR1464778](#)

General Routing

- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- IRB over VTEP unicast traffic might get dropped on EX9200 platforms. [PR1436924](#)
- The MAC pause frames will increment in the receive direction if half-duplex mode on 10-Mbps or 100-Mbps speed is configured. [PR1452209](#)
- Link-up delay and traffic drop might be seen on mixed service provider Layer 2 or Layer 3 type and enterprise style Layer 2 type configurations. [PR1456336](#)
- MAC addresses learned on an RTG might not be aged out after the aging time. [PR1461293](#)
- The RTG link is nearly 20 seconds down when the backup node is rebooting. [PR1461554](#)
- On EX Series switches with ELS, some command lines to disable MAC learning are not working. [PR1464797](#)
- The jdhcpd might consume high CPU and no further subscribers can be brought up if there are more than 4000 DHCP relay clients in the MAC move scenario. [PR1465277](#)
- The MAC move message might have an incorrect **from** interface when rapid MAC moves happen. [PR1467459](#)
- FPCs might get disconnected from the EX3400 Virtual Chassis briefly after image upgrade or reboot. [PR1467707](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- SSH session closes while you check for the **show configuration | display set** command for both local and non-local users. [PR1470695](#)
- EX3400 is advertising only 100 Mbps when configured with 100-Mbps speed with autonegotiation enabled. [PR1471931](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- CoS 802.1p bits rewrite might not happen in Q-in-Q mode. [PR1472350](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)

- On EX4300 switches, the output of the **show security macsec statistics** command shows incorrect high values. [PR1476719](#)
- The dhcpd process might crash in a Junos fusion environment. [PR1478375](#)
- Core files are generated at **cassis_alloc_list_timed_free** in **cassis_free_thread_entry**. [PR1478392](#)
- TFTP installation from the loader prompt might not succeed on the EX Series switches. [PR1480348](#)
- ARP request packets for unknown hosts might get dropped in the remote PE device in an EVPN-VXLAN scenario. [PR1480776](#)
- On EX2300, SNMP traps are not generated when the MAC addresses limit is reached. [PR1482709](#)
- DHCP binding fails while verifying DHCPv4 snooping functionality in a private VLAN with a firewall to block or allow certain IPv4 packets. [PR1490689](#)
- Traffic loss might be seen under MC-LAG scenario on EX4650. [PR1494507](#)

Infrastructure

- Continuous dcpfe error messages and eventd process hog might be seen in an EX2300 Virtual Chassis scenario. [PR1474808](#)
- Kernel core file might be generated if you deactivate daemon on EX2300 and EX3400 platforms. [PR1483644](#)

Interfaces and Chassis

- Executing commit might hang because of a stuck dcd process. [PR1470622](#)

Junos Fusion for Enterprise

- The SDPD generates core files at **vfpc_all_eports_deletion_complete** **vfpc_dampen_fpc_timer_expiry**. [PR1454335](#)
- Loop detection might not work on extended ports in a Junos fusion scenario. [PR1460209](#)

Junos Fusion Satellite Software

- Temperature sensor alarms are seen on EX4300 switches in a Junos fusion scenario. [PR1466324](#)

Layer 2 Features

- The LLDP function might fail when a Juniper Networks device connects to a non-Juniper device. [PR1462171](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)

Layer 2 Ethernet Services

- Member links state might be asynchronized on a connection between a PE device and a CE device in an EVPN A/A scenario. [PR1463791](#)
- Confirm and reply packets might not get processed correctly because of issues with DHCPv6 relay processing. [PR1496220](#)

MPLS

- BGP session might keep flapping between two directly connected BGP peers because of the incorrect usage of the TCP-MSS. [PR1493431](#)

Platform and Infrastructure

- The IRB traffic might drop after a mastership switchover. [PR1453025](#)
- The OSPF neighbor might go down when mDNS or PTP traffic is received at a rate higher than 1400 pps. [PR1459210](#)
- Traffic loss might be observed for more than 20 seconds when performing NSSU on EX4300 Virtual Chassis. [PR1461983](#)
- IGMP reports are dropped with mixed enterprise/SP configuration styles on EX4300 switches. [PR1466075](#)
- The switch might not be able to learn MAC addresses with **dot1x** and **interface-mac-limit** configured. [PR1470424](#)
- On an EX4300, the input firewall filter attached to isolated or community VLANs is not matching dot1p bits on the VLAN header. [PR1478240](#)
- The traffic destined to a VRRP VIP might be dropped after the IRB interface is disabled on the initial VRRP master. [PR1491348](#)

Routing Protocols

- BGP IPv4 or IPv6 convergence and RIB install or delete time degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The MUX state in an LACP interface does not go to **collecting and distributing** and remains **attached** after enabling the aggregated Ethernet interface. [PR1484523](#)
- FPC might go to "NotPrsnt" state after upgrading with non-tvp image in Virtual Chassis or Virtual Chassis Fabric setup. [PR1485612](#)

User Interface and Configuration

- The **umount: unmount of /.mount/var/val/chroot/packages/mnt/jweb-ex32-d2cf6f6b failed: Device busy** message is seen when Junos OS is upgraded with the **validate** option. [PR1478291](#)

Virtual Chassis

- Disabling one of the VCP ports might result in other VCP port to flap. [PR1469257](#)

SEE ALSO

What's New	 33
What's Changed	 43
Known Limitations	 47
Open Issues	 47
Documentation Updates	 55
Migration, Upgrade, and Downgrade Instructions	 56

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides](#) | 56

This section lists the errata and changes in Junos OS Release 19.4R2 for the EX Series switches documentation.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

What's New	 33
What's Changed	 43
Known Limitations	 47
Open Issues	 47
Resolved Issues	 51
Migration, Upgrade, and Downgrade Instructions	 56

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 56

This section contains the upgrade and downgrade support policy 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

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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://support.juniper.net/support/eol/software/junos/>.

SEE ALSO

What's New 33
What's Changed 43
Known Limitations 47
Open Issues 47
Resolved Issues 51
Documentation Updates 55

Junos OS Release Notes for JRR Series

IN THIS SECTION

● What's New 58
● What's Changed 59
● Known Limitations 59
● Open Issues 60
● Resolved Issues 60
● Documentation Updates 61
● Migration, Upgrade, and Downgrade Instructions 62

These release notes accompany Junos OS Release 19.4R2 for JRR Series. They describe new and changed features, limitations, and known and resolved problems in the hardware.

You can find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- [What's New in 19.4R2 Release | 58](#)
- [What's New in 19.4R1 Release | 58](#)

Learn about new features introduced in the Junos OS main release and the maintenance releases for JRR Series.

What's New in 19.4R2 Release

- There are no new features or enhancements to existing features for JRR Series in Junos OS Release 19.4R2.

What's New in 19.4R1 Release

Hardware

- **JRR200 Route Reflector**—Starting with Junos OS Release 19.4R1, JRR200 Route Reflector a 1U form factor appliance with a multicore x86 CPU and preinstalled vRR software that can host one route reflector instance is available. JRR200 is suitable for large enterprises, data centers and service providers for hosting vRR software to scale up to 30 million routing information base (RIB) entries. The JRR200 route reflector comes with eight 1/10 Gigabit Ethernet SFP+ ports, 64 GB of DDR4 memory, and two 240 GB solid-state drives (SSDs) in a RAID1 configuration. It is available in both AC and DC models which support Zero Touch Provisioning mode (ZTP) to ensure seamless insertion into the network and provide operational simplicity.
[See [JRR200 Route Reflector Hardware Guide](#) and [JRR200 Route Reflector Quick Start](#)]
- **ZTP Support for JRR200 Route Reflector**—Starting in Junos OS Release 19.4R1, ZTP can automate the provisioning of the device configuration and software image on JRR200 Route Reflector. ZTP supports self image upgrades and automatic configuration updates using ZTP DHCP options. In this release, ZTP

supports revenue ports em2 thru em9, in addition to management port em0 which is supported in Junos OS Releases before 19.4R1.

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

SEE ALSO

What's Changed	 59
Known Limitations	 59
Open Issues	 60
Resolved Issues	 60
Documentation Updates	 61
Migration, Upgrade, and Downgrade Instructions	 62

What's Changed

There are no changes in behavior and syntax for JRR Series in Junos OS Release 19.4R2.

SEE ALSO

What's New	 58
Known Limitations	 59
Open Issues	 60
Resolved Issues	 60
Documentation Updates	 61
Migration, Upgrade, and Downgrade Instructions	 62

Known Limitations

There are no known limitations for JRR Series in Junos OS Release 19.4R2.

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

SEE ALSO

What's New	 58
What's Changed	 59
Open Issues	 60
Resolved Issues	 60
Documentation Updates	 61
Migration, Upgrade, and Downgrade Instructions	 62

Open Issues

There are no open issues in hardware and software for JRR Series in Junos OS Release 19.4R2.

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

SEE ALSO

What's New	 58
What's Changed	 59
Known Limitations	 59
Resolved Issues	 60
Documentation Updates	 61
Migration, Upgrade, and Downgrade Instructions	 62

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.4R2](#) | [61](#)
- [Resolved Issues: 19.4R1](#) | [61](#)

Learn which issues were resolved in Junos OS main release and the maintenance releases for JRR Series devices.

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

Resolved Issues: 19.4R2

General Routing

- USB install image does not work for JRR200 platform. [PR1471986](#)
- Link state of virtual em interfaces in Junos OS might not reflect the true link status of corresponding physical interfaces in the Linux host. [PR1492087](#)

Resolved Issues: 19.4R1

There are no fixed issues in Junos OS Release 19.4R1 for JRR Series.

SEE ALSO

What's New	 58
What's Changed	 59
Known Limitations	 59
Open Issues	 60
Documentation Updates	 61
Migration, Upgrade, and Downgrade Instructions	 62

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides](#) | 62

This section lists the errata and changes in Junos OS Release 19.4R2 documentation for JRR Series.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

[What's New](#) | 58

[What's Changed](#) | 59

[Known Limitations](#) | 59

[Open Issues](#) | 60

[Resolved Issues](#) | 60

[Migration, Upgrade, and Downgrade Instructions](#) | 62

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 62

This section contains the upgrade and downgrade support policy for Junos OS for the JRR Series Route Reflector. Upgrading or downgrading Junos OS might take several minutes, 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

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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 about software installation and upgrade, see the [Installation and Upgrade Guide](#) and [JRR200 Route Reflector Quick Start](#).

SEE ALSO

What's New	58
What's Changed	59
Known Limitations	59
Resolved Issues	60
Open Issues	60
Documentation Updates	61

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

●	What's New	64
●	What's Changed	64
●	Known Limitations	65
●	Open Issues	65
●	Resolved Issues	66
●	Documentation Updates	67
●	Migration, Upgrade, and Downgrade Instructions	68

These release notes accompany Junos OS Release 19.4R2 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

NOTE: For a complete list of all hardware and software requirements for a Junos Fusion Enterprise, including which Juniper Networks devices can function as satellite devices, see [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#) .

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

There are no new features or enhancements to existing features in Junos OS Release 19.4R2 for Junos fusion for enterprise.

NOTE: For more information about the Junos Fusion Enterprise features, see the [Junos Fusion Enterprise User Guide](#).

SEE ALSO

What's Changed	 64
Known Limitations	 65
Open Issues	 65
Resolved Issues	 66
Documentation Updates	 67
Migration, Upgrade, and Downgrade Instructions	 68

What's Changed

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 19.4R2 for Junos fusion for enterprise.

SEE ALSO

What's New	 64
Known Limitations	 65
Open Issues	 65
Resolved Issues	 66
Documentation Updates	 67
Migration, Upgrade, and Downgrade Instructions	 68

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 19.4R2 for Junos fusion for enterprise.

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

SEE ALSO

What's New	 64
What's Changed	 64
Open Issues	 65
Resolved Issues	 66
Documentation Updates	 67
Migration, Upgrade, and Downgrade Instructions	 68

Open Issues

IN THIS SECTION

- [Junos Fusion for Enterprise](#) | 66

Learn about open issues in this release for Junos fusion for enterprise. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Junos Fusion for Enterprise

- In a Junos fusion for enterprise environment, when traffic originates from a peer device connected to the aggregation device and the ICL is a LAG, there might be a reachability issue if the cascade port is disabled and traffic has to flow through the ICL LAG to reach the satellite device. As a workaround, use single interface as the ICL instead of a LAG. [PR1447873](#)
- On a Junos fusion for enterprise system, intermediate traffic drop is sometime seen between the aggregation device and satellite device when sFlow is enabled on the ingress interface. [PR1450373](#)
- In Junos fusion for enterprise, the dpd process generate a core file on satellite devices running SNOS. [PR1460607](#)

SEE ALSO

What's New	 64
What's Changed	 64
Known Limitations	 65
Open Issues	 65
Documentation Updates	 67
Migration, Upgrade, and Downgrade Instructions	 68

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 19.4R2](#) | [67](#)

This section lists the issues fixed in Junos OS Release 19.4R2.

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 19.4R2

- The SDPD process generates a core file at `vfpc_all_eports_deletion_complete` `vfpc_dampen_fpc_timer_expiry`. [PR1454335](#)
- Loop detection might not work on extended ports in a Junos fusion scenario. [PR1460209](#)
- The temperature sensor alarm is seen on EX4300 in a Junos fusion scenario. [PR1466324](#)

SEE ALSO

What's New	 	64
What's Changed	 	64
Known Limitations	 	65
Open Issues	 	65
Documentation Updates	 	67
Migration, Upgrade, and Downgrade Instructions	 	68

Documentation Updates

This section lists the errata and changes in Junos OS Release 19.4R2 documentation for Junos Fusion for enterprise.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

What's New	 	64
What's Changed	 	64
Known Limitations	 	65
Open Issues	 	65

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading Junos OS on an Aggregation Device | 68
- Upgrading an Aggregation Device with Redundant Routing Engines | 70
- Preparing the Switch for Satellite Device Conversion | 71
- Converting a Satellite Device to a Standalone Switch | 72
- Upgrade and Downgrade Support Policy for Junos OS Releases | 72
- Downgrading from Junos OS | 73

This section contains the procedure to upgrade or downgrade Junos OS and satellite software for a Junos fusion for enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos Fusion Enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the **junos-install** 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 **junos-install** package and details of the installation process, see the [Installation and Upgrade Guide](#).

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the [Junos OS Administration Library](#).

To download and install Junos OS:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list on the right of the page.
5. Select the **Software** tab.
6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **junos-install** package on the aggregation 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 commands:

```
user@host> request system software add validate reboot source/package.tgz
```

All other customers, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot source/package-limited.tgz
```

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

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

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to minimize disrupting network operations as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.

3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos Fusion Enterprise. See [Configuring or Expanding a Junos Fusion Enterprise](#).

For satellite device hardware and software requirements, see [Understanding Junos Fusion Enterprise Software and Hardware Requirements](#).

Use the following command to install Junos OS on a switch before converting it into a satellite device:

```
user@host> request system software add validate reboot source/package-name
```

NOTE: The following conditions must be met before a Junos switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The switch running Junos OS can be converted only to SNOS 3.1 and later.
- Either the switch must be set to factory-default configuration by using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

When the interim installation has completed and the switch is running a version of Junos OS that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device using the console port.
2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

NOTE: The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port port-number
```

For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 0
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 1
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 2
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 3
```

This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos Fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, or preconfiguration. See [Configuring or Expanding a Junos Fusion Enterprise](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Switch

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove it from the Junos Fusion topology. For more information, see [Converting a Satellite Device to a Standalone Device](#).

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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>

Downgrading from Junos OS

Junos Fusion Enterprise is first supported in Junos OS Release 16.1, although you can downgrade a standalone EX9200 switch to earlier Junos OS releases.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

To downgrade a Junos Fusion Enterprise from Junos OS Release 19.4R2, follow the procedure for upgrading, but replace the 19.2 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

[What's New | 64](#)

[What's Changed | 64](#)

[Known Limitations | 65](#)

[Open Issues | 65](#)

[Resolved Issues | 66](#)

[Documentation Updates | 67](#)

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- [What's New | 74](#)
- [What's Changed | 75](#)
- [Known Limitations | 75](#)
- [Open Issues | 76](#)
- [Resolved Issues | 76](#)
- [Documentation Updates | 77](#)
- [Migration, Upgrade, and Downgrade Instructions | 78](#)

These release notes accompany Junos OS Release 19.4R2 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.4R2.

SEE ALSO

- | |
|--|
| What's Changed 75 |
| Known Limitations 75 |
| Open Issues 76 |
| Resolved Issues 76 |
| Documentation Updates 77 |

What's Changed

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 19.4R2 for Junos fusion for provider edge.

SEE ALSO

[What's New | 74](#)[Known Limitations | 75](#)[Open Issues | 76](#)[Resolved Issues | 76](#)[Documentation Updates | 77](#)[Migration, Upgrade, and Downgrade Instructions | 78](#)

Known Limitations

There are no known limitations for Junos Fusion Provider Edge in Junos OS Release 19.4R2.

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

SEE ALSO

[What's New | 74](#)[What's Changed | 75](#)[Open Issues | 76](#)[Resolved Issues | 76](#)[Documentation Updates | 77](#)[Migration, Upgrade, and Downgrade Instructions | 78](#)

Open Issues

There are no open issues in the Junos OS Release 19.4R2 for Junos Fusion Provider Edge.

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

SEE ALSO

[What's New | 74](#)

[What's Changed | 75](#)

[Known Limitations | 75](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

[Migration, Upgrade, and Downgrade Instructions | 78](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.4R2 Release | 76](#)
- [Resolved Issues: 19.4R1 Release | 77](#)

Learn which issues were resolved in Junos OS main release and the maintenance releases for Junos Fusion Provider Edge.

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

Resolved Issues: 19.4R2 Release

Junos Fusion for Provider Edge

- The `sdpd` process might continuously crash if there are more than 12 cascade-ports configured to a satellite device. [PR1437387](#)

- The aggregated Ethernet interface might flap whenever a new logical interface is added. [PR1441869](#)

Resolved Issues: 19.4R1 Release

There are no fixed issues in the Junos OS Release 19.4R1 for Junos Fusion Provider Edge.

SEE ALSO

[What's New | 74](#)[What's Changed | 75](#)[Known Limitations | 75](#)[Open Issues | 76](#)[Documentation Updates | 77](#)[Migration, Upgrade, and Downgrade Instructions | 78](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 77](#)

This section lists the errata and changes in Junos OS Release 19.4R2 for Junos Fusion Provider Edge.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

What's New 74
What's Changed 75
Known Limitations 75
Open Issues 76
Resolved Issues 76
Migration, Upgrade, and Downgrade Instructions 78

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 78
- Upgrading an Aggregation Device with Redundant Routing Engines | 81
- Preparing the Switch for Satellite Device Conversion | 81
- Converting a Satellite Device to a Standalone Device | 83
- Upgrading an Aggregation Device | 85
- Upgrade and Downgrade Support Policy for Junos OS Releases | 86
- Downgrading from Junos OS Release 19.4 | 86

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

Basic Procedure for Upgrading an Aggregation Device

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 [Installation and Upgrade Guide](#).

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. See the [Installation and Upgrade Guide](#).

The download and installation process for Junos OS Release 19.4R2 is different from that for earlier Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the page.
5. Select the **Software** tab.
6. Select the software package for the release.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the aggregation 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 commands.

- For 64-bit software:

NOTE: We recommend that you use 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

```
user@host> request system software add validate reboot
source/jinstall64-19.4R2.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-19.4R2.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

NOTE: We recommend that you use 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

```
user@host> request system software add validate reboot
source/jinstall64-19.4R2.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-19.4R2.SPIN-export-signed.tgz
```

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

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

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

The **validate** option validates the software package against the current configuration as a prerequisite for adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is for a different release.

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 19.4R2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately as follows to minimize disrupting network operations:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite

device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see [Understanding Junos Fusion Software and Hardware Requirements](#)

NOTE: The following conditions must be met before a standalone switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The switch can be converted to only SNOS 3.1 and later.
- Either the switch must be set to factory-default configuration by using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

Customers with EX4300 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-ex-4300-14.1X53-D43.3-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add reboot
source/jinstall-qfx-5-14.1X53-D43.3-domestic-signed.tgz
```

When the interim installation has completed and the switch is running a version of Junos OS on one line that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device by using the console port.
2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

NOTE: The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device by using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose your connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port port-number
```

For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 0
```

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 1
```

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 2
```

```
user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 3
```

This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos Fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See [Configuring Junos Fusion Provider Edge](#) for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Device

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove the satellite device from the Junos Fusion topology.

NOTE: If the satellite device is a QFX5100 switch, you need to install a PXE version of Junos OS. The PXE version of Junos OS is software that includes *pxe* in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D43 is named `install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz`. If the satellite device is an EX4300 switch, you install a standard `jinstall-ex-4300` version of Junos OS.

The following steps explain how to download software, remove the satellite device from Junos fusion, and install the Junos OS software image on the satellite device so that the device can operate as a standalone device.

1. Using a Web browser, navigate to the Junos OS software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads>
2. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. Select **By Technology > Junos Platform > Junos Fusion** from the drop-down list and select the switch platform series and model for your satellite device.
4. Select the Junos OS Release 14.1X53-D30 software image for your platform.
5. Review and accept the End User License Agreement.
6. Download the software to a local host.
7. Copy the software to the routing platform or to your internal software distribution site.
8. Remove the satellite device from the automatic satellite conversion configuration.

If automatic satellite conversion is enabled for the satellite device's member number, remove the member number from the automatic satellite conversion configuration. The satellite device's member number is the same as the FPC slot ID.

[edit]

```
user@aggregation-device# delete chassis satellite-management auto-satellite-conversion
satellite member-number
```

For example, to remove member number 101 from Junos Fusion:

[edit]

```
user@aggregation-device# delete chassis satellite-management auto-satellite-conversion
satellite 101
```

You can check the automatic satellite conversion configuration by entering the **show** command at the **[edit chassis satellite-management auto-satellite-conversion]** hierarchy level.

9. Commit the configuration.

To commit the configuration to both Routing Engines:

[edit]

```
user@aggregation-device# commit synchronize
```

Otherwise, commit the configuration to a single Routing Engine:

```
[edit]
user@aggregation-device# commit
```

10. Install the Junos OS software on the satellite device to convert the device to a standalone device.

```
[edit]
user@aggregation-device> request chassis satellite install URL-to-software-package fpc-slot  
member-number
```

For example, to install a PXE software package stored in the `/var/tmp` directory on the aggregation device onto a QFX5100 switch acting as the satellite device using FPC slot 101:

```
[edit]
user@aggregation-device> request chassis satellite install  
/var/tmp/install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz fpc-slot 101
```

For example, to install a software package stored in the `var/tmp` directory on the aggregation device onto an EX4300 switch acting as the satellite device using FPC slot 101:

```
[edit]
user@aggregation-device> request chassis satellite install  
/var/tmp/jinstall-ex-4300-14.1X53-D30.3-domestic-signed.tgz fpc-slot 101
```

The satellite device stops participating in the Junos Fusion topology after the software installation starts. The software upgrade starts after this command is entered.

11. Wait for the reboot that accompanies the software installation to complete.
12. When you are prompted to log back into your device, uncable the device from the Junos Fusion topology. See *Removing a Transceiver from a QFX Series Device* or *Remove a Transceiver*, as needed. Your device has been removed from Junos Fusion.

NOTE: The device uses a factory-default configuration after the Junos OS installation is complete.

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 19.4R2, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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>.

Downgrading from Junos OS Release 19.4

To downgrade from Release 19.4 to another supported release, follow the procedure for upgrading, but replace the 19.4 **jinstall** package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

SEE ALSO

[What's New | 74](#)

[What's Changed | 75](#)

[Known Limitations | 75](#)

[Open Issues | 76](#)

[Resolved Issues | 76](#)

[Documentation Updates | 77](#)

Junos OS Release Notes for MX Series 5G Universal Routing Platform

IN THIS SECTION

- [What's New | 87](#)
- [What's Changed | 110](#)
- [Known Limitations | 116](#)
- [Open Issues | 119](#)
- [Resolved Issues | 135](#)
- [Documentation Updates | 168](#)
- [Migration, Upgrade, and Downgrade Instructions | 169](#)

These release notes accompany Junos OS Release 19.4R2 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- [What's New in 19.4R2 Release | 88](#)
- [What's New in 19.4R1 Release | 89](#)

Learn about new features introduced in the Junos OS main release and the maintenance releases for MX Series routers.

What's New in 19.4R2 Release

Hardware

NOTE: The MX2K-MPC11E line card is not supported in any Junos OS 19.4 releases. It is supported in Junos OS 19.3R2 and later 19.3 releases and in Junos OS 20.1R1 and later Junos OS releases.

Class of Service (CoS)

- **L2 CoS classifiers and rewrite feature on MX Series routers (MX240, MX480, MX960) with the MPC10 line card**—Starting in Junos OS Releases 19.3R2 and 19.4R1, support is provided for L2 Class of Service (CoS) classifiers and rewrite on MX Series routers (MX240, MX480, MX960) with the MPC10 line card.

[See [Protocols and Applications Supported by the MPC10E.](#)]

Interfaces and Chassis

- **Transparent forwarding of CFM packets over VPLS (MX Series)**—In Junos OS Release 19.4R2, an MX Series router continues to forward CFM packets over a VPLS even if the packets contain more VLAN tags than the number configured on an interface. In earlier releases of Junos OS, irrespective of the number of tags in a packet, the router forwards or drops the packet based on the interface configuration. In those releases, all CFM messages are parsed at the provider edge. But as a result of this feature support, only those CFM packets are parsed for which the number of VLAN tags matches the number configured on the provider edge interface.

NOTE: We do not support transparent forwarding on untagged and triple-tagged CFM packets.

[See [Example: Configuring Ethernet CFM over VPLS.](#)]

MPLS

- **Delegate segment routing LSPs to a PCE (MX Series)**—Starting in Junos OS Release 19.4R2, you can enable a Path Computation Client (PCC, ingress MX Series router) to delegate locally configured IPv4 non-colored segment routing LSPs to a Path Computation Element (PCE) controller. The PCE controls the delegated LSPs and can modify LSP attributes for traffic steering. A PCC with delegation capability can take back control of the delegated segment routing LSPs from the PCE when the PCEP session goes down; the LSPs would otherwise be deleted from the PCC. You can thus ensure LSP data protection by averting a situation where packets are silently discarded or dropped (also known as a traffic black-hole condition).

[See [Example: Configuring Path Computation Element Protocol for SR-TE LSPs.](#)]

Subscriber Management and Services

- **RADIUS-sourced connection status updates to CPE devices (MX Series)**—Starting in Junos OS Release 19.4R2, you can use RADIUS-sourced messages to convey information, such as upstream bandwidth or connection rates, that the BNG transparently forwards to CPE devices. Configure RADIUS to send the router the Juniper Networks Connection-Status-Message VSA (26-4874–218) in Access-Accept or CoA messages. Include the **lcp-connection-update** PPP option in the client dynamic profile to enable PPP to send the VSA contents to the CPE device in the Connection-Status-Message option of an LCP Connection-Update-Request message.

[See [RADIUS-Sourced Connection Status Updates to CPE Devices.](#)]

What’s New in 19.4R1 Release

Hardware

NOTE: The MX2K-MPC11E line card is not supported in any Junos OS 19.4 releases. It is supported in Junos OS 19.3R2 and later 19.3 releases and in Junos OS 20.1R1 and later Junos OS releases.

- **MX-SPC3 Services Card (MX240, MX480, and MX960)**—Starting with Junos OS Release 19.3R2, the MX-SPC3 Services Card is available on MX240, MX480, and MX960 routers. The MX-SPC3 card provides additional processing power to run Next Gen Services. The MX-SPC3 contains two Services Processing Units (SPUs) with 128 GB of memory per SPU. Line cards such as DPCs, MICs, and MPCs intelligently distribute all traffic traversing the router to the SPUs to have services processing applied to it.

Next Gen Services provide the best of both routing and security features on MX Series routers MX240, MX480, and MX960. All Next Gen Services are provided by the MX-SPC3 services card. Next Gen Services provide capabilities for manipulating traffic before it’s delivered to its destination. Next Gen Services features run on the MX Series, and are based on a different software architecture than legacy MX Series services. You can run Next Gen Services on MX240, MX480 and MX960 routers. Some Next Gen Services features use different Junos CLI statements than the equivalent legacy service.

NOTE: The only services card that supports Next Gen Services is the MX-SPC3. Next Gen Services use their own software architecture, which is not compatible with legacy services.

[Table 1 on page 91](#) summarizes the Next Gen Services supported in this release.

Table 1: Next Gen Services Summary

Next Gen Services Supported by MX-SPC3 Services Card	
Carrier Grade NAT	6rd Softwires
	Deterministic NAT
	Dynamic Address-Only Source NAT
	Global System Logging
	IPv4 Connectivity Across IPv6-Only Network Using 464XLAT
	Network Address Port Translation
	Port Forwarding
	Static Source NAT
	Stateful NAT64
	Static Destination NAT
	Stateless Source Network Prefix Translation for IPv6
	Twice NAPT
	Twice Static NAT
	Class of Service
Stateful Firewall Services	
Intrusion Detection Services	
Traffic Load Balancing	
DNS Request Filtering	
Aggregated Multiservices Interfaces	
Inter-chassis High Availability	NAT, Stateful Firewall, and IDS Flows

Table 1: Next Gen Services Summary (continued)**Next Gen Services Supported by MX-SPC3 Services Card**

See [Protocols and Applications Supported by MX-SPC3 Services Card](#) for information about the protocols and applications that this SPC3 supports.

The MX-SPC3 services card is compatible end-to-end with the MX Series Switch Fabrics, Routing Engines and MS-MPC line cards. See [Table 2 on page 92](#):

Table 2: MX-SPC3 Services Card Compatibility with MX Series Switch Fabrics, Routing Engines and MPC Line Cards

Switch Fabric	Route Engine	MPC Line Cards
SCBE	RE-S-1800X4-16G-UPG-BB	MPC2E-3D
	RE-S-1800X4-32G-UB	MPC2-3D-NG
		MPC3E and MPC3E-3D-NG
		MPC4E-3D
		MPC-3D-16XGE
SCBE2	RE-S-1800X4-16G-UPG-BB	MPC2E-3D
	RE-S-1800X4-32G-UB	MPC2-3D-NG
	RE-S-X6-64G-UB	MPC3E and MPC3E-3D-NG
		MPC4E-3D
		MPC5E and MPC5EQ
		MPC7E, MPC7EQ, and MPC-3D-16XGE
		MPC-3D-16XGE

Refer to our [TechLibrary](#) for all MX router documentation. For Next Gen Services, refer to the following documentation: See

- [Next Gen Services Interfaces Overview for Routing Devices](#)
- [Next Gen Services Interfaces User Guide for Routing Devices](#)
- [Broadband Subscriber Services Feature Guide](#)
- [Monitoring, Sampling, and Collection Services Interfaces Feature Guide](#)
- [MX240 Universal Routing Platform Hardware Guide](#)

- [MX480 Universal Routing Platform Hardware Guide](#)
- [MX960 Universal Routing Platform Hardware Guide](#)
- [MX Series 5G Universal Routing Platform Interface Module Reference](#)

Class of Service (CoS)

- **Support for seamless MPLS Layer 3 features (MX Series with MPC10E line cards)**—Starting in Junos OS Release 19.4R1, the following MPLS Layer 3 features are supported on MX Series routers with MPC10E line cards:
 - Redundant logical tunnel interfaces.
 - Pseudowire subscriber interfaces using logical tunnel or redundant logical tunnel interfaces as anchor point.

[See [Redundant Logical Tunnels Overview](#), and [MPLS Pseudowire Subscriber Logical Interfaces](#).]

EVPN

- **Support for EVPN routing policies (ACX5448, EX4600, EX4650, EX9200, MX Series, QFX Series, and vMX)**—Starting in Junos OS Release 19.4R1, Junos OS has expanded routing policy support to include the creation and application of policy filters specific to EVPN routes. You can create policies and apply policy filters to import and export EVPN routes at the routing-instance level or at the BGP level. Junos OS supports the following matching criteria for EVPN routes:
 - Route distinguisher ID
 - NLRI route type
 - EVPN Ethernet tag
 - BGP path attributes
 - Ethernet segment identifier
 - MAC address on EVPN Type 2 routes
 - IP address on EVPN Type 2 and EVPN Type 5 routes
 - Extended community

[See [Routing policies for EVPN](#).]

- **Exclusion list with MAC pinning in an EVPN network (EX9200 and MX Series)**—When you enable **mac-pinning** on an interface, all MAC addresses that are learned on that interface will be pinned and cannot be relearned on the other interfaces in the EVPN network. Starting in Junos OS Release 19.4R1, you can create a list of MAC addresses that would be excluded from being pinned and the MAC address can be moved and relearned on another interface within the EVPN network. While MAC pinning is configured on the interface, the exclusion list is configured for the device. To create an exclusion list, include a list of MAC addresses with the **exclusive-mac** parameter at the **[edit protocols l2-learning global-mac-move]** hierarchy level.

[See [Creating exclusion list for MAC Pinning](#).]

- **Support for EVPN functionality (MX Series with MPC10 line card)**—Starting in Junos OS 19.4R1, you can configure MPC10 line cards on a MX Series router to support single-homed devices on an EVPN-MPLS network.

[See [EVPN Multihoming Overview](#).]

Forwarding and Sampling

- **Inline monitoring services (MX Series with MPCs excluding MPC10E)**—Starting in Junos OS Release 19.4R1, you can configure a new monitoring technology that provides the flexibility to monitor different streams of traffic at different sampling rates on the same interface. You can also export the packet up to the configured clip length to a collector in an IP Flow Information Export (IPFIX) format. The IPFIX format includes important metadata information about the monitored packets for further processing at the collector.

The inline monitoring services overcome the limitations of traditional sampling technologies, such as JFlow, sFlow, and port mirroring, thereby providing you the benefit of effective sampling and troubleshooting processes.

[See [Inline Monitoring Services Configuration](#).]

- **Improved failover in conjunction with consistent load balancing for ECMP groups (MX Series routers with MPCs)**—Starting in Junos OS Release 19.4R1, we've added new functionality to prevent the reordering of flows to already active paths in an equal-cost multipath routing (ECMP) group if one or more path next-hops go down. Before this feature, when a server in the ECMP path failed, the flows directed to that server were redistributed to other, active links. If a second server in the ECMP path also went down, the newly redistributed traffic would be redistributed again, even though the ECMP path is active. The improved failover and traffic rebalancing introduced in this release minimize the traffic redistribution when multiple servers in the ECMP path fail.

[See [Configuring Consistent Load Balancing for ECMP Groups](#) and [Load Balance Traffic on MX Series Routers](#).]

General Routing

- **Optimized BGP peer reestablishment (MX Series, PTX Series, and QFX Series)**—Starting with Junos OS Release 19.4R1, BGP peers in different groups can close in parallel. The connect/retry algorithm makes more frequent attempts to reestablish BGP peers, which reduces downtime. The connect/retry algorithm makes 16 attempts instead of 5 to reestablish BGP peers in the first 256 seconds after they go down. Peers can reestablish while cleanup of the Adj-RIB-In routes is in progress. If a peer comes back up before its route has been deleted from the routing table, that route is not deleted. The **DeletePending** flag in the **show route detail** and **show route extensive** command output indicates that a BGP route needs to be processed. **PurgePending**, **PurgeInProgress**, and **PurgeImpatient** flags in the **show bgp neighbor** command output show the status of the purge of routing table entries.

[See [Understanding External BGP Peering Sessions](#), [show bgp neighbor](#), [show route detail](#), and [show route extensive](#).]

High Availability (HA) and Resiliency

- **View ISSU status during an upgrade (MX240, MX480, MX960, MX2010, MX2020, PTX3000, and PTX5000)**—Starting in Junos OS Release 19.4R1, you can use the **request system software in-service-upgrade status** command to display the status of a unified ISSU. You will need to run this command on the Routing Engine where the unified ISSU was triggered to display the correct unified ISSU log file.

[See [request system software in-service-upgrade](#).]

Interfaces and Chassis

- **New universal PSM and PDM (MX2008, MX2010, and MX2020)**—Starting in Release 19.4R1, Junos OS supports the high-voltage second-generation universal power supply module (PSM; model number: MX2K-PSM-HV) and power distribution module (PDM; model number: MX2K-PDM-HV). The PSM has a main output and a standby output. The main output provides up to 3000 W power with a single feed, and up to 3400 W power with dual feeds. The standby output provides up to 30 W power. The PSM accepts either a AC input (voltage range: 180 VAC through 305 VAC) or DC input (voltage range: 190 VDC through 410 VDC). Each universal PDM has nine HVAC/HVDC inputs.

NOTE: We recommend that you use MX2K-PSM-HV PSM only with MX2K-PDM-HV PDM.

[See [MX2010 Power System Description](#) and [MX2020 Power Subsystem Description](#).]

- **High-capacity second-generation AC PSM (MX960)**—Starting in Release 19.4R1, Junos OS supports the new high-capacity second-generation AC power supply module (PSM; model number: MX960-PSM-5K-AC-S) on MX960 routers. An enhanced version of the existing PSM used in the MX960 chassis, the new high-capacity PSM provides a maximum output power of 5100 W with dual feeds, and 2550 W with a single feed. The PSM supports a minimum input voltage of 180 VAC and a maximum input voltage of 264 VAC. The PSM supports 1+1 redundancy.

[See [MX960 Power System Overview](#).]

- **Smart SFP transceivers for encapsulating and transporting PDH traffic (MX Series routers)**—Starting in Junos OS Release 19.4R1, on MX Series routers with MPCs (MPC1, MPC2, and MPC3) and MICs, you can configure and manage the following smart SFP transceivers to encapsulate PDH traffic:
 - DS3 smart SFP (SFP-GE-TDM-DS3)
 - E1 smart SFP (SFP-GE-TDM-E1)
 - T1 smart SFP (SFP-GE-TDM-T1)

The transceivers encapsulate PDH (E1 or T1 or DS3) packets as Ethernet frames while transporting legacy time division multiplexing (TDM) traffic over packet switched networks (PSNs). At the receiver end of the emulated circuit, another smart transceiver, paired with the first one and preconfigured to carry packets that are in the same multicast MAC address group, de-encapsulates the Ethernet frames, rebuilds the TDM data stream, and forwards it onto the local TDM interface.

- **Support for 1-Gbps speed on 10-Gbps port (JNP10K-LC2101 line card on MX10008 and MX10016)**—Starting in Junos OS Release 19.4R1, you can configure the 10-Gigabit Ethernet port on the JNP10K-LC2101 line card to operate at 1-Gbps speed by using the **speed** statement at the **[edit interfaces *interfacename* *gether-options*]** hierarchy level. After you commit the configuration, the operating speed of the 10-Gbps port changes to 1-Gbps speed.

To view the speed configured for the interface, use the **show interfaces extensive** command. The **SpeedConfiguration** field in the command output indicates the current operational speed of the interface. If the interface is configured with 1-Gbps speed, then the value of the **SpeedConfiguration** field is displayed as **1G**; if the interface is configured with 10-Gbps speed, then **SpeedConfiguration** displays **AUTO**.

Autonegotiation is supported when the interface speed is configured for 1-Gbps speed.

NOTE: On the JNP10K-LC2101 line card, rate selectability at PIC level and port level does not support 1-Gbps speed.

[See [Introduction to Rate Selectability](#).]

- **Support for monitoring link degradation (MX Series routers with MPC7E, MPC8E, and MPC9E)**—Starting in Junos OS Release 19.4R1, you can monitor the quality of physical links on Ethernet Interfaces and take corrective action when the link quality degrades beyond a certain value. To enable your device to monitor the links, use the **link-degrade-monitor** statement at the **[edit interfaces *interface-name*]** hierarchy level. This feature monitors the bit error value (BER) of the link and initiates corrective action when the BER value crosses a user-configured threshold.

Starting in Junos OS Release 19.4R1, the following line cards support link degrade monitoring:

- MPC7E (MPC7E-MRATE and MPC7E-10G (non-MACsec mode))

- MPC8E (MIC-MRATE MICs)
- MPC9E (MIC-MRATE MICs)

NOTE: Link degrade monitoring is not supported on the MACsec-enabled MPC7E-10G and MIC-MACSEC-MRATE.

[See [Link Degrade Monitoring Overview](#).]

- **Optimize fabric path to prevent traffic hop (MX2008, MX2010, and MX2020 with MPC9E)**—Starting in Junos OS Release 19.4R1, you can optimize the fabric path of the traffic flowing over abstracted fabric (af) interfaces between two guest network functions (GNFs) by configuring a fabric optimization mode. This feature reduces fabric bandwidth consumption by preventing any additional fabric hop (switching of traffic flows from one Packet Forwarding Engine to another because of abstracted fabric interface load balancing) before the packets eventually reach the destination Packet Forwarding Engine.

To configure fabric optimization mode, use the following CLI commands at the base system (BSYS): **set chassis network-slices guest-network-functions gnf *id* collapsed-forward <monitor | optimize>**.

[See [Optimizing Fabric Path for Abstracted Fabric Interface](#).]

- **SCBE3-MX interoperates with MPC 3D 16x10GE (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.4R1, the Enhanced Switch Control Board SCBE3-MX (model number: SCBE3-MX-S) supports the 16-port 10-Gigabit Ethernet MPC (MPC 3D 16x10GE) on the MX240, MX480, and MX960 routers with enhanced midplane. The SCBE3-MX-S supports a pluggable Routing Engine and provides a control plane and data plane interconnect to each line card slot. The MPC 3D 16x10GE supports a fabric bandwidth of 160 Gbps.

[See [SCBE3-MX Description](#) and [16x10GE MPC](#).]

- **Enhancement to increase the threshold of corrected single-bit errors (MPC7E, MPC8E, and MPC9E on MX Series)**—In Junos OS Release 19.4R1, the threshold of corrected single-bit errors is increased from 32 to 1024, and the alarm severity is changed from Major to Minor for those error messages. There is no operational impact upon corrected single-bit errors. Also, a log message is added to display how many single-bit errors have been corrected between the reported events as follows:

EA[0:0]: HMCIF Rx: Link0: Corrected single bit errordetected in HMC 0 - Total count 25

EA[0:0]: HMCIF Rx: Link0: Corrected single bit errordetected in HMC 0 - Total count 26

[See [Alarm Overview](#).]

Junos OS, XML, API, and Scripting

- **Python 3 support for commit, event, op, and SNMP scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, you can use Python 3 to execute commit, event, op, and SNMP scripts on devices running Junos OS. To use Python 3, configure the **language python3** statement at the **[edit system scripts]** hierarchy level. When you configure the

language python3 statement, the device uses Python 3 to execute scripts that support this Python version and uses Python 2.7 to execute scripts that do not support Python 3 in the given release.

The Python 2.7 end-of-support date is January 1, 2020, and Python 2.7 will be EOL in 2020. The official upgrade path for Python 2.7 is to Python 3. As support for Python 3 is added to devices running Junos OS for the different types of onbox scripts, we recommend that you migrate supported script types from Python 2 to Python 3, because support for Python 2.7 might be removed from devices running Junos OS in the future.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

- **Automation script library upgrades (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, devices running Junos OS that support the Python extensions package include upgraded Python modules. Python scripts can leverage the upgraded versions of the following modules:

- **idna** (2.8)
- **jinja2** (2.10.1)
- **jnpr.junos** (Junos PyEZ) (2.2.0)
- **lxml** (4.3.3)
- **markupsafe** (1.1.1)
- **ncclient** (0.6.4)
- **packaging** (19.0)
- **paho.mqtt** (1.4.0)
- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

[See [Overview of Python Modules Available on Devices Running Junos OS.](#)]

- **Support for 64-bit architecture added for use of management interface in a nondefault routing instance in op scripts and JET applications (MX Series)**—Junos OS Release 19.4R1 supports 64-bit architecture for Junos operating scripts and on-box JET applications being able to use the function `set_routing_instance()` to program the protocol software (TCP/UDP) to use a nondefault routing instance instead of the default management routing interface.

[See [set_routing_instance\(\) Function \(Python\).](#)]

Junos Telemetry Interface

- **Transceiver sensor support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000)**—In Junos OS Release 19.4R1, you can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) or gRPC Network Management Interface (gNMI) services to export transceiver statistics from MX960, MX2010, MX2020, PTX1000 and PTX5000 routers to outside collectors. This feature supports OpenConfig transceiver model `openconfig-platform-transceiver.yang` 0.5.0.

Both streaming and ON-CHANGE statistics are supported using the following base path:

- `/components/components/transceiver/`

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#) and [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

- **Physical Ethernet interface sensor support on JTI (MX960, MX2020, PTX1000, PTX5000)**—Starting in Junos OS Release 19.4R1, you can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) services or gRPC Network Management Interface (gNMI) services to export physical Ethernet interface statistics from MX960, MX2020, PTX1000, and PTX5000 routers to outside collectors. This feature supports OpenConfig model `openconfig-if-ethernet.yang` (physical interface level) version 2.6.2 (no configuration). Both streaming and ON-CHANGE statistics are supported using the following resource paths:

- `/interfaces/interface/ethernet/state/mac-address` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/auto-negotiate` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/duplex-mode` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/port-speed` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/enable-flow-control` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/hw-mac-address` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/negotiated-duplex-mode` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/negotiated-port-speed` (with ON_CHANGE support)
- `/interfaces/interface/ethernet/state/counters/in-mac-control-frames`
- `/interfaces/interface/ethernet/state/counters/in-mac-pause-frames`
- `/interfaces/interface/ethernet/state/counters/in-oversize-frames`
- `/interfaces/interface/ethernet/state/counters/in-jabber-frames`
- `/interfaces/interface/ethernet/state/counters/in-fragment-frames`
- `/interfaces/interface/ethernet/state/counters/in-8021q-frames`
- `/interfaces/interface/ethernet/state/counters/in-crc-errors`
- `/interfaces/interface/ethernet/state/counters/in-block-errors`
- `/interfaces/interface/ethernet/state/counters/out-mac-control-frames`
- `/interfaces/interface/ethernet/state/counters/out-mac-pause-frames`

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

- **ON_CHANGE support for component sensors on JTI (MX960)**—Junos OS Release 19.4R1 supports ON_CHANGE statistics for the following component sensors using Junos telemetry interface (JTI) and either remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services.

Junos OS releases before Release 19.4R1 support these component sensors on the MX960 router only to stream statistics.

- `/components/component`
- `/components/component/name/`
- `/components/component/state/type`
- `/components/component/state/id`
- `/components/component/state/description`
- `/components/component/state/serial-no`
- `/components/component/state/part-no`

Streaming telemetry data through gRPC or gNMI requires the OpenConfig for Junos OS module.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Layer 2 Features

- **Support for an increase in MAC table size on the MPC10E-15C-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.4R1, MX Series routers with MPC10E-15C-MRATE line cards support MAC table size of upto 1 million entries per PFE for Layer 2 services.

You can configure the MAC limit size at global level at the `[edit protocols l2-learning global-mac-limit]` hierarchy level.

You can also configure the MAC table size using bridge domains at the `[edit bridge-domains bridge-domain-name bridge-options mac-table-size]` hierarchy level.

[See [Understanding Layer 2 Bridge Domains](#) , [Understanding Layer 2 Learning and Forwarding](#) .]

Layer 2 VPN

- **Support for VPLS (MX series with MPC10 line card)**—Starting in Junos OS 19.4R1, you can configure VPLS on the MPC 10 line card in a MX Series router.

[See [Introduction to VPLS](#) and [VPLS Configuration Overview](#).]

MPLS

- **Distributed CSPF for segment routing LSPs (MX Series)**—Starting in Junos OS Release 19.4R1, you can compute a segment routing LSP locally on the ingress device according to the constraints you have configured. With this feature, the LSPs are optimized based on the configured constraints and metric type. The LSPs are computed to utilize the available ECMP paths to the destination.

Prior to Junos OS Release 19.4R1, for traffic engineering of segment routing paths, you could either explicitly configure static paths, or use computed paths from an external controller.

[See [Enabling Distributed CSPF for Segment Routing LSPs](#).]

- **update-threshold statement modified to generate IGP update for lower bandwidth reservation (MX Series)**—Starting in Junos OS Release 19.4R1, you can configure the threshold value of the **update-threshold** statement to accept:
 - an integer or floating point values up to 3 significant digits after decimal point using the **threshold-percent** option
 - an absolute value of bandwidth threshold which generates an IGP update using the **threshold-value** option

These options are mutually exclusive and can be used for generating an IGP update for lower bandwidth reservations.

[See [update-threshold](#).]

- **Color-based mapping of VPN services over SRTE (MX Series)**—Starting in Junos OS Release 19.4R1, you can specify a color attribute along with an IP protocol next hop to resolve transport tunnels over static colored and BGP segment routing traffic-engineered (SRTE) label-switched paths (LSPs). This is called the color-IP protocol next hop resolution, where you are required to configure a resolution-map

and apply it to the VPN services. Prior to this release, the VPN services were resolved over IP protocol next hops only.

With this feature, you can enable color-based traffic steering of Layer 2 and Layer 3 VPN services.

[See [Color-Based Mapping of VPN Services Overview](#).]

- **Support for static adjacency segment identifier for aggregated Ethernet member links on MPC10E-15C-MRATE line cards (MX240, MX480, and MX960)**—Starting with Junos OS Release 19.4R1, you can configure a transit single-hop static label-switched path (LSP) for a specific member link of an aggregated Ethernet (ae) interface. The label for this route comes from the segment routing local block (SRLB) pool of the configured static label range. Configure the ae member interface name using the **member-interface** statement option at the **[edit protocols mpls static-label-switched-path name transit name]** hierarchy level. This feature is supported for ae interfaces only.

[See [transit](#) and [Configuring Static Adjacency Segment Identifier for Aggregate Ethernet Member Links Using Single-Hop Static LSP](#).]

Multicast

- **Next-generation multicast VPN supported on MPC10E-15C-MRATE line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.4R1, the MPC10E-15C-MRATE line card supports next-generation MVPN.

[See [Multicast Overview](#).]

- **Continuous, persistent IGMP and MLD statistics (MX Series)**—Starting in Junos OS Release 19.4R1, in addition to collecting statistics on IGMP and MLD control traffic for currently active subscribers, MX Series routers also collect and maintain cumulative and persistent statistics that account for both past and current subscribers. The device stores these statistics and copies them to the backup Routing Engine at a configurable interval, so this information is preserved across routing daemon restarts, graceful Routing Engine switchovers (GRES), in-service software upgrade (ISSU) operations, or line card reboots. Use the **continuous** option with the **show igmp statistics** or **show mld statistics** command to view continuous statistics; without this option, you see default statistics only for currently active subscribers.

[See [show igmp statistics](#) or [show mld statistics](#).]

Network Management and Monitoring

- **Packet mirroring with Layer 2 headers for Layer 3 forwarded traffic (MX Series routers with MPCs or MICs)**—Starting in Junos OS Release 19.4R1, you can enable port mirroring at packet level along with Layer 2 headers even if the filters are installed with Layer 3 match actions. Use the new firewall-filter action **I2-mirror** at the `[edit firewall family inet|inet6 filter filter-name term tcp-flags then]` hierarchy level to request Layer 2 header reporting.

OAM

- **Support for Ethernet OAM and Metro Ethernet services over segment routing (ACX5448-D, ACX5448-M, MX Series)**—Starting with Junos OS Release 19.4R1, ACX5448-D, ACX5448-M and MX Series routers support Ethernet OAM and Metro Ethernet services over segment routing.

[See Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING for IS-IS Protocol, Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS, Ethernet OAM Connectivity Fault Management .]

Routing Policy and Firewall Filters

- **Support for firewall forwarding on MPC10E line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.4R1, the following traffic policers are fully supported on MX240, MX480, and MX960 routers with MPC10E line cards:

- GRE tunnels, including encapsulation (**family any**), decapsulation, GRE-in-UDP over IPv6, and the following sub-options: sample, forwarding class, interface group, and no-ttl-decrement
- Input and output filter chains
- Actions, including policy-map filters, do-not-fragment, and prefix
- Layer 2 policers
- Policer overhead adjustment
- Hierarchical policers
- Shared bandwidth
- Percentages
- Logical interfaces

[See [Traffic Policer Types](#).]

- **GTP load balancing on MPC10E-15C-MRATE line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.4R1, the MPC10E-15C-MRATE line card supports GPRS tunneling protocol (GTP) load balancing.

[See [Understanding Per-Packet Load Balancing](#).]

Routing Protocols

- **Integrating RIFT protocol into Junos OS (MX240, MX480, MX960, QFX5100, QFX5110, QFX5120-32C, QFX5120-48Y, QFX5120-48YM, QFX5200, QFX5210, QFX10008, and VMX virtual routers)**—Starting in Junos OS Release 19.4R1, you can integrate a new IGP protocol, Routing in Fat Tree (RIFT), into Junos OS to route packets in variants of CLOS-based and fat tree network topologies (also called the spine and leaf model).

The RIFT protocol is capable of automatic construction of fat-tree topologies, providing you the benefit of having a close to zero necessary configuration. RIFT makes networks resilient, extensively traceable, and simpler to manage, thereby overcoming the deployment limitations of evolving IP fabrics.

[See [RIFT Overview and Set Up](#).]

- **Bidirectional Forwarding Detection (BFD) Strict Mode for OSPF (MX Series)**—Starting in Release 19.4R1, Junos OS supports BFD strict mode for OSPF. The BFD strict mode for OSPF enables a router to prevent establishing OSPF adjacency until a BFD session is established. This helps in faster and more reliable connection with the peer devices. To enable this feature, both the devices should support BFD strict-mode.

To configure BFD strict-mode, use `set strict-bfd` at the `[edit protocols ospf area area_id interface interface_name]` hierarchy level.

You can also configure a hold down interval to delay the sending of session UP notification to the BFD client which helps in achieving a more stable connection. To configure a hold down interval, use `set holddown-interval holddown-interval` at the `[edit protocols ospf area area_id interface interface_name bfd-liveness-detection]` hierarchy level.

- **Support for BGP Update Threading (MX Series and VRR)**—Starting in Junos OS Release 19.4R1, the BGP protocol work to do Update message generation for peers in a BGP group is moved out from the main BGP thread to its own new set of pthreads, called BGP Update I/O threads. Each Update I/O thread is responsible for generating updates for one or more BGP peer groups. BGP Update threads construct updates for groups in parallel and independent of other groups that are being serviced by different update threads. This might offer significant convergence improvement in a write-heavy workload that involves advertising to many peers spread across many groups. BGP Update I/O threads can be configured independent of RIB sharding feature but are mandatory to use with RIB sharding as they help improve packing of prefixes in outbound BGP update messages and thus help improve performance.

BGP update thread is disabled by default. If you configure update-threading on a routing engine, RPD creates update threads. By default, the number of update threads created is the same as the number of CPU cores on the routing engine. Update threading is only supported on a 64 bit routing protocol process (rpd). Optionally, you can specify the number-of-threads you want to create by using `set update-threading <number-of-threads>` statement at the `[edit system processes routing bgp]` hierarchy level. The range is currently 1 through 128.

See [\[update-threading\]](#) and [\[Understanding BGP UPDATE IO Thread\]](#).

- **Support for BGP RIB Sharding (MX Series and VRR)**—Starting in Junos OS Release 19.4R1, the BGP process is split into different threads so that they can run concurrently on a multicore routing engine through RIB sharding which results in reduced convergence time and faster performance. BGP RIB

sharding splits a BGP RIB into several sub RIBs and each sub RIB handles a subset of BGP routes. Each sub RIB is served by a separate RPD thread to achieve parallel processing.

BGP RIB sharding is disabled by default. This feature is supported only on 64-bit routing protocol process (rpd) where the Routing Engine has at least 4 CPU cores and 16 GB of memory.

If you configure rib-sharding on a routing engine, RPD will create sharding threads. By default the number of sharding threads created is same as the number of CPU cores on the routing engine. Optionally, you can specify the number-of-shards you want to create. The range is currently 1 through 31.

The **show route command** shows the aggregate data from main and all shards to provide the unified view of the RIBs.

NOTE: BGP RIB sharding is supported for inet.0 and inet6.0 RIBs only. All the other RIBs are still processed without sharding.

To enable this feature, you can configure **rib-sharding** at the **[edit system processes routing bgp]** hierarchy level. Sharding is dependent on the update I/O thread feature. Therefore, update I/O thread feature is essential and mandatory when you configure sharding. To enable update I/O, you need to configure **update-threading** at the **[edit system processes routing bgp]** hierarchy level for **rib-sharding** configuration to pass **commit check**.

See [\[rib-sharding.\]](#)

- **Optimizing the static route configuration processing during commit (MX Series)**— Starting in Release 19.4R1, Junos OS optimizes the static route configuration processing during commit by managing only the new, modified and deleted routes instead of all the routes. The processing of these static route configurations are optimized:
 - Local SRGB
 - Global SRGB
 - Node-segment implementation of 256 label block
- **VRRP support for MPC10E-15C-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.4R1, MPC10E-15C-MRATE line cards support VRRP.

[See [Understanding VRRP.](#)]

- **Unnumbered interface support for IS-IS and OSPFv2 with topology-independent loop-free alternate (ACX Series, MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, you can enable IPv4 processing on a point-to-point interface without assigning it an explicit IPv4 address. The router borrows the IPv4 address of another Ethernet or loopback interface already configured on the router and assigns it to the unnumbered interface to conserve IPv4 addresses.

To enable IPv4 processing for unnumbered interfaces include **unnumbered-address source** at the **[edit interfaces [name] unit [name] family inet]** hierarchy level.

[See [Configuring an Unnumbered Interface](#).]

- **Support for flexible algorithm in IS-IS for segment routing–traffic engineering (MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, you can thin slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path based on traffic engineering metric to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

To define a flexible algorithm, include **flex-algorithm** statement at the **[edit routing-options]** hierarchy level.

To configure participation in a flexible algorithm include the **flex-algorithm** statement at the **[edit protocols isis segment routing]** hierarchy level.

[See [Understanding IS-IS Flexible Algorithm for Segment Routing](#).]

- **Support for disable-4byte-as and minimum-hold-time configurations (MX Series)**—Starting in Junos OS Release 19.4R1, you can use the **minimum-hold-time** and **disable-4byte-as** configurations. By configuring **minimum-hold-time**, you can now prevent BGP session establishment toward BGP peers that attempt to negotiate a lower BGP session hold-time than the configured **minimum-hold-time**, which helps reduce the load on a router by avoiding sending constant keepalive messages at a high frequency. You can use **disable-4byte-as** configuration to enable a BGP peer that uses a 4-Byte to interact with another BGP peer old speaker that uses 2-Byte.

NOTE:

- We recommend using Bidirectional Forwarding Detection (BFD) rather than lowering BGP hold timers and also recommend configuring a meaningful **minimum-hold-time** value (for example, 20 seconds or higher) for all BGP peers (for example, at the BGP group level). If a BGP remote node does not support BFD, and therefore a reduced BGP hold-time is easier for the quicker discovery of a BGP neighbor failure, you can configure a lower **minimum-hold-time** value. Use it with caution and only for a limited number of BGP peers.
- We recommend that you configure the **disable-4byte-as** configuration only if the BGP peer does not support or ignores the capability advertisement of **4byte-as**, and brings up the session as a 2byte AS.

[See [disable-4byte-as](#) and [minimum-hold-time](#)]

- **Support for BGP PIC Edge with BGP labeled unicast (MX Series and PTX Series)**—Starting with Junos OS Release 19.4R1, MX Series and PTX Series routers support BGP PIC Edge with BGP labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

[See [Load Balancing for a BGP Session](#).]

- **Decouple RSVP for IGP-TE (MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, a device can advertise selective **traffic-engineering** attributes such as **admin-color** and **maximum-bandwidth**, without enabling RSVP, for segment routing and interior gateway protocol (IGP) deployments.

Services Applications

- **Port Mirroring support (MPC10E line card on MX240 MX480, and MX960 routers)**—Starting in Junos OS Release 19.4R1, Junos OS supports port mirroring on the MPC10E line card for VPLS.

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

- **Programmable DNS error code in response to DNS query (MX240, MX480 and MX960 routers)**—Starting in Junos OS Release 19.4R1, for the DNS queries for blocklisted domains which are of SRV and TXT query types, you can specify a TXT or SRV response code in the DNS response with an empty answer section. To specify the response code, configure the **txt-resp-err-code** or **srv-resp-err-code** option at the **[edit services web-filter profile profile-name dns-filter-template template-name]** hierarchy level. For both the options, if you configure **Noerror** as the value, the error code is sent as **0** with an empty response; whereas, if you set **Refusederror** as the value, the error code is sent as **5**.

[See [DNS Request Filtering for Blacklisted Website Domains](#).]

[dns-filter](#)

Software-Defined Networking

- **Map PCE-initiated P2MP LSPs to MVPN (MX Series)**—Starting in Junos OS Release 19.4R1, you can associate a single or range of MVPN multicast flows (S,G) to a dynamically created PCE-initiated point-to-multipoint label-switched path (LSP). You can specify only selective types of flows, which include a route distinguisher (RD), (S,G) address, and LSP name. When the incoming traffic matches the specified flows, it is mapped to the point-to-multipoint PCE-initiated LSP.

With this feature, you can benefit from reduced configuration as the PCE-initiated point-to-multipoint LSPs are dynamically mapped, thereby eliminating the need to statically enable MVPN and point-to-multipoint LSPs.

[See [Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs](#).]

- **Tunnel templates for PCE-initiated segment routing LSPs (MX Series)**—Starting in Junos OS Release 19.4R1, you can configure a tunnel template for Path Computation Element (PCE)-initiated segment routing LSPs and apply it through policy configuration. These templates enable dynamic creation of segment routing tunnels with two additional parameters – Bidirectional forwarding detection (BFD) and LDP tunneling.

With the support for tunnel configuration, the LSPs that you would configure statically can now be automatically created from the PCE, thereby providing the benefit of reduced configuration on the device.

[See [Understanding Static Segment Routing LSP in MPLS Networks](#).]

Software Licensing

- **Subscriber Access Licensing (MX Series and vMX)**— Starting in Junos OS Release 19.4R1, you need one license per subscriber interface created on subscriber access model.

You need only one license if the DHCP dual stack session running with a single SDB session. To configure the single SDB session, use the **classification-key** option in the `[edit system services dhcp-local-server]` hierarchy .

[See [Subscriber Access Licensing Overview](#) and [classification-key \(DHCP Relay Agent\)](#).]

Subscriber Management and Services

- **Support for GRES and anchor PFE redundancy on Junos Multi-Access User Plane (MX240, MX480, MX960)**—Starting with Junos OS Release 19.4R1, Junos Multi-Access User Plane supports graceful Routing Engine switchover (GRES) and anchor PFE 1:1 hot-standby redundancy to preserve sessions and bearers in the event of failure.

[See [GRES on Junos Multi-Access User Plane](#) and [Anchor PFEs and Redundancy in Junos Multi-Access User Plane](#).]

- **Automatic removal of the redirect service after a one-time redirect (MX Series)**—Starting in Junos OS Release 19.4R1, you can configure the router to redirect a subscriber only once when the subscriber logs in. This enables you to easily provide notifications or advertisements to your subscriber base when subscribers log in. The initial HTTP-GET request from the subscriber triggers the removal of the redirect service. After the temporary redirect to the captive portal, subscribers can reach the specified URL without being redirected for the duration of the session. Automatic removal enables you to avoid using an external policy server, such as a RADIUS CoA message, to remove the redirect service.

[See [How to Automatically Remove the HTTP Redirect Service After the Initial Redirect](#).]

- **Support for charging and usage reports on Junos Multi-Access User Plane (MX240, MX480, MX960)**—Starting with Junos OS Release 19.4R1, Junos Multi-Access User Plane supports volume based Usage Reporting Rules (URRs) in accordance with 3GPP TS 23.203, Policy and charging control architecture.

[See [CUPS Session Creation and Data Flow with Junos Multi-Access User Plane](#).]

- **Junos Multi-Access User Plane (MX240, MX480, MX960)**—With Junos OS Release 19.4R1, we introduce Junos Multi-Access User Plane, a software solution that turns your MX router into a high-capacity user plane function called a System Architecture Evolution Gateway-User Plane (SAEGW-U). This MX SAEGW-U interoperates with a third-party SAEGW-C (control plane function), per 3GPP Release 14 Control User Plane Separation (CUPS) architecture, to provide high-throughput 4G and 5G fixed-wireless access service with support for 5G non-stand-alone (NSA) mode. CUPS enables independent scaling of the user and control planes, network architecture flexibility, operational flexibility, and an easier migration path from 4G to 5G services. The CUPS architecture is optional for 4G but inherent in 5G architecture.

To transform your MX240, MX480, or MX960 router into an SAEGW-U, all you need is at least one MPC7 linecard, a routing engine with at least 32GB memory, and Junos OS Release 19.4R1.

[Junos OS Release 19.4R1 is the first release to support Junos Multi-Access User Plane functionality. We recommend you use this release for lab testing & early field qualification. Full deployment support is available in a later release. Documentation for Junos Multi-Access User Plane is included in the JUNOS Release 19.4R1 documentation here: [Junos Multi-Access User Plane User Guide](#).]

- **Support for Lawful Intercept on Junos Multi-Access User Plane (MX240, MX480, MX960)**—Starting with Junos OS Release 19.4R1, Junos Multi-Access User Plane supports Lawful Intercept in accordance with 3GPP TS 33.107, Lawful interception architecture and functions.

[See [MX Series Router As SAEGW-U](#).]

- **CoA messages support Session-Timeout attribute (MX Series)**—Starting in Junos OS Release 19.4R1, you can apply a session timeout for subscriber sessions with a RADIUS CoA message that includes the Session-Timeout attribute (27). This capability is useful, for example, when subscribers purchase Internet access for a specific period of time and must log out when the session expires. In earlier releases, the router does not recognize the attribute if it is included in a CoA message.

[See [Understanding Session Options for Subscriber Access](#).]

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:
 - Additional information is now logged for MGD-FFP intermodule communication.
 - Commit errors that previously were only shown onscreen are now logged.

We provide a new operational command, **request debug information**, to speed up the initial information-gathering phase of debugging.

[See [request debug information](#).]

SEE ALSO

[What's Changed | 110](#)

[Known Limitations | 116](#)

[Open Issues | 119](#)

[Resolved Issues | 135](#)

[Documentation Updates | 168](#)

[Migration, Upgrade, and Downgrade Instructions | 169](#)

What's Changed

IN THIS SECTION

- [What's Changed in 19.4R2-S4 | 111](#)
- [What's Changed in 19.4R2-S2 Release | 111](#)
- [What's Changed in 19.4R2-S1 Release | 111](#)
- [What's Changed in 19.4R2 Release | 111](#)
- [What's Changed in 19.4R1 Release | 112](#)

Learn about what changed in Junos OS main and maintenance releases for MX Series routers.

What's Changed in 19.4R2-S4

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the `cspf-backoff-time` statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

See [cspf-backoff-time](#).

What's Changed in 19.4R2-S2 Release

General Routing

- **Command to view summary information for resource monitor (MX Series routers and EX9200 line of switches)**—You can use the `show system resource-monitor` command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

[See [show system resource-monitor](#).]

What's Changed in 19.4R2-S1 Release

Routing Protocols

- **Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.

What's Changed in 19.4R2 Release

Class of Service

- We've corrected the output of the `"show class-of-service interface | display xml"` command. Output of the following sort: `<container><leaf 1> data <leaf 2> data <leaf 3> data <leaf 1> data <leaf 2> data <leaf 3> data` will now appear correctly as: `<container> <leaf 1> data <leaf 2> data <leaf 3> data <container> <leaf 1> data <leaf 2> data <leaf 3> data`.

General Routing

- **Precision Time Protocol (PTP) interface configuration (MX2020, MX2010, MX480, MX960, and MX240)**—

Remove the aggregated Ethernet interface association and upgrade the device when configuring PTP interface.

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:
 - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.
 - ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Network Management and Monitoring

- **Enhancement to the show SNMP mib command**—In Junos OS Release 19.4R2, and later, a new option, **hex**, is supported to display the SNMP object values in the hexadecimal format. In earlier releases, the **show snmp mib** command displays the SNMP object values in ASCII and decimal format only.

See [show snmp mib](#)

Platform and Infrastructure

- **NTP Boot Server configuration (MX204, MX960, MX10003, MX10002, MX10016, MX10000, MX480, MX104, MX10008, MX240, MX2010, MXTSR80, MX80, MX2008, MX150, and MX2020)**—Use **set ntp server <address>** command to set the correct time when we boot the router instead of boot-server <address>.

See [Synchronizing and Coordinating Time Distribution Using NTP](#).

Software Installation and Upgrade

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

What's Changed in 19.4R1 Release

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, MX Series, PTX Series, and SRX Series)**—Starting with Junos OS Release 19.4R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

[See [commit \(System\)](#).]

- **Automatic installation of YANG-based CLI for RIFT protocol (MX Series, QFX Series, and vMX with 64-bit and x86-based servers)**—In RIFT 1.2 Release, installation of the CLI for RIFT protocol occurs automatically along with the installation of the junos-rift package. In the pre-1.0 releases of the junos-rift package, the RIFT CLI had to be installed separately using **request system yang** command after installation of the junos-rift package.

Interfaces and Chassis

- **Logical Interface is created along with physical interface by default (MX Series, QFX Series, EX Series)**—Starting in Junos OS Release 19.4R1, logical interfaces are created on ge, et, and xe interfaces along with the physical interface, by default. In earlier Junos OS releases, by default, only physical interfaces are created.

For example, for ge interfaces, previously when you viewed the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

- **Change in error severity (MX960, MX240, MX2020, MX480, MX2008, and MX2010)**—Starting in Junos OS Release 19.4R1, the severity of the CRC errors (XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR) has been reduced from Fatal to Major. Earlier, these errors caused the line card to be reset, if the **interasic-linkerror-recovery-enable** command was configured. Now, these errors will only disable the Packet Forwarding Engines that are affected. With this change, the **interasic-linkerror-recovery-enable** command has no effect in these errors because severity of these errors has been reduced to Major.

NOTE: This behavior change is applicable to the following line cards only: MPC5E, MPC6 MPC7, MPC8, and MPC9.

MPLS

- **Root XML tag change for show rsvp pop-and-forward | display xml command (MX480)**—We have changed the root XML tag for the **show rsvp pop-and-forward | display xml** command to rsvp-pop-and-fwd-information to make it consistent with the XML tag convention. In earlier releases, the command output displays **rsvp-pop-and-fwd-info** XML tag. Update the scripts with the rsvp-pop-and-fwd-info XML tag to reflect the new rsvp-pop-and-fwd-information XML tag.

[See [Junos XML API Explorer - Operational Tags](#).]

Network Management and Monitoring

- **SSHD process authentication logs timestamp (MX Series)**—Starting in Junos OS Release 19.4R1, the SSHD process authentication logs use only the time zone defined in the system time zone. In the earlier releases, the SSHD process authentication logs sometimes used the system time zone and the UTC time zone.

[See [Overview of Junos OS System Log Messages](#).]

Routing Protocols

- **XML RPC equivalent included for the show bgp output-scheduler | display xml rpc CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.4R1, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler](#).]

Services Applications

- **Update to CLI option for configuring the version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03**—In Junos OS Release 19.4R1, the **version-3** option under the **[edit services softwire softwire-concentrator map-e]** hierarchy for configuring the version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 is optional. In the earlier Junos OS releases, if you did not configure the **version-3** option, the configuration resulted in an error.

[See [map-e](#).]

Software-Defined Networking

- **Increase in the maximum value of delegation-cleanup-timeout (MX Series)**—Starting in Junos OS Release 19.4R1, you can configure a maximum of 2,147,483,647 seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

With the increase in maximum value of **delegation-cleanup-timeout** from 600 to 2,147,483,647 seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that might disrupt the PCEP session with the main active stateful PCE.

[See [delegation-cleanup-timeout](#).]

Subscriber Management and Services

- **Enhancement to commands to display reason for Routing Engine disconnect (MX Series)**—Starting in Junos OS Release 19.4R1, several commands display the reason when the master and standby Routing Engines disconnect because of a DRAM size mismatch error. On a chassis with two Routing Engines, this error can result when both of the following are true:
 - The Routing Engines have different amounts of DRAM.
 - A 64-bit Junos OS image is loaded on the chassis.

You can avoid this problem by doing either of the following:

- Ensure that both Routing Engines have the same amount of DRAM.
- Load a 32-bit image.

[See [show system subscriber-management summary](#), [show database-replication summary](#), [request chassis routing-engine master](#), and [show chassis routing-engine](#).]

- **Prevent queue-based throttling from stopping subscriber login (MX Series)**—Starting in Junos OS Release 19.4R1, you can specify a value of 0 with the **high-cos-queue-threshold** statement. This value prevents any subscriber from being throttled by queue-based throttling.

[See [Throttling Subscriber Load Based on CoS Resource Capacity](#).]

System Logging

- **Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

What's New 87
Known Limitations 116
Open Issues 119
Resolved Issues 135
Documentation Updates 168
Migration, Upgrade, and Downgrade Instructions 169

Known Limitations

IN THIS SECTION

- General Routing | 116
- Interfaces and Chassis | 118
- Platform and Infrastructure | 119
- Routing Protocols | 119

Learn about known limitations in this release for MX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- PTX Series platform with FPC Model FPC-PTX-P1-A or FPC2-PTX-P1A might encounter single event upset (SEU) event that can cause a linked-list corruption of the TQCHIP. The following syslog message will get reported: **Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002.** Junos OS chassis management error handling does detects such a condition, raises an alarm, and performs **disable-pfe** action for the affected Packet Forwarding Engine entity. [PR1254415](#)
- First packet pertaining to J-Flow Packet Forwarding Engine sensor in UDP mode is missing after line card reboot on NFX150 platform. [PR1344755](#)
- Traffic on GRE interface both ingress and egress cannot be Layer 2 mirrored. [PR1462375](#)

- Applying and removal of 1Gbps speed results in channel being down. [PR1456105](#)
- The control peer PFCP heartbeat request timeout window must be greater than 90 seconds. [PR1459135](#)
- Load balancing does not work as expected when tested with NAPT44 case twice. [PR1477670](#)
- aftd hogged on executing clear VPLS table and MACs are not learned for less than 5 minutes. [PR1473334](#)
- If MTU is configured to a value higher than 9500, which is the maximum permissible value, configuration is done successfully. However, the actual value is set back to 1518 without any error. [PR1372690](#)
- The MIC-MACSEC-20G supports 10-Gigabit speed through the **set chassis fpc x pic y pic-mode 10G** configuration that is applied to both the PICs in that MIC. Any other PIC mode configuration must be removed and the 10-Gigabit PIC mode configuration must be applied. [PR1374680](#)
- In USF and non-USF cases, the monitor interface is MS (or) VMS interface. When chassisd restarts, all FPCs restart. SRD also restarts and ICCP connection goes down. If the FPC hosting the ICL goes down first before SRD receives physical interface IFD down for the monitored interface, the switchover does not happen immediately. [PR1416064](#)
- JSD generates core files when aggressively subscribing and unsubscribing both gRPC and gNMI subscriptions from multiple sessions. [PR1433744](#)
- The SPC3 cards are not supported with RE-2000. Even if the RE-2000 is the backup Routing Engine. [PR1435790](#)
- Traffic error does not get policed after it is locally switched for VLAN 100 and 101, while verifying the selective local-switching functionality with 4000 VLANs. [PR1436343](#)
- Routing Engine interprets any input from the console port as interrupts. Depending on the frequency, console noise impacts the Routing Engine interrupt handling to different extents even with the current throttling mechanism. When the interrupt frequency is too high for the Routing Engine to handle, the impact might vary from line-card reboot (partial impact) to Routing Engine reboot (chassis-wide impact). [PR1436386](#)
- In a large-scale setup such as large number of routing-instances or interfaces, if there are frequent changes in the configuration and interface flapping when the rpd restarts through deactivate or activate of logical-system or restart routing, the rpd might crash. [PR1438049](#)
- Whenever the primary path goes down for the SRTE-tunnel, dynamic tunnel module (DTM) starts an expiry timer of 15 minutes. If the primary path comes up within this timer period, the tunnel comes up again. After the timer expires and the primary path is still not up, DTM asks SR-TE to remove the tunnel. Also, if there are multiple paths to reach the tunnel endpoint, BGP routes resolve over the other route, for example L-ISIS path. Later, even if the primary path comes up, BGP routes remain resolved over the other secondary route and do not change. No re-resolution happens because the SRTE-tunnel is being resolved with more than one indirection. For example, SR-TE over MPLS over IS-IS in this case. The same issue occurs in RSVP tunnels. The issue is applicable to uncolored tunnels only. [PR1439557](#)
- Sampling applications like **port-mirror** and **inline-jflow** are not supported on VPLS tunnel interfaces in ingress direction where ingress packets are sent to the IRB interface for routing. Configuration of sampling

application on VPLS tunnel interfaces in such scenario causes packet to drop in ingress direction.

[PR1444849](#)

- If Sx Modification-Request has an Update FAR Apply Action that has the DUPL and DROP bits set, the traffic is dropped as expected. However, the packets are not duplicated to the SX3LIF/MD. This happens for both upstream and downstream traffic. [PR1450859](#)
- When 32000 inetcolor and 32000 inet6color are programmed together, the jsd process is hit. [PR1452464](#)
- In a scaled scenario where the Routing Engine pushes a lot of routes to the Packet Forwarding Engine in the presence of the dynamic tunnel configuration, FIB convergence might take more time, leading to traffic drops. [PR1454817](#)
- Member of It interface of a rlt interface must have same bandwidth configured. Bandwidth mismatch might lead to unexpected behavior. Changes to It or rlt interfaces must not be done if a ps interface is anchored over these tunnel interfaces. [PR1458951](#)
- The It interface Scheduler remains in the invalid state under egress IFD list after changing the It tunnel to a different Packet Forwarding Engine. [PR1458955](#)
- Changes to rlt interface with ps anchored over is not recommended. For more information, refer to: [anchor-point \(Pseudowire Subscriber Interfaces\)](#). [PR1460898](#), [PR1460910](#)
- The traffic on GRE interface on both ingress and egress cannot be Layer 2 mirrored. [PR1462375](#)
- In AFT-based platforms, all Layer 2 inject packets will hit egress feature list and filters/policer are also part of it. A couple of scenarios are listed here: 1. Having reject as terminating action in the firewall will respond back with ICMP with a specific reject code. Because this acts as L3 inject for AFT, these packets are subjected to egress feature list (example: egress filter, policers, and CoS configuration). So, this is causing excess packet count in egress filters. 2. Scenario 1 also applies to routing-instance terminating action through which destination IP is not reachable. While configuring the filters these scenarios must be considered so as to classify these packets not to hit the policers/filter counters if required. This is not an issue, but it is a "Feature As Designed" in AFT-based platforms. [PR1477638](#)

Interfaces and Chassis

- In a large-scale subscriber environment, changing aggregated Ethernet member link configuration might cause two Routing Engines to generate core files. [PR1375638](#)
- When you use centralized mode for VRRP and if there are scaled VRRP instances, when the VRRP master side fails, such as ungraceful Routing Engine switchover, the traffic might drop for a short time. [PR1451704](#)

Platform and Infrastructure

- On all Junos OS platforms, execution of Python scripts through enhanced automation does not work on veriexec images. [PR1334425](#)

Routing Protocols

- Three BGP replication flaps are seen on a new master Routing Engine after GRES. The route synchronization issue is also seen between Routing Engines without GRES. [PR1441925](#)
- When you scale RIB to 80 million after FPC restarts, it is not able to scale on the backup Routing Engine. [PR1444073](#)

SEE ALSO

What's New 87
What's Changed 110
Open Issues 119
Resolved Issues 135
Documentation Updates 168
Migration, Upgrade, and Downgrade Instructions 169

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 120](#)
- [EVPN | 120](#)
- [Forwarding and Sampling | 121](#)
- [General Routing | 121](#)
- [Interfaces and Chassis | 128](#)
- [Layer 2 Ethernet Services | 130](#)
- [Layer 2 Features | 130](#)
- [MPLS | 130](#)
- [Network Management and Monitoring | 131](#)

- Platform and Infrastructure | 131
- Routing Protocols | 133
- Services Applications | 134
- VPNs | 134

Learn about open issues in this release for MX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- The tag changes are intentionally added. [PR1475179](#)

EVPN

- Duplicate packets in EVPN scenario are seen because a nondesignated forwarder is sending an inclusive multicast packet to the PE-CE interface after MAC lookup. [PR1245316](#)
- With Junos OS Release 19.3R1, VXLAN OAM host-bound packets are not throttled with DDoS policers. [PR1435228](#)
- When DHCP is used with EVPN, the Layer 2 learning daemon adds a destination route to the kernel with a "permanent remote" flag while the DHCP process adds a destination route with a "permanent remote" flag. There could be a race condition where the Layer 2 learning destination route is overwritten by the DHCP route, causing the remote flag to get deleted. This subsequently leads to the ARP route to age out in the kernel. To ensure that DHCP routes are not added to the kernel, you must configure the **forward-only** statement under **forwarding-options dhcp-relay**. [PR1439568](#)
- EVPN-VXLAN core isolation does not work when the system reboots or the routing restarts. [PR1461795](#)
- When dynamic list nexthop is referenced by more than 1 route it could result in an early delete of the next hop from kernel, thereby observing nhindex as 0 ("Next hop type: Dynamic List, Next hop index: 0" in **show route** command). As a workaround, restart routing. [PR1477140](#)
- In EVPN active-active mode, the Layer 2 Address Learning Daemon (l2ald) memory might slowly come up when the local multi-homing CE interfaces continuously flap. If the memory of l2ald is exhausted, it will cause the l2ald to crash. [PR1498023](#)

Forwarding and Sampling

- If IPv4 prefix is added to a prefix-list referred by IPv6 firewall filter, the following log message is not seen: **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes, Match [from prefix-list Block-Host] might be optimized.** [PR1395923](#)
- Restart the firewall process in both Routing Engines when firewall error logs are noticed along with **SSD hardware failure** logs. [PR1397171](#)
- After you restart routing, the remote mask, which indicates from which remote PE devices MAC IP addresses are learned, that the routing daemon sends might be different from the existing remote mask that the Layer 2 learning daemon had prior to restart. This causes a mismatch between Layer 2 learning and routing daemons interpretation as to where the MAC IP address entries are learned, either local or remote, leading to the MAP IP table being out of synchronization. [PR1452990](#)
- Verify event CP down/up long enough to trigger EP timeout for CoS hierarchy model 2 , failing as expected DHCP subscribers are not bound. [PR1505409](#)

General Routing

- On an MX104 platform, when using "snmpbulkget" or "snmpbulkwalk" (for example, used by the SNMP server) on chassisd related component (for example, jnxOperatingEntry), a hardware limitation might result in high CPU usage and slow response by the chassisd process, which might also lead to query timeout on the SNMP client. In addition, the issue might not be seen when using SNMP query for interface statistics. [PR1103870](#)
- SIP session fails when the IPv4 SIP client in public network initiates SIP call with the IPv6 SIP client in the private network. [PR1139008](#)
- Configuration committed by a logical system user will not be propagated to backup Routing Engine. As a workaround, if **set system commit no-delta-synchronize** is configured, it gets synchronized. [PR1160759](#)
- On dual Routing Engines with GRES enabled, after performing GRES, if the configuration synchronization on the backup Routing Engine fails when it becomes the new master Routing Engine, then in rare conditions, some interfaces cannot be deleted or configuration changes cannot be committed. [PR1179324](#)
- On a vMX platform, performance of the Intel X710 NIC is lower compared to the performance of Intel 82599 NIC. [PR1281366](#)
- If vmhost snapshot is taken on alternate disk and there is no further vmhost software image upgrade, the expectation is that if the current vmhost image get corrupted, the system boots with the alternate disk so the user can recover the primary disk to restore the state. However, under the condition where corruption is with the host root file system, the node is booting with previous vmhost software instead of booting from the alternate disk. [PR1281554](#)
- Chassisd might report a read back error for MPC 3D 16x10G FPCs when it tries to access 0x3a register. The syslog error has no impact and can be safely ignored. [PR1299594](#)

- In some scenarios with MPC, the following major alarm and following messages are generated: **messages log: fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) alarmd[3158]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC x Major Errors Major alarm set, FPC x Major Errors fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) cli> show chassis alarms 1 alarms currently active Alarm time Class Description 2019-01-25 15:18:03 UTC Major FPC x Major Errors.** Despite the major alarm set, this error is due to the Unknown Error Address logged in hardware to the DQ underrun. This message is harmless and has no service impact. [PR1303489](#)
- The chain-composite statement does not bring in a lot of gain, because TCNH is based on ingress rewrite premise. [PR1318984](#)
- In Message Queue Telemetry Transport (MQTT) scenario, the memory leakage (about 4000 memory leakage every 30 seconds) might be seen. However, on very long runs, this uses up high memory, which can indirectly impact other daemons running. [PR1324531](#)
- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections. Therefore, the reactions to failure situations might not be handled in a graceful way. For example, TCP connection timeout because of jlock hog crossing boundary value (5 seconds) might cause bad consequences in MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling marker infrastructure in a MX Series Virtual Chassis setup. [PR1332765](#)
- In some cases, online insertion and removal (OIR) of a MIC on an FPC might lead to traffic destined to the FPC to be silently dropped or discarded. The only way to recover from this is to restart the FPC. The issue will not be seen if you use the corresponding CLI commands to offline and then online the MIC. [PR1350103](#)
- For configurations of bridging routing instances with aggregated Ethernet logical interfaces (6400 logical interfaces) and IRB instances, all from a single FPC, the CPU utilization of the FPC stays at 100 percent for 4 minutes. The behavior from PFEMAN of the FPC has the processing time spiked on IF IPCs, and this seems to be the case of MPC7E from Junos OS Release 16.1R1 (or even earlier). After 4 minutes, the CPU utilization comes down and the FPC is normal. Therefore, scaled configurations on MPC7E takes settling time of more than 4 minutes. [PR1359286](#)
- In rare circumstances, a faulty SFP transceiver installed in an MX104 might cause the AFEB to go offline. The backup Routing Engine and fan tray might also show an alarm. [PR1360426](#)
- If any of the log messages continue to appear in the MPC console, it indicates that the presence of a faulty SFP/SFP+ transceiver is causing I2C transaction from the main board CPU. There is no software recovery available to recover from this situation. These logs also indicate potential Layer 2 circuit transaction failure with any of the 10 ports available with GMIC2 in PIC 0, resulting in an unexpected behavior. For example, links do not come up or the MIC does not boot when restarted. **I2C Failed device: group 0xa0 address 0x70 Failed to enable PCA9548(0x70):grp(0xa0)->channel(0) mic_sfp_select_link:MIC(0/0) - Failed to enable PCA9548 channel, PCA9548 unit:0, channel ID: 0, SFP link: 0 mic_sfp_id_read: Failed to select link 0.** As a workaround, detect and replace the faulty SFP/SFP+ transceiver plugged into the GMIC2 ports. [PR1375674](#)

- On an EX9208 switch, a few xe-interfaces are going down with the error message **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error**. [PR1377840](#)
- In a subscriber management environment, multiple error messages **shmlog: argcnt 309 not enough memory** can be generated every hour. These messages are harmless, and there is no service impact. [PR1384371](#)
- In a subscriber management scenario, the bbe-smgd process might crash when two subscribers log-in with the same framed-route prefix and preference values returned from RADIUS. [PR1387690](#)
- The virtio throughput remains the same for multi-queue and single-queue deployments. [PR1389338](#)
- On MX Series platforms, the DHCP/PPP subscribers might fail to bind. The reason is that when installing new software images, the shared memory (created by the previously running image) might not be cleared out. The issue will persist until the previous values in shared memory are removed, and the daemons affected by the data in shared memory might continue to crash and thus not be able to function properly. [PR1396470](#)
- On MX10003 platform, after removing the FPC from a slot, when a new FPC is plugged in, not only does the chassis show the old serial number for this new FPC, but the entire FPC ID EEPROM data is retained. All the fields show old values. [PR1409930](#)
- The MX104 router has the following limitations in error management: The show chassis fpc error CLI command is not available for MX104 in the Junos OS Releases 13.3R7, 15.1R2, 14.1R5, 14.2R4, 13.3R8, and later. Junos OS does not initiate restart of the system on encountering a fatal error. Although you can configure the action to disable Packet Forwarding Engine for major errors, Junos OS does not disable its only Packet Forwarding Engine on encountering a major error. [PR1413314](#)
- The **show services hybrid-access sessions**, **show services hybrid-access statistics**, and **show services hybrid-access tunnels** commands display values of zero for hybrid access gateway traffic statistics even when traffic is active in the gateway sessions and the tunnels. [PR1419529](#)
- A new statement **nh-detail** is added in **show route detail** and **show route extensive** that prints next hop detail in extensive format with proper and formatted XML tags. [PR1424930](#)
- Dynamic tunnel summary is displaying the incorrect count of up and total tunnels after multiple iterations of activating and deactivating the dynamic tunnel configuration. [PR1429949](#)
- On MX Series platforms, if the clock frequency is slowly changing on CB0 (slow drift), the clock source for MPC-3D-16XGE-SFPP might not be changed to CB1, which might cause interfaces on it to go down and remain in the down state. [PR1433948](#)
- In gRIBI, programmed routes references a next-hop group ID, which in turn points to one or more next-hop IDs. Each next-hop ID contains details of the actual next-hop. Next-hop group ID and next-hop ID are mapped to an IPv6 prefix (for example, FC01::<GRP-ID>). In the case of an IPv4 indirect next hop, gRIBI needs to resolve IPv6 through IPv4 next hop over three levels of indirection. Junos OS does not support IPv6 over IPv4 multilevel next-hop resolution. Therefore, gRIBI cannot resolve next-hop GRPID FC01::<grp-id> and next hop ID <FC01> through an actual indirect IPv4 gateway address. This is a Junos OS limitation. [PR1434050](#)

- On dual Routing Engine MX Series platforms with subscriber management scenario, the repd process might crash after booting for the first time with a newly installed Junos OS release. The replication daemon (repd) is a process to synchronize subscriber information across Routing Engines. The repd crash has no impact on the live service. [PR1434363](#)
- On MPC10E 3D MRATE-15xQSFP the L2 over GRE is not supported. Though the configuration get committed, the feature will not work. [PR1435855](#)
- On routers running Junos OS and serving as EVPN gateways, FPC core files available at heap_block_log due to NULL entries are also seen in the ifbd level list, which are typically added for flush list. This occurs because of the relink logic failure flush logic for MACs when there is ifbd/bd delete. [PR1441824](#)
- FPC19(MPC9E) might not go offline due to unreachable destinations in phase 3 stage. [PR1443803](#)
- Routing Engine generated jumbo frames might get dropped due to incorrect MTU setting on the internal switch. [PR1444963](#)
- Push label is missing in **show route** command output for colored tunnels. It does not have any impact on functionality. [PR1447900](#)
- The **show ddos-protection protocols arp statistics |display xml** command does not show APR violation packets and is also not incremented. [PR1449968](#)
- No functional impact seen because of the error message. The message comes during loading of RLI config (triggers are route addition/deletion). [PR1451213](#)
- OIDs-related service-set module might not work when the following command is performed: **show snmp mib walk enterprises.2636.3.32.1.3.1.4** **show snmp mib walk enterprises.2636.3.32.1.3.1.6** **show snmp mib walk enterprises.2636.3.32.1.3.1.8**. This is expected behavior. **show snmp mib walk 1.3.6.1.4.1.2636.3.32** or **show snmp mib walk jnxSpSvcSet** OID access (one time good enough) results in creating the service-set SNMP database needed. Once **show snmp mib walk 1.3.6.1.4.1.2636.3.32** or **show snmp mib walk jnxSpSvcSet** is accessed, above OIDs might successfully return proper values. [PR1452153](#)
- This is a cosmetic issue that affects only the CLI. RADIUS, L2tp etc. are unaffected. CLI issue is seen after ANCP restart and before ANCP neighbor is re-established and port-ups are received. Under normal working conditions, after ANCP restart, the port-ups should be received right away and the CLI issue will be never seen. [PR1453837](#)
- The behavior has been modified to display the correct protocol number instead of 255 whenever an unknown protocol number is encountered. [PR1454792](#)
- IPv6 accounting stop attributes are not correct for MLPPP subscribers. [PR1455175](#)
- If the command **persist-groups-inheritance** is enabled, when a user executes a **delete** operation (for example, to delete the entire configuration) and selects no, later when the user tries to commit the configuration changes related to groups, multiple daemons might crash. [PR1455960](#)
- With logical-system configuration, filter-based GRE encapsulation is not working. [PR1456762](#)

- After more than 2 million multicast subscribers are activated without performing GRES or bbe-smgd restart, further multicast subscribers might be unable to log in. [PR1458419](#)
- With the scale filter-based-forwarding (FBF) configuration, two FBF instances seem to be unable to forward the traffic to the respective routing instance. It appears that the FBF programming is incorrect for these two FBF instances. [PR1459340](#)
- Occasional warning messages such as **TCP connect error** can be seen during FPC reboot. These are generally inconsequential and have no impact on the FPC or the line-card software functionality. [PR1460153](#)
- NAT performance is impacted with remote syslog enabled. [PR1460211](#)
- The **show dynamic-tunnel database** Junos OS CLI command output did not filter IP-IP tunnels based on user-provided destination address/range. [PR1461659](#)
- A BFD session might flap when it is moving to aggressive interval after coming up with slow/nonaggressive interval. This issue is mainly seen in a scaled setup. [PR1462775](#), [PR1465285](#)
- Backport jemalloc profiling CLI support to all releases where jemalloc is present. [PR1463368](#)
- The traffic stops when volume quota is reached but resumes wrongly after APFE failover. Threshold and quota values are not updated to the secondary APFE. If quota is hit on primary APFE and traffic starts dropping due to quota and switchover happens, traffic will continue to flow until quota is hit. [PR1463723](#)
- The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline/online scenarios: [Oct 3 08:48:35.836 LOG: Err] ifl ps240.1 (1712): child ifl lt-1/0/0.32767 (7709) already there [Oct 3 08:48:35.836 LOG: Err] IFRT: 'Aggregate interface ifl add req' (opcode 87) failed [Oct 3 08:48:35.836 LOG: Err] ifl 1712, child ifl 7709; agg add failed. [PR1464524](#)
- Commit script does not apply changes in private mode unless a commit full is performed. [PR1465171](#)
- The following syslog error messages are harmless and expected during FPC offline/restart scenarios with PS-RLT (with/without link protection) configuration: Nov 12 15:02:00 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x0 delete failed for ifl lt-3/0/0.32767 with err=2 Nov 12 15:02:00 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x1 delete failed for ifl lt-3/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x1 delete failed for ifl lt-5/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x0 delete failed for ifl lt-5/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_lp_handle_event: LP event = 6, child lt-5/0/0 err = 22. The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline. **Online scenarios:** Nov 12 15:08:37 cleansing fpc3 user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIfName:ps1.0 memberIfName:lt-3/0/0.32767 Nov 12 15:08:37 cleansing fpc3 user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIfName:ps1.0 memberIfName:lt-5/0/0.32767. [PR1466531](#)
- 100 percent traffic drop is observed on MX2010 with BGP rib-sharding and update-threading in BGP Layer 3 VPN streams post configuration by removal/restoration. [PR1469873](#)

- If redundant APFEs simultaneously fail/reboot while sessions are bound, inconsistencies could occur between the APFEs. This inconsistency can in rare situations lead to an rmpsd core on the backup Routing Engine following additional subsequent APFE failovers. [PR1471580](#)
- On multicore next-generation Routing Engines on MX240/MX480/MX960 platforms with Universal Services Framework (USF) mode enabled and USF-based services configuration, subsequent Junos VMhost upgrade might fail with an error. The reason is that the next-generation Routing Engine never supported verification of optional packages, but as a feature, USF requires such verification as a necessity because USF is added as an optional package (feature). USF is a common services architecture for router-based services and the security services. All the services/plugins from mspmand will be ported onto flowd to achieve this functionality. [PR1472287](#)
- When L2-Bridge domain is configured with all four channelized interfaces into different bridge domains of different VLAN IDs. Now sending traffic to one of the VLAN IDs, traffic flows correctly but statistics are sometimes updated incorrectly on the other three channelized interfaces. [PR1472464](#)
- An occasional aftd core file is observed on setups with GNF configuration after rebooting another line card such as MPC8 within the same GNF. [PR1473403](#)
- When both MSTP and ERP are enabled on the same interface, then ERP will not come up properly. [PR1473610](#)
- IS-IS and micro BFD sessions are not coming up during baseline for MPC10E card. [PR1474146](#)
- On rebooting the external server, the SNMP values configured within `/etc/snmp/snmpd.conf` at the server get overwritten with content from the JDM SNMP stanza. Trap configuration changes get completely removed. JDM restart or stop/start does not change the host `/etc/snmp/snmpd.conf` file, only system reboot of the server does. [PR1474349](#)
- Changing framing modes on a CHE1T1 MIC between E1 and T1 on a MPC3E NG HQoS line card will cause the PIC to go offline. [PR1474449](#)
- During -ve scenario like clear MAC table, sometimes the expected number of 512K MAC addresses are not relearned. [PR1475205](#)
- Error messages **[Error] L2alm : l2alm_mac_process_hal_delete_msg:667 Ignoring MAC delete with ifl index 355, fwd_entry has 7888** are seen after performing configuration removal/restore with IPMPLS configurations in the MX480. [PR1475785](#)
- In VPLS configurations, ARP resolution over an IRB interface could fail if the hosts are behind a vt-tunnel. In such scenario, as a workaround, the **protocols vpls no-tunnel-services** command could be used. [PR1477005](#)
- With Junos OS Release 19.3R2 and later, syslogs such as **spmb0 cmtty_sfb_temp_check: sfb[0] is powered OFF** and **spmb0 cmtty_sfb_voltage_check_one: sfb[0] is powered OFF** are flooding even though SFB2s are online in MX2008/MX2010/MX2020. [PR1477924](#)
- The traffic loss is seen for 10 seconds when switching from secondary to primary path (even with disabling sbfd configuration). [PR1478299](#)

- The next-generation MX-SPC3 services card can exhibit inconsistent behavior when the vmhost image is installed on the next-generation Routing Engine (NG-RE): RE-S-X6-64G-UB. To use the next-generation Routing Engine, a workaround is available. See the "Problem" and "Workaround" sections for more detail. Other Routing Engines that are compatible with next-generation services do not experience this problem. These Routing Engine are RE-S-1800X4-16G-UPG-BB, RE-S-1800X4-32G-UB, RE-S-1800X4-16G-UPG-BB, and RE-S-1800X4-32G-UB. [PR1482334](#)
- After multiple login/logout (create/delete) of sessions, there can be a few sessions stuck in INVALID state. [PR1484855](#)
- On MX Series platforms with MS-MPC/MS-MIC, if there are lot of self-generated packets like TCP-tickle and UDP-logging, there might be data congestion on the data path due to no throttling functionality for such types of packets. When the data path is blocked, prolonged flow-control might happen with the service interfaces being brought down and the PIC reboot. At the same time, the mspmand core file will be generated if **dump-on-flow-control** is enabled. [PR1489942](#)
- Support for upgrading PSMs firmware on the MX2000 line of devices. [PR1489939](#), [PR1489967](#)
- User configured MTU might be ignored after upgrading the unified ISSU using **request vmhost software in-service-upgrade**. [PR1491970](#)
- The delta PSM Firmware upgrade status is incorrectly displayed. [PR1493045](#)
- Missing firmware image file in **usr/share/pfe/firmware**. [PR1494557](#)
- In node slicing setup with MPC5 or newer cards after GRES, RADIUS interim updates might not carry the actual statistics. [PR1494637](#)
- On an MX Series platform in a node slicing setup, if subscriber services and targeted distribution feature are enabled on an aggregated Ethernet interface, when a FPC with aggregated Ethernet member ports comes online, subscribers might be disconnected. [PR1498024](#)
- **no-normalization** configuration statement is enabled under routing-instances stanza for a routing-instance. The **commit check** fails after adding the logical interface with any of the following configuration statements into the routing-instance: **vlan-id all**, **vlan-id inner-all**, **input-vlan-map**, **output-vlan-map**, **vlan-id-list**, **vlan-id-range**, **vlan-tags inner-list**, **vlan-tags inner-range**. [PR1499265](#)
- On MX2020 and MX2010, the **pem_tiny_power_remaining** message will be continuously logged in chassisd log. [PR1501108](#)
- After rollback the VLAN-ID under EVPN is observing traffic drop. [PR1505849](#)
- OPENSSL_KATS core file is found at **fips_run_kats**. [PR1501849](#)
- EIM and APP mapping are stuck after session are timed out gracefully. [PR1501995](#)
- On deactivating or activating routing instance, packet from non existing source on GRE/UDP designated tunnel are accepted where they are supposed to be dropped. [PR1503421](#)

- **show bridge statistics** might not display the statistics information for pseudowire subscriber interfaces. [PR1504409](#)
- This issue is related to CPCD service on spc3 card. Http-get request is not redirect to even after service is attached to subscriber. [PR1505438](#)

Interfaces and Chassis

- After GRES, the 1-Gigabit Ethernet changes to 10-Gigabit Ethernet. [PR1326316](#)
- Some routers index the SFP transceivers starting at 1, while interface numbering starts from 0; thus, reading the Packet Forwarding Engine-level output can be confusing. [PR1412040](#)
- L2IFL configuration is now decoupled from bridge / EVPN configuration. A L2IFL can now be configured without being assigned to a bridge / EVPN. [PR1438172](#)
- When user checks for interface level statistics, issue is coming for IPv6 counters. at the originating router IPv6 local stats counters not updating due to that IPv6 local stats counters not incrementing. IPv6 transit stats are derived from Total stats and Local stats (Transit stats = Total - Local). As the Local stats are not updating, total stats and transit stats will be the same. And this issue specific to MPC10E platforms. [PR1467236](#)
- MPC10 line card runs on newer version of FPC software. Currently convergence number for MPC10 is not at par with legacy MPC cards for high scale. Following are few recommendation to achieve better convergence numbers with VRRP on MPC10 card. Configure failover delay to 5 seconds. This will help in quick action without impacting RE CPU usage. set protocols vrrp failover-delay 5000 If system has inherit sessions configured then its recommended to reduce inherit advertisement interval timer to 6 seconds. Default of inherit is 120s which causes slower convergence for inherit sessions. set protocols vrrp inherit-advertisement-interval 6 With above two knobs configured worst case convergence for a scale of 8000 can be expected to be at 38 seconds. [PR1474656](#)
- When FPC is restarted all vrrp transmit sessions anchored on that fpc gets redistributed to other available fpc. Tx gets disrupted during this time causing flap at peer end. Because peer router takes mastership traffic gets redirected to peer for some time. But when transmit sessions are up again on the DUT peer router moves to backup. For active vrrp sessions traffic revert back to original master quickly. Its primarily because advertisement interval of active vrrp session is 1 second. But traffic for inherit session does not revert back to original router quickly. This is because advertisement interval of inherit session is 120s. Even though peer router has moved to backup intermediate switch still point towards peer for VMAC. It changes only after getting packet from original master. This may take up to 120s for inherit sessions hence causing blackhole for 120+ seconds. This issue is present for non MPC10E line cards as well. Similar traffic drops can be seen when fpc4 is restarted. This can be avoided by reducing advertisement interval of inherit sessions with following command set protocols vrrp inherit-advertisement-interval 6 After configuring this worst case loss were observed to be around 8.5 seconds. [PR1474694](#)
- MPC10 line card runs on newer version of FPC software. Currently convergence number for MPC10 is not at par with legacy MPC cards for high scale. Following are few recommendation to achieve better

convergence numbers with VRRP on MPC10 card. Configure failover delay to 5 seconds. This will help in quick action without impacting RE CPU usage. set protocols vrrp failover-delay 5000 If system has inherit sessions configured then its recommended to reduce inherit advertisement interval timer to 6 seconds. Default of inherit is 120s which causes slower convergence for inherit sessions. set protocols vrrp inherit-advertisement-interval 6 With above two knobs configured worst case convergence for a scale of 8000 can be expected to be at 38 seconds. For 19.4r1.5 - with suggested knob track interface ,local interface enabling back for dut to restore mastership time is 38.2sec - track interface, remote side interface enabling back for dut to restore mastership time is 38.3sec For 19.4r1.5 - ,with suggested knob track route ,local interface enabling back for dut to restore mastership time is 14.5 sec - track route, remote side interface enabling back for dut to restore mastership time is 15 sec [PR1475140](#)

- Some of the Demux Vlan over AE configured statically(from CLI configuration) are not programmed with child legs. All the traffic on these IFLs will be dropped. [PR1476465](#)
- On all Junos OS platforms with logical-systems setup, a local IP address might not be cleaned properly and become stale after a specific order of configuration changes as follows. set logical-systems JNPR interfaces ge-0/0/0 unit 0 family inet address 10.10.10.1/24 commit deactivate logical-systems JNPR interfaces ge-0/0/0 unit 0 family inet set interfaces ge-0/0/0 unit 0 family inet address 10.10.10.10/24 commit. [PR1477084](#)
- When priority is increased for all 4k vrrp sessions some of the vrrp sessions doesnt take over mastership. Peer router continues to become master for those sessions. There wont be any traffic traffic impact as one router is still master. Possible workaround will be to deactivate and activate vrrp sessions in backup state. [PR1478349](#)
- dual-stack subscribers login, but delayed/incomplete dhcpv6 binding causes errors with DCD and **DCD_CONFIG_WRITE_FAILED** errors which brings down all subscribers and cannot rebind as DCD believes these are still present. [PR1478215](#)
- PPPoE subscribers are not UP while verifying static V4 subscriber in passive mode. [PR1483395](#)
- Input and output bytes count mismatch in IPv6 traffic stats while issuing show interface extensive. [PR1505100](#)

Layer 2 Ethernet Services

- After performing febr restart, jdhcpd process might generate a core file on the other PE device in multi-home topology. [PR1504365](#)

Layer 2 Features

- When "input-vlan-map" with a push operation is enabled for dual-tagged interfaces in "enhanced-IP" mode, there is a probability that the broadcast, unknown unicast, and multicast (BUM) traffic might be dropped or discarded silently on some of the child interfaces of the egress ae interfaces or on some of the equal-cost multipath (ECMP) core links. [PR1078617](#)

MPLS

- In Resource Reservation Protocol (RSVP) Label Switched Path (LSP) with loose or undefined path scenario, the LSP might stay in down state due to loop detection after the link in the path flaps. [PR1384929](#)
- With ldp-tunneling over RSVP LSP where RSVP LSP has link-protection, LDP route might flap when the interface on the bypass is brought down. [PR1450516](#)
- Multiple non-stop attempts to clear IGP database can result in routing daemon core when locally computed SRTE LSPs are configured in the order of thousands. [PR1456212](#)
- While setting LSP BW constraint and signalling LSP the remaining BW can be less than expected like in the example below not allowing further reservation through this link **user@router> show rsvp interface**
RSVP interface: 2 active Active Subscr- Static Available Reserved Highwater Interface State resv iption
BW BW BW mark et-0/1/0.0 Up 1 90% 100Gbps 4.99999Gbps 85Gbps 90Gbps <<<<<< So if you try
to signal a new LSP with BW 5G through this link it will fail. [PR1458527](#)
- On all Junos OS platforms with distributed CSPF under Segment Routing Traffic Engineering (SR-TE) scenario, if SR protocols are deactivated and then activated, restart routing, etc, the rpd process might crash. [PR1493721](#)
- The rpd process might crash generating a core when SNMP polling is done using oid jnxMplsTeP2mpTunnelDestTable. [PR1497641](#)
- Flood next-hop ID is not same in both master and backup Routing Engine. [PR1495925](#)
- Label is deleted from MPLS table for MVPN generated sessions after switchover phase of unified ISSU leading to traffic loss in MVPN stream. [PR1500615](#)
- On all Junos OS platforms running with the Resource Reservation Protocol (RSVP) configured, when some socket error occurs, RSVP might be unable to send message. However, the counter for sent messages has already been increased due to the software design defect. The incorrect counter might lead to the rpd process crash. It is a rare timing issue and it might be more likely to happen in a scaled scenario. [PR1505834](#)

iso.3.6.1.2.1.80.1.3.1.7 -> pingResultsProbeResponses iso.3.6.1.2.1.80.1.3.1.9 -> pingResultsRttSumOfSquares. [PR1458983](#)

- When traffic is received from 1000 different vrf instance on PF from CE, then few flows (4 to 5) are dropped at PE. That is the reason why 0.1 to 0.9 percent drop has seen. The reason for the drop is identified. [PR1460471](#)
- After performing a unified ISSU with scaled configuration, high CPU utilization is observe in crash MPC 16x10GE card . [PR1461715](#)
- In NTP with the boot-server scenario, when the router or switch boots, the NTP daemon sends a ntpdate request to poll the configured NTP boot-server to determine the local date and time. If the ntpdate is not activated correctly while the device booting, the ntpdate might not work successfully. Then some cosmetic error messages of time synchronization might be seen, but there is no impact with time update since ntp daemon will update the time eventually. [PR1463622](#)
- Line card errors found at **HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invlaid IRB topo/ IRB ifl zero in l2 nh 40495 add IRB**. [PR1472222](#)
- Line card Errors found at **RT-HAL,rt_mesh_group_delete_check,1599: Deletion of a non-existant mesh-group : proto 35 rtt 60grp-index 0 with the error PFE_ERROR_NOT_FOUND: route check failed, entry not found with steady state measurement check**. [PR1472454](#)
- Few OAM sessions are not established with scale EVPN etree and CFM configuration. [PR1478875](#)
- Remote mep are not coming up as expected while verifying MIP functionality with bridge domains. [PR1484303](#)
- Traffic loss is observed after unified ISSU, when you enable and activate the interface and when you disable and deactivate the interface. [PR1493723](#)
- ARP packets are getting dropped at certain IRB interfaces after switchover phase in unified ISSU in MX Series routers. [PR1500183](#)
- Multiple NGMPC4, RMPC9, fpc8-smpc core has observed on MX2010 platform and FPC's are toggling offline and online. [PR1503947](#)
- LACP state machine will not converge to CD on PEER device, because of which traffic drop is seen on DUT. [PR1505465](#)

Routing Protocols

- LDP OSPF are 'in sync' state and the reason observed for this is "IGP interface down" with LDP synchronization enabled for OSPF. `user@host> show ospf interface ae100.0 extensive` Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. "IGP interface down" is observed as the reason because although LDP notified OSPF that LDP synchronization is achieved, OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up. [PR1256434](#)
- From Junos OS Release 16.1 and later, when both sides of a BGP session are doing NSR Routing Engine switchover simultaneously (double failures), depending on the configuration and scale, there is a chance the BGP session might get stuck and BGP PDUs cannot be exchanged. The permanently stuck OutQ are seen which is a typical symptom for this issue. This is because both sides are waiting for socket record boundary. Both sides are waiting to drain their partially written PDU. Neither side read at this state, leading to permanent stuck. [PR1323306](#)
- It is possible for a GNF with rosen6 multicast to display stuck KRT queue entries after recovery from a dual Routing Engine reboot at the BSYS. [PR1367849](#)
- Performance improvement with `addpath-optimization` statement configured will vary across releases due to variability of baseline convergences without the configuration statement. [PR1395684](#)
- When configuring an alternate incoming interface for a PIM RPF check using `rpf-selection`, you might find that additional groups outside the configured range switch to the alternate incoming interface. [PR1443056](#)
- Even when `protocols mpls traffic-engineering bgp-igp` statement is configured, UDP tunnel routes are not added to inet.0. The UDP tunnel routes are added only to inet.3 table whether the configuration statement is configured or not. [PR1457426](#)
- NSR with larger scales of 4K+ IPv6 peer is not qualified or supported. [PR1461436](#)
- BGP graceful restart might have some traffic loss when sharding is enabled. [PR1475773](#)
- In next-generation MVPN setup, using MPC10 on egress PE device with load balance join of multiple groups in C_VPN, egress PE might not receive multicast traffic. [PR1476969](#)
- On all Junos OS platforms running in the Multicast Source Discovery Protocol (MSDP) scenario, if there are multiple overlapping mesh groups existing in the topology, for example, node A and B belong to mesh group M, node A and D belong to mesh group N, while node B and D belong to a different mesh group O, the Source Active (SA) messages sent from node A might get continuously looped due to this issue. This might lead to the rpd memory leak and 100 percent CPU utilization. [PR1485206](#)
- The last route from RIB inet6.x will not be deleted if there is another configuration present under RIB configuration (for example, `set routing-options rib inet6.0 static defaults active`). As a workaround,

remove any other configuration present in RIB inet6.x configuration and then delete the last route.

[PR1495477](#)

- On all Junos OS platforms with BGP max-acceptance prefix limit (accepted-prefix-limit with teardown action) configured, the rpd might crash when a policy change causes the number of imported routes cross the configured threshold value and teardown action that is being performed by the router.

[PR1499977](#)

- EVPN MLD snooping group limit is not working as expected. [PR1501485](#)

Services Applications

- Memory corruption leads to L2tp process to crash. [PR1407885](#)
- In Layer 2 tunneling protocol (L2TP) subscriber environment with Juniper Network's L2TP tunnel switch (LTS) and L2TP network server (LNS) are used, if client negotiates Link Control Protocol (LCP) with no ppp-options to L2TP Access Concentrator (LAC). It might cause some problems but it has no service impact. These ppp-options are Address and Control Field Compression (ACFC), PPP Protocol Field Compression (PFC) and Async Control Character Map (ACCM). The reason is that when MX Series router functions as L2TP LTS or LNS, it will initiate LCP re-negotiation (ppp-options) if Last Received LCP CONFREQ AVP (attribute-value pair) is not included in Incoming-Call-Connected (ICCN) message received from LAC. And this might cause some problems for peers which do not support these options and do not want to negotiate with them. [PR1426164](#)
- This issue is seen while attaching cpd service on subscriber. This is because of the unwanted inline gencfg message that is pushed from the Routing Engine to the Packet Forwarding Engine. Therefore, the Packet Forwarding Engine crashes. [PR1502527](#)

VPNs

- In MVPN environment with SPT-only option, if the source or receiver is connected directly to c-rp PE and the MVPN data packets arrive at the c-rp PE before its transition to SPT, the MVPN data packets might be dropped. [PR1223434](#)
- P2MP LSP replication to backup Routing Engine is not proper. [PR1453900](#)
- MPC10E: NG-MVPN GTM for remote IPv6 source is not working. [PR1454163](#)
- On MVPN scenario the LSP might stay down on removing all VT interfaces from 1 hop egress. [PR1474830](#)
- After NSR switchover, sometimes the selective tunnel on the new master Routing Engine might fall back to the inclusive tunnel. After some time, the traffic will be migrated back to the selective tunnel. Some traffic loss will be seen during this migration. [PR1475204](#)
- When there is a layer 2 circuit local switching configuration along with connection protection and backup neighbour configuration and layer 2 circuit trace logs are enabled. If you remove the layer 2 circuit configuration then we hit this issue. As a result, rpd process crashes. [PR1502003](#)

SEE ALSO

[What's New | 87](#)[What's Changed | 110](#)[Known Limitations | 116](#)[Resolved Issues | 135](#)[Documentation Updates | 168](#)[Migration, Upgrade, and Downgrade Instructions | 169](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.4R2 | 135](#)
- [Resolved Issues: 19.4R1 | 147](#)

Learn which issues were resolved in Junos OS main release and the maintenance releases for MX Series routers.

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

Resolved Issues: 19.4R2

Application Layer Gateways (ALGs)

- SIP messages that need to be fragmented might be dropped by the SIP ALG. [PR1475031](#)
- FTPS traffic might get dropped on MX Series platforms if FTP ALG is used. [PR1483834](#)

Authentication and Access Control

- The LLDP packets might get discarded on all platforms running Junos OS. [PR1464553](#)

Class of Service (CoS)

- MX Series routers generated OAM/CFM LTR message sent with a different priority than the incoming OAM/CFM LTM message. [PR1466473](#)
- Unexpected traffic loss might be discovered in certain conditions under Junos fusion scenario. [PR1472083](#)

- MX10008 and MX100016 might generate cosd core file after executing **commit** or **commit check** command if **policy-map** configuration is set. [PR1475508](#)

EVPN

- Deleting a Layer 2 logical interface generates an error if the interface is not deleted first from EVPN. [PR1482774](#)
- The ESI of IRB interfaces does not update after autonomous-system number change if the interface is down. [PR1482790](#)
- The ARP entry is gone from kernel after adding and deleting the virtual-gateway-address. [PR1485377](#)

Forwarding and Sampling

- A problem with statistics on some interfaces of a router might be observed after FPC or PIC reboot. [PR1458143](#)
- Type 1 ESI/AD route might not be generated locally on EVPN PE device in the all-active mode. [PR1464778](#)
- Traffic might not be forwarded into the right queue but the default queue when VPLS traffic has three or more VLAN tags with VLAN priority 5. [PR1473093](#)
- The filter might not be installed if the "policy-map xx" is present under it. [PR1478964](#)

General Routing

- Syslog error messages **PFEIFD: Could not decode media address with length 0** might be generated by the Packet Forwarding Engine. [PR1341610](#)
- UID might not release properly in some scenarios after service session deactivation. [PR1188434](#)
- Reduce **XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR** from fatal to major. [PR1390333](#)
- NAPT66 split is not supported with AMS. Therefore, commit fails with IPv6 pool in AMS. [PR1396634](#)
- The nonexistent subscribers might appear in the **show system resource-monitor subscribers-limit chassis extensive** output. [PR1409767](#)
- Egress monitored traffic is not mirrored to destination for analyzers on MX Series routers. [PR1411871](#)
- **FPC x Voltage Tolerance Exceeded** alarm is raised and cleared upon bootup of JNP10K-LC2101. [PR1415671](#)
- **Resetting the playback engine** log messages are seen on MPC5E. [PR1420335](#)
- PF core voltage is not set as per the required e-fuse value and remains at default value (0.9V) on JNP10008-SF and JNP10016-SF. [PR1420864](#)
- PTP might not work on MX104 if phy-timestamping is enabled. [PR1421811](#)
- Disable PTP and show warning when hyper mode is configured. [PR1429527](#)
- Error **dfw_abstract_issu_stats_counters_restore:2222 Failed to find Index = 4613734?** is seen during unified ISSU with 19.3I-20190409_dev_common.0.2212. [PR1429879](#)

- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- **ZF interrupts [MAJOR] Out-of-range Dest PFE INTR for Gnt** seen during MPC6 or MPC9 line card startup. [PR1436148](#)
- IRB over VTEP unicast traffic might get dropped on MX Series platforms. [PR1436924](#)
- Unified ISSU is failing from Junos OS Release 19.1R1 legacy Junos OS release images. [PR1438144](#)
- The EX Series ports might stay in up state even if the EX4600 and QFX5100 lines of devices are rebooted. [PR1441035](#)
- The interface might go into administrator down state after FPC restart with PTP configuration enabled. [PR1442665](#)
- Irregular traffic drop might be seen when traffic is ingress from MPC3E and egress to MPC10E. [PR1445649](#)
- IPv6 throughput numbers for NAT with HTTP traffic is not at par with IPv4. [PR1449435](#)
- **Mixed Master and Backup RE types** alarm occurs as MX2008 with RE-MX2008-X8-128G detect backup Routing Engine as RE-MX2008-X8-64G. [PR1450424](#)
- Main chassisd thread at a JNS GNF could experience stalls upon GNF SNMP polling for hardware-related OIDs. [PR1451215](#)
- Support for drop flows when packet is dropped. [PR1451921](#)
- Interfaces shutdown by 'disable-pfe' action might not be up using MIC offline or online command. [PR1453433](#)
- Add **syslog** configuration statement to stateful firewall rule then condition. [PR1453502](#)
- When scale configurations are applied from approximately 10 minutes, the chassisd process CLI will either have a delay in response or will time out. [PR1454638](#)
- On 4x1GE using QSFP28 optics, continuous logging in chassisd process occurs when "speed 1g" is configured: **pic_get_nports_inst** and **ch_fru_db_key**. [PR1456253](#)
- The PID is the same before and after installation of JSU package. [PR1457304](#)
- LSP statistics are not getting reset after routing is restarted. [PR1458107](#)
- Inline S-BFD packets are dropped on MPC6E MIC1/PIC1 ports: 0-11. [PR1459529](#)
- Subscriber statistics can be broken after the unified ISSU. [PR1459961](#)
- Multiple leaf devices and prefixes are missing when LLDP neighbor is added after streaming is started at the global level. [PR1460347](#)
- Support of del_path for the LLDP neighbor change at various levels. [PR1460621](#)
- Explicit deletion notification (del_path) is not received when LLDP neighbor is lost as a result of disabling local interface on the DUT through CLI (gNMI). [PR1461236](#)
- On MPC10E more output packets are seen than expected when ping function is performed. [PR1461593](#)

- Traffic drop might be seen in scaled scenario with VRRP sessions configured on aggregated Ethernet interfaces. [PR1462310](#)
- The Routing Engine switchover might not be triggered when the master CB clock fail. [PR1463169](#)
- MVPN traffic might be dropped after performing switchover. [PR1463302](#)
- The subscribers might not pass traffic after doing some changes to the dynamic-profile filter. [PR1463420](#)
- The IPoE subscriber route installation might fail. [PR1464344](#)
- The bbe-smgd process might generate a core file (0x000000000088488c in bbe_autoconf_delete_vlan_session_only (session_id=918) at `../../../../src/junos/usr.sbin/bbe-svcs/smd/plugins/autoconf/bbe_autoconf_plugin.c:3115`). [PR1464371](#)
- The CPU utilization on mgd daemon might get stuck at 100 percent after the NETCONF session is interrupted by flapping interface. [PR1464439](#)
- The MS-MIC might not work when it is used on specific MPC. [PR1464477](#)
- Constant messages flooding in log **summit_pic_port_profile_isvalid: VALID Port profile**. [PR1464879](#)
- The jdhcpd process might consume high CPU use, and no further subscribers can be brought up if there are more than 4000 dhcp-relay clients in the MAC-MOVE scenario. [PR1465277](#)
- On MPC10E line cards, the bandwidth-percent with shaping-rate might not work as expected on aggregated Ethernet interfaces after shaping-rate change. [PR1465766](#)
- ICMP error messages are still unreceived after enabling the **enable-asymmetric-traffic-processing** configuration statement. [PR1466135](#)
- The PPPoE subscribers get stuck because the PPPoE inline keepalives do not work properly. [PR1467125](#)
- The DOM MIB alarm for the channelized 10 Gigabit Ethernet interface is not showing any alarm for LF/RF. [PR1467446](#)
- The error **the user-ad-authentication subsystem is not responding to management requests** might be seen when executing the **show services user-identification identity-management status** command. [PR1467991](#)
- Not able to get the service sessions when NAT64 is configured with destination-prefix length 32. [PR1468058](#)
- Daemons might not be started if **commit** is executed after **commit check**. [PR1468119](#)
- Benign logs might show in Junos OS Release 19.3R2 when switching between configurations using **load-override** with GRES and **commit-synchronize**. [PR1468234](#)
- IPv6 dynamic subscribers might be unable to access on Junos OS Release 18.2R3 and later. [PR1468414](#)
- Optic measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- The rpd process might crash if BGP sharding is enabled. [PR1468676](#)

- The tcp-log connections fail to reconnect and get stuck in "Reconnect-In-Progress" state. [PR1469575](#)
- A hierarchical-scheduler should not be configured on a "ps" interface. [PR1470049](#)
- Traffic might be not policed properly on MPC10E interfaces. [PR1470629](#)
- SNMP interface-mib stops working for PPPoE clients. [PR1470664](#)
- Unable to set up 26 million sessions (NAPT44) at 900 Kpps per second. [PR1470833](#)
- In rare occasions the router might send out one extra URR quota value for a bearer. [PR1470890](#)
- Message **fpcX user.notice logrotate: ALERT exited abnormally with [1]** pops at 04:02:01. [PR1471006](#)
- DHCP relay with forward-only might fail to send offer messages when DHCP client is terminated on logical tunnel interface. [PR1471161](#)
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card might cause major alarms on other FPCs in the system. [PR1471372](#)
- In cRPD platform, license violations are captured as nagging log messages and no alarm is raised. [PR1471455](#)
- The clksyncd crash might be seen when PTP over aggregated Ethernet is configured on MX104 platform. [PR1471466](#)
- ARP suppression (default enabled) in EVPN does not work on MX10008 and MX10016 line cards. [PR1471679](#)
- The pccd core file and PCEP session flap might be seen in PCE-initiated or PCE-delegated LSP scenario. [PR1472051](#)
- Chassis alarm on BSYS: RE0 to one or many FPCs is shown in em1: backup Routing Engine. [PR1472313](#)
- Service accounting statistics do not get updated after changing to firewall filters. [PR1472334](#)
- The kernel might crash and vmcore might be observed after configuration change is committed. [PR1472519](#)
- Restarting the rpd back-to-back crashes the rpd process. [PR1472643](#)
- Active error counts are not increasing for Layer 2 circuit in SYNCE cards. [PR1472660](#)
- SDB goes down frequently if the **reauthenticate lease-renewal** statement is enabled for DHCP. [PR1473063](#)
- Dynamic-profile for VPLS-PW pseudowire incorrectly is reporting Dynamic Static Subscriber Base Feature license alarm. [PR1473412](#)
- JET certificate request utility is broken in JET sb. [PR1473892](#)
- Ingress multicast replication does not work with GRES configuration. [PR1474094](#)
- Subnet information might get corrupted if it is passed by a RADIUS server. [PR1474097](#)
- MX150 core files are not seen under **show system core-dumps**. [PR1474118](#)

- QSA adapter Lane 0 port might be also brought down when disabling one of the other lanes. [PR1474231](#)
- A newly added LAG member interface might forward traffic even though its micro BFD session is down. [PR1474300](#)
- Memory leak might occur on MX Series routers with SAEGW and usage reporting rule (URR) report enabled. [PR1474306](#)
- The output of the **show services sessions** and **show services sessions extensive** commands do not display member interface of the AMS where the session got landed. [PR1474313](#)
- The clksyncd process generates core file after the GRES. [PR1474987](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- The RADIUS accounting updates of service session have incorrect statistics data. [PR1475729](#)
- Traffic loss might be seen as backup Routing Engine takes around 20 seconds to acquire mastership. [PR1475871](#)
- Traffic drop might be observed while performing a unified ISSU on the MX2000 line of devices. [PR1476505](#)
- The bbe-mibd might be crashed on MX Series platform in subscriber environment. [PR1476596](#)
- Traffic loss might be observed to the LNS subscribers in case the **routing-service** configuration statement is enabled under the dynamic-profile. [PR1476786](#)
- **Attribute Length Error** error notification is observed. [PR1477089](#)
- Traffic loss might be seen in SAEGW scenario after SAEGW daemon restart or both GRES operation. [PR1477461](#)
- In NAT-T scenario the IKE version 2 IPsec tunnel might flap if the tunnel initiator is not behind NAT. [PR1477483](#)
- The rpd process might crash when the JET RIB API is used to set the "bandwidth" attribute. [PR1477745](#)
- The following error log messages are observed: **chassisd[7836]: %DAEMON-3-CHASSISD_IOCTL_FAILURE: acb_get_fpga_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device) after every commit.** [PR1477941](#)
- Packet Forwarding Engine might be disabled because of the major error on MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9. [PR1478028](#)
- The **show evpn statistics instance** command gets stuck on multihomed scenario. [PR1478157](#)
- At scale logins of both default and dedicated bearers might require retries from the control plane. [PR1478191](#)
- FPC memory leak might happen after executing **show pfe route**. [PR1478279](#)
- Output chain filter counters are not proper. [PR1478358](#)

- MX Series-based MPC linecard might crash when there is bulk route update failure in a corner case. [PR1478392](#)
- FPC with **vpn-localization vpn-core-facing-only** configuration might get stuck in ready state after configuration removal or restoration because vt logical interface under MVPN is not cleaned up (physical interface cleanup failed for vt-ifl under MVPN instance). [PR1478523](#)
- The protocol MTU might not be changed on lt- interface from the default value. [PR1478822](#)
- TCP-log sessions might be in established state but no logs get sent out to the syslog server. [PR1478972](#)
- Mobile-edge sessions might be lost if GRES is being performed while sessions are logged in with URR enabled. [PR1478985](#)
- The SCBE3 fabric plane gets into check state in MX Series Virtual Chassis. [PR1479363](#)
- After kmd restart IPsec SA comes up but traffic fails for some time in certain scenarios. [PR1480692](#)
- The rpd process might crash when executing **show route protocol l2-learned-host-routing** or **show route protocol rift** CLI command on a router. [PR1481953](#)
- Logging in some PPPoE subscribers through aggregated Ethernet interface might cause the device to reboot. [PR1482431](#)
- Fragmentation limit and reassembly timeout configuration under services option is missing for SPC3. [PR1482968](#)
- Packet loss might be observed after device is rebooted or l2ald is restarted in EVPN-MPLS scenario. [PR1484468](#)
- ARP entry might not be created in the EVPN-MPLS environment. [PR1484721](#)
- The logical tunnel interface might not work on MPC10 line card. [PR1484751](#)
- Interface input error counters are not increasing on MX150. [PR1485706](#)
- The **krt-nexthop-ack-timeout** might not automatically be picked up on restarting the rpd process. [PR1485800](#)
- MPC10E installed in FPC slot 4 might drop host-outbound traffic. [PR1485942](#)
- Command completion help text for LLDP-MED coordinate configuration statement contains spelling errors. [PR1486327](#)
- The aftd process might crash when MPC10 line card is installed. [PR1487416](#)
- With 4 member AMS used in the service-set, commit check fails when /30 subnet address is used as NAT pool IP. [PR1489885](#)
- The syslog error **Failed to connect to the agentx master agent (/var/agentx/master): Unknown host (/var/agentx/master) (No such file or directory)** is continuously generated with dns-sinkholing. [PR1490487](#)

- When NAT/SFW rule is configured with application-set with multiple applications having different TCP inactivity-timeout, sessions are not getting TCP inactivity-timeout as according to the configured application order. [PR1491036](#)
- The unified ISSU is not supported on next-generation MPC cards. [PR1491337](#)
- FPCs might stay down or restart when swapping MPC7, MPC8, and MPC9 with MPC10 or vice versa in the same slot. [PR1491968](#)
- Port numbers logged in ALG syslog are incorrect. [PR1497713](#)
- Few of DHCP INFORM packets specific to particular VLAN might be taking the incorrect resolve queue. [PR1467182](#)
- MPC11 is not supported in Junos OS 19.4 Release. [PR1503605](#)

High Availability (HA) and Resiliency

- Unified ISSU might fail on MX204 and MX10003 Virtual Chassis with an error message. [PR1480561](#)

Infrastructure

- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. [PR1462986](#)
- The scheduled tasks might not be executed if "cron" daemon goes down without restarting automatically. [PR1463802](#)

Interfaces and Chassis

- Restarting chassisd with GRES disabled might cause FPC to restart and some demux interfaces to be deleted. [PR1337069](#)
- When the logical interface is associated to a routing instance inside a LR is removed from the routing instance, the logical interface is not added to the default routing instance. [PR1444131](#)
- Continuous VRRP state transition (VRRP master/backup flap) is seen when one device drops VRRP packets. [PR1446390](#)
- Interface descriptions might be missing under logical systems CLI. [PR1449673](#)
- **MIC Error code: 0x1b0002** alarm might not be cleared for MIC on MPC5E when the voltage has returned to normal. [PR1467712](#)
- When you configure ESI on a physical interface and disable a logical interface, traffic drop will be seen under the physical interface. [PR1467855](#)
- Executing commit might hang up due to stuck dcd process. [PR1470622](#)
- Commit error is not thrown when member link is added to multiple aggregation group with different interface specific options. [PR1475634](#)
- The interface on MIC3-100G-DWDM might go down after performing an interface flap. [PR1475777](#)

- When you delete and add a logical interface (both the logical interfaces with same VLAN ID) in a single commit, the configuration check fails with the error **duplicate VLAN-ID**. [PR1477060](#)
- MC-AE interface might be shown as unknown status if adding the sub-interface as part of the VLAN on the peer MC-AE node. [PR1479012](#)
- Seeing commit failure on MX Series routers with **VPI number 0 for child at-4/2/4.0 not set at the IFD** when **allow-any-vci** attribute is configured under the logical interface. [PR1479153](#)
- The **vrmp-inherit-from** change operation leads to packet loss when traffic is forwarded to the VIP gateway. [PR1489425](#)
- Traffic is not forwarded properly when traffic-control-profiles with logical interface queues are configured. [PR1475350](#)

J-Web

- Junos OS: Security vulnerability in J-Web and web based (HTTP/HTTPS) services. [PR1499280](#)

Junos Fusion Enterprise

- SDPD core file is found **vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry**. [PR1454335](#)
- Loop detection might not work on extended ports in Junos fusion scenarios. [PR1460209](#)

Layer 2 Ethernet Services

- The jdncpd process might go into infinite loop and might use 100 percent CPU. [PR1442222](#)
- On MX2010 and MX2020, no alarm is generated when FPC is connected to master Routing Engine through backup Routing Engine/CB. [PR1461387](#)
- Member links state might be asynchronized on a connection between a PE device and a CE device in EVPN active-active scenario. [PR1463791](#)
- Telemetry data for relay/bindings/binding-state-v4relay-binding and relay/bindings/binding-state-v4relay-bound is not correct. [PR1475248](#)
- Issues with DHCPv6 relay processing confirm and reply packets. [PR1496220](#)

Layer 2 Features

- Connectivity is broken through LAG because of the members configured with hold-time and force-up. [PR1481031](#)

MPLS

- Traffic loss might be seen if p2mp with NSR enabled. [PR1434522](#)
- The traffic might silently get dropped and discarded after the LACP timeout. [PR1452866](#)
- P2MP LSP might flap after VT interface in MVPN routing instance reconfigured. [PR1454987](#)
- The device might use the local-computed path for the PCE-controlled LSPs after link/node failure. [PR1465902](#)

- Fast reroute detour next-hop down event might cause primary LSP down in a particular scenario. [PR1469567](#)
- The rpd process might crash during shutdown. [PR1471191](#)
- The rpd crash might be seen after doing some commit operations that could affect RSVP ingress routes. [PR1471281](#)
- Errors continuously flood in backup Routing Engine **JTASK_IO_CONNECT_FAILED: RPD/./var/run/rpdtmd_control: Connecting to 128.0,255.255,255.255,0.0.0,0.0.0, failed: No such file or directory.** [PR1473846](#)
- LDP/BFD sessions are not coming up in a scaled setup. [PR1474204](#)
- RSVP LSPs might not come up in a scaled network with a high number of LSPs if NSR is used on transit router. [PR1476773](#)
- Kernel crashes and device might restart. [PR1478806](#)
- RPD 100 percent CPU load and rpd core files are generated on the backup Routing Engine. [PR1479249](#)
- High CPU utilization for rpd might be seen if RSVP is implemented. [PR1490163](#)
- The rpd might crash when flapping BGP with FEC 129 VPWS enabled. [PR1490952](#)
- Backup Routing Engine might crash unexpectedly because of a rare timing issue. [PR1493053](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. [PR1493431](#)

Platform and Infrastructure

- Jcrypto syslog help package and events are not packaged even when errmsg is compiled. [PR1290089](#)
- With CNH enabled, the MPLS CoS rewrite does not work for 6PE traffic. [PR1436872](#)
- Traffic loss might be seen in case of Ethernet frame padding with VLAN. [PR1452261](#)
- The MPC might drop packets after enabling firewall fast-lookup filter. [PR1454257](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- On MX204 platform, Packet Forwarding Engine errors might occur when incoming GRE tunnel fragments (1) get sampled and (2) undergo inline reassembly. [PR1463718](#)
- EVPN-VXLAN Type-5 tunnel might not work properly on MX Series platforms. [PR1466602](#)
- All the subscriber services might be unavailable on vBNG running on MX150 or vMX running in payg mode. [PR1467368](#)
- The Layer 2 traffic over aggregated Ethernet interfaces sent from one member to another is corrupted on MX Series Virtual Chassis setup. [PR1467764](#)
- JNH memory leaks might be seen after CFM session flap for LSI/VT interfaces. [PR1468663](#)

- The switch might not be able to learn MAC address with dot1x and **interface-mac-limit** configured. [PR1470424](#)
- SSH login might hang and the TACAS + server closes the connection without sending any authentication failure response. [PR1478959](#)
- The convergence time for MVPN fast upstream failover might be more than 50ms. [PR1478981](#)
- The **show system buffer** command displays all zeros in the MX104 chassis. This is a cosmetic issue and there is no service impact reported. [PR1484689](#)
- MAC malformation might happen in a rare scenario under MX Series Virtual Chassis setup. [PR1491091](#)
- In node slicing setup MPLS TTL might be set to zero when the packet goes through af interface configured with CCC family. [PR1492639](#)
- Routing Engine crash might be seen when a large number of next hops are quickly deleted and readded in large ARP/ND scale scenario. [PR1496429](#)
- Python or Slax script might not be executed. [PR1501746](#)

Routing Policy and Firewall Filters

- The router-id from martian address range cannot be committed even if the range is allowed by configuration. [PR1480393](#)

Routing Protocols

- The BGP sessions might flap with the configuration statement **keep none** when a VRF is deleted. [PR1439560](#)
- The CPU utilization on rpd process spins at 100 percent once the same external BGP route is learned in different vrf tables. [PR1442902](#)
- The TI-LFA backup path for adj-sid is broken in OSPF. [PR1452118](#)
- SSH login might fail if a user account exists in both local database and RADIUS/TACACS+. [PR1454177](#)
- MoFRR with MLDP inband signaling do not work. [PR1454199](#)
- TI-LFA might be unable to install backup path in the routing table in a specific case. [PR1458791](#)
- The rpd might crash if IPv4 routes are programmed with IPv6 next hop through JET APIs. [PR1465190](#)
- BGP peers might flap if the parameter of hold-time is set small. [PR1466709](#)
- The configured BGP damping policy might not take effect after BGP is disabled and then enabled followed by **commit**. [PR1466734](#)
- The BGP multipath does not work if multipotology routing is used. [PR1467091](#)
- The rpd might crash after configuring **independent-domain** under the master routing-instance. [PR1469317](#)
- The mcsnoopd might crash when the STP moves the mrouter port to the blocked state. [PR1470183](#)

- The BFD client session might flap when removing BFD configuration from the peer end (from another vendor) of the BFD session [PR1470603](#)
- The rpd might stop when both instance-import and instance-export policies contain the as-path-prepend action. [PR1471968](#)
- The rpd process might crash with BGP multipath and damping configured. [PR1472671](#)
- Removing cluster from BGP group might cause prolonged convergence time. [PR1473351](#)
- Adjacency SID might be missed and not be advertised to peer/controller/BMP monitor in BGP-LS NLRI. [PR1473362](#)
- The rpd process might crash with BGP multipath and route withdraw occasionally. [PR1481589](#)
- The rpd process might crash when deactivating logical systems. [PR1482112](#)
- The rpd might be crashed after BGP peer flap. [PR1482551](#)
- The rpd crashes if the same neighbor is set in different RIP groups. [PR1485009](#)
- The BGP-LU routes do not have the label when BGP sharding is used. [PR1485422](#)
- Removal of BGP and **rib-sharding** configuration can cause routing protocols to become unresponsive. [PR1485720](#)
- The rpd crashes if BGP LLGR with RIB sharding and traceoptions for graceful-restart configured. [PR1486703](#)
- The rpd might crash when you perform GRES with MSDP configured. [PR1487636](#)
- High CPU utilization might be observed when the outgoing BGP updates are sending slowly. [PR1487691](#)
- The rpd process might generate core file after **always-compare-med** is configured for BGP path-selection. [PR1487893](#)
- BGP RIB sharding feature cannot be run on a system with a single CPU. [PR1488357](#)
- The rpd crashes when reset OSPF neighbours. [PR1489637](#)

Services Applications

- MX Series L2tp LTS fails to forward the agentCircuitId and agentRemotId AVP toward the LNS. [PR1472775](#)
- The kmd might crash due to the incorrect IKE SA establishment after the remote peer's NAT mapping address has been changed. [PR1477181](#)

Subscriber Access Management

- No volume attribute in Accounting Stop for Service session when Activated Services session by Configuration [PR1470434](#)
- The logical interfaces might be missing in the NAS Port ID. [PR1472045](#)

- The authd might crash after unified ISSU from Junos OS Release 18.3 or earlier to Release 19.4 or later. [PR1473159](#)
- Some address relevant fields are missing when executing "test aaa ppp" command [PR1474180](#)
- Syslog messages occur **pfe_tcp_listener_open_timeout: Peer info msg not received from addr: 0x6000080. Socket 0xfffff804ad23c2e0 closed.** [PR1474687](#)
- Process of verifying deleting services through CoA when the specified family-type has been deactivated failed because an incorrect number of service sessions are active. [PR1479486](#)
- The CoA request might not be processed if it includes "proxy-state" attribute. [PR1479697](#)
- The **mac-address** CLI option is accessible under the **access profile <profile-name> radius options calling-station-id-format** stanza. [PR1480119](#)

User Interface and Configuration

- On MX Series platforms, the J-Web page might not get redirected to login once the session expires due to idle timeout. [PR1459888](#)

VPNs

- Traffic loss might be observed when the inter-AS NG-MVPN VRF is disabled on one of the ASBRs. [PR1460480](#)
- The l2circuit displays MM status, which might cause traffic loss. [PR1462583](#)
- The l2circuit connections might be stuck in OL state after changing the l2circuit community and flapping the primary LSP path [PR1464194](#)
- The rpd might crash when "link-protection" is added/deleted from LSP for MVPN ingress replication selective provider tunnel [PR1469028](#)
- The MPC10E-15C-MRATE next-generation MPVN ingress replication flushing out is not proper when in egress the ingress replication configuration is deactivated. [PR1475834](#)

Resolved Issues: 19.4R1

Class of Service (CoS)

- Unexpected traffic loss might be discovered in certain conditions under fusion scenario. [PR1472083](#)

EVPN

- Asynchronous results between ARP table and Ethernet switching table occurs if EVPN ESI link flaps multiple times. [PR1435306](#)
- EVPN or MPLS IRB logical interface might not come up when local Layer 2 interface is down. [PR1436207](#)
- The specific source ports of UDP packet are dropped on EVPN or VXLAN setup. [PR1441047](#)
- The rpd might crash or consume 100 percent of CPU after flapping the routes. [PR1441550](#)

- Restarting Layer 2 learning might cause some remote MAC addresses to move into forwarding dead state. [PR1441565](#)
- Traffic drop might be seen in EVPN Layer 3 Gateway. [PR1442319](#)
- Core-isolation feature does not work after you set or delete the **no-core-isolation** command on MX Series router. [PR1442973](#)
- The EVPN type 2 routes might not have advertised properly in the logical systems. [PR1443798](#)
- The local host address is missing from the EVPN database and mac-ip-table. [PR1443933](#)
- The bridge mac-table age timer does not expire for rbeb interfaces. [PR1453203](#)
- Instance type is changed from VPLS to EVPN, resulting in loss of packet. [PR1455973](#)
- Preference-based DF Election algorithm does not work on LT interface. [PR1458056](#)
- The rpd crash might be seen if BGP route is resolved over the same prefix protocol next-hop in inet.3 table that has both RSVP and LDP routes. [PR1458595](#)
- The DF router might send ARP request or NS to the local segment. [PR1459830](#)
- In EVPN scenario, memory leak might be observed when **proxy-macip-advertisement** is configured. [PR1461677](#)
- Traffic received from VTEP is dropped if the VNI value used for type-5 routes is greater than 65535. [PR1461860](#)
- Rpd might crash with EVPN-related configuration changes in static VXLAN to MPLS stitching scenario. [PR1467309](#)

Forwarding and Sampling

- You might not be able to apply the firewall filter configuration change after ISSU upgrades to release 16.1R1 or later. [PR1419438](#)
- The following syslog error messages are seen at **pfed: rtplib: ERROR received async message with no handler: 28** [PR1458008](#)
- On the MX Series and QFX Series devices, the Layer 2 ald process might leak memory. [PR1455034](#)
- The rt-delay-threshold can be set below 1 second but rt-marker-interval is limited to 1 second. [PR1425544](#)
- The high CPU utilization of Layer 2 ald is seen after replacing EVPN configuration. [PR1446568](#)
- On MX Series routers with MPC10 line cards, the incoming packets might get dropped. [PR1446736](#)
- On MX204, input/output counters of aggregated Ethernet bundle or member links configured on non-default logical systems are not updated. [PR1446762](#)
- ARP packets gets dropped by Packet Forwarding Engine after chassis-control in the MX Series routers. [PR1450928](#)
- Commit error and dfwd core file might be observed when you apply a firewall filter with the **then traffic-class** or **then dscp** action. [PR1452435](#)

- The following false warning message is seen on commit (commit check) after upgrading to Junos OS Release 19.2R2-S1.4: **warning: vxlan-overlay-load-balance configuration for forwarding options has been changed.....** [PR1459833](#)
- On MX Series router, the following logs are seen: **L2ALD_MAC_IP_LIMIT_REACHED_IF: Limit on learned MAC+IP bindings reached for .local.1048605; current count is 1024.** [PR1462642](#)
- The EA WAN SerDes gets into a stuck state, leading to continuous "DFE tuning timeout" errors and link staying down. [PR1463015](#)
- An output bandwidth-percent policer with logical-bandwidth-policer applied to an aggregated Ethernet bundle along with an output-traffic-control-profile has incorrect effective policing rate. [PR1466698](#)
- Type 1 ESI/ or AD route are not generated locally on EVPN PE in all-active mode. [PR1464778](#)

General Routing

- Load balancing is uneven across aggregated Ethernet member links when the aggregated Ethernet bundle is part of an ECMP path. The aggregated Ethernet member links must span the Virtual Chassis members. [PR1255542](#)
- Unable to configure **pic-mode** when MPC10E is inserted. [PR1452467](#)
- Basic circuit cross-connect traffic flow does not occur with the logical systems. [PR1474983](#)
- Service accounting statistics do not get updated after changes are made to the firewall. filters [PR1472334](#)
- System reboot is required when GRES is enabled or disabled with the **mobile-edge** configuration. [PR1444406](#)
- Agentd memory might leak and crash when the RPD session closes without releasing memory. [PR1455384](#)
- Active error counts are not increasing for Layer 3 circuit in SYNCE cards. [PR1472660](#)
- The PTP function might consume the kernel CPU for a long time. [PR1461031](#)
- Not able to get the service sessions when configure NAT64 with destination-prefix length is 32. [PR1468058](#)
- Inner-list functionality with dual tag does not work. Traffic gets dropped at the ingress port. [PR1469396](#)
- Memory leak on Layer 2 cpd process causes Layer 2 cpd to crash. [PR1469635](#)
- On MPC10 interfaces, certain configuration steps might cause traffic to not get policed properly. [PR1470629](#)
- The interfaces on MPC-3D-16XGE-SFPP card does not get created after upgrading the system to Junos OS Release 18.1 and later. [PR1471429](#)
- In cRPD platform, license violations are captured as nagging log messages and no alarm is raised. [PR1471455](#)
- PCC tries to send a report to PCE but the connection between PCC and PCE is not in the up state especially in the case of MBB in PCE provisioned or controlled LSP. [PR1472051](#)

- Active error counts are not increasing for I2C in the SYNCE cards. [PR1472660](#)
- MX10000 QSA adapter lane 0 port goes in the down state when disabling one of the other lanes. [PR1474231](#)
- The **show services sessions** and **show services sessions extensive output** command does not display the member interface of the AMS where the session got landed. It displays only the AMS interface name. [PR1474313](#)
- **request system [halt | power-off]** reboots the system instead of halting the system. [PR1474985](#)
- The physical interface of aggregated Ethernet might take time to come up after disabling or enabling the interface. [PR1465302](#)
- Observing **bbe-smgd-core (0x000000000088488c** in **bbe_autoconf_delete_vlan_session_only (session_id=918)** at
 `../..../src/junos/usr.sbin/bbe-svcs/smd/plugins/autoconf/bbe_autoconf_plugin.c:3115`). [PR1464371](#)
- ZT VPLS: The **native-vlan-id** functionality does not work and an untagged traffic does not pass with the **native-vlan-id** configuration. [PR1463544](#)
- Traffic might be impacted due to fabric hardening being stuck. [PR1461356](#)
- The **SmiHelperd** process is not initialized in the Junos OS PPC Releases. [PR1455667](#)
- Queue data might be missing from the following path: **/interfaces/interface/state** [PR1456275](#)
- Interface with Tri-rate Copper SFP (P/N:740-01311) in MIC 3D 20x 1GE(LAN)-E,SFP stops forwarding traffic after unified ISSU. [PR1379398](#)
- The **vehosd** application fails to generate a minor alarm. [PR1448413](#)
- IPv6 throughput numbers for NAT with HTTP traffic is not at par with IPv4. [PR1449435](#)
- JFLOW: reducing the maximum flow table size when you use **Flex-flow-sizing**. [PR1413513](#)
- The severity of the following error is reduced from fatal to major: **XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR**. [PR1390333](#)
- The **high-cos-queue-threshold** range is changed to [uint 0 .. 90;]. [PR1390424](#)
- The PPPoE subscribers are not able to reconnect after FPC reboots. [PR1397628](#)
- The rpd generates the following core files: **cmgr_if_route_exists_condition_init**, **ctx_handle_node**, **task_reconfigure_complete**. [PR1401396](#)
- Change the default parameters for resource-monitor rtt parameters. [PR1407021](#)
- When you use the inline J-Flow application, the FPC crashes and slows the convergence upon HMC fatal error condition. [PR1407506](#)
- For the initial packet, which is specific to MPC10 and onward, the ICMP redirect s are not seen at the source and packets are sent to the better next hop. [PR1409346](#)

- On MX150, the log severity level changes. [PR1411846](#)
- On platforms running Junos OS Evolved, the redirect IP supports BGP flowspec filters. [PR1413371](#)
- Behavior issues occur with SR-TE Junos telemetry interface sensors when IS-IS sensors are also enabled and the route nexthops are aggregated Ethernet interfaces. [PR1413680](#)
- On PowerPC based MX Series platforms, the DHCP/DHCPv6 subscribers might fail to establish sessions. [PR1414333](#)
- cRPD does not restrict the number of simultaneous JET API sessions. [PR1415802](#)
- The JSU package installation might fail. [PR1417345](#)
- The rpd core files are seen when you restart the rpd or when the logical system is deactivated. [PR1418192](#)
- Changing CAK and CKN multiple times within a short interval (around 5 minutes) sometimes show the security MACsec connection's inbound and outbound channel display with more than one AN active. But on the Packet Forwarding Engine hardware side, the correct AN and SAK is programmed and MKA protocol from both ends transmit correct and latest AN on each hello packet. You should not see any traffic drop due to this display issue. [PR1418448](#)
- The **ROUTING_LOOP_DETECTED** subcode is not generated under **PATHERR_RECV** code when a strict path loop is created for LSP event telemetry notifications. [PR1420763](#)
- The jnxFruState shows value as 10 for Routing Engine instead of 6 in response to .1.3.6.1.4.1.2636.3.1.15.1.8.9.1.0.0. [PR1420906](#)
- MX Series router LNS might fail to forward the traffic on the subscriber access route. [PR1421314](#)
- After the control plane event, a few IPsec tunnels fail to send traffic through the tunnel. [PR1421843](#)
- RSI bloat occurs due to VM host-based log collection. [PR1422354](#)
- The XML output might be not hierarchically structured if you use the **show security group-vpn member ipsec statistics** command. [PR1422496](#)
- The **show system subscriber-management summary** command should include the failure reason for standby disconnect when primary and back Routing Engine memories are not matched. [PR1422976](#)
- Ports might get incorrectly channelized if they are already of 10-Gigabit Ethernet and they are channelized to 10-Gigabit Ethernet again. [PR1423496](#)
- Configuration commit might fail when the file system gets into full state. [PR1423500](#)
- Even when disk-failure-action reboot or disk-failure-action halt are configured, the system does not reboot or halt when disk error is encountered. [PR1424187](#)
- The rpd keeps crashing after changing configuration. [PR1424819](#)
- The mspmand process might crash and restart with a mspmand core file created after doing a commit change to deactivate and activate the service set. [PR1425405](#)
- One hundred percent of CPU usage is seen on route monitor of static routes after the client is disconnected from prpd server. [PR1425559](#)

- On MX204 or MX10003, MPC reboot or Routing Engine mastership switchover might occur. [PR1426120](#)
- Observing NPC core at trinity_rtt_hw_bulk_helper, trinity_rt_delete, rt_entry_delete_msg_proc (rt_params=0x48803bd8) at ../../../../src/pfe/common/applications/route/hal/rt_entry.c:5210. [PR1427825](#)
- On MX Series platforms with PPP configuration, when something abnormal happens such as the user dialup router is abnormally powered off or the keepalive packet is dropped due to network problem, the PPP session ages out. In a rare case, the PPP session does not get deleted, which prevents the new session from being created. So the new session is not able to log in. The PPP traffic might be dropped because of the duplicate-protection feature on the interface. And the IP address of the PPP interface cannot be pinged. [PR1428212](#)
- Incorrect display of MAC/MAC+IP and count values are seen after setting **global-mac-limit** and **global-mac-ip-limit**. [PR1428572](#)
- On MX10003 platform, fabric drops might be seen when two FPCs come online together. [PR1428854](#)
- The aggregated Ethernet interface does not come up after rebooting the FPC or device although the physical member link is up. [PR1429917](#)
- The routers that are configured with protect core might send ipfix sampling packets with the wrong next-hop information. [PR1430244](#)
- Performance degradation is observed for about 20 seconds after the fabric board on MX10008 or MX100016 is taken offline. [PR1430739](#)
- Error might occur when you use a script to load the configuration. [PR1431198](#)
- The l2cpd process might crash and generate a core dump file when interfaces are flapping. [PR1431355](#)
- Dual stack subscriber accounting statistics are not baselined when one stack logs out. [PR1432163](#)
- Traffic might be sent on the standby link of an aggregated Ethernet bundle and get lost with LACP fast-failover enabled. [PR1432449](#)
- After you delete the CLI configuration chassis license bandwidth, the bandwidth value does not default to maximum bandwidth value. [PR1433157](#)
- The rpd generates core files during the route flash when the policy is removed. [PR1434243](#)
- Packet Forwarding Engine memory leak might be seen if MLPPP links are flapped. [PR1434980](#)
- MicroBFD 3x100ms flap is seen upon inserting a QSFP in another port. [PR1435221](#)
- Traffic drops when session key rolls over between primary and fallback nodes for more than 10 times. [PR1435277](#)
- The mc aggregated Ethernet interface might get stuck in the waiting state after a device reboot. [PR1435874](#)
- The local route in the secondary routing table gets stuck in the KRT. [PR1436080](#)

- The ifHCInOctets counter on aggregated Ethernet interface shows the zero value when SNMP MIB walk is executed. [PR1436201](#)
- When you reboot or power off the backup Routing Engine, a trap message is displayed. [PR1436212](#)
- A few static PPP subscribers are stuck in the initialization state permanently and the **Failed to create client session, err=SDB data corrupted** error is seen. [PR1436350](#)
- The subscriber interim statistics might reset to zero and idle-timeout might not work in the MX Series Virtual Chassis setup. [PR1436419](#)
- Not able to reach the router after downgrading from Junos OS Release 18.2-20190513.0 to 18.2R2.6. [PR1436832](#)
- On MPC10, the micro-BFD sessions do not come up in centralized mode. [PR1436937](#)
- Ping fails on logical interfaces with dual tag. [PR1437608](#)
- The CPU utilization on a daemon might be around 100 percent or the backup Routing Engine might crash in race conditions. [PR1437762](#)
- ISSU fails from 19.1R1 legacy Junos release images. [PR1438144](#)
- RPD might generate a core file during router boot up due to file pointer issue because there are two code paths that can close the file. [PR1438597](#)
- On MX Series Virtual Chassis platforms, subscriber flows might not be synchronized between aggregated Ethernet members. [PR1438621](#)
- The syslog server over TCP-based-syslog does not receive carrier-grade NAT logs when data traffic is sent at 10,000 sessions/sec. [PR1438928](#)
- Incorrect values are observed in the **JUNIPER-TIMING-NOTFNS-MIB** table. [PR1439025](#)
- On platforms running Junos OS Evolved, the **show jdaf service cmd statistics/clients** command is not available. [PR1439118](#)
- In an MX Series Virtual Chassis, FPC on Virtual Chassis backup router might reboot. [PR1439170](#)
- Interface-specific filters do not have any effect on MPC10E line cards. [PR1439327](#)
- When a group is applied at non-root level, updating commands inside the group does not update the hierarchies where they are applied. [PR1439805](#)
- IPv6 throughput is not on par with IPv4. [PR1439917](#)
- PRPD flexible tunnel profile queries do not return DMAC when set to all zeros by client. [PR1439940](#)
- The following syslog error message might appear: **UI_SCHEMA_MISMATCH_SEQUENCE: Schema header sequence numbers**. [PR1440141](#)
- On VMware/ESXi in a multiple FPCs chassis, the interfaces assignment is incorrect and some physical interfaces are not visible. [PR1440360](#)
- CoS-related errors are seen and subscribers are not able to get service. [PR1440381](#)

- On MX Series, CPU might hang or interface might stop working on 100-Gigabit Ethernet port. [PR1440526](#)
- In some situations when too many statistics need to be collected from the Packet Forwarding Engine level at the same time, the bulk manager thread of the FPC microkernel level might be continuously busy and cause permanent 100 percent FPC CPU utilization. [PR1440676](#)
- DHCP offer packets toward IRB over LT interface are getting dropped in DHCP relay environment. [PR1440696](#)
- The Layer 2 dynamic VLANs miss when an interface is added to or removed from an aggregated Ethernet bundle. [PR1440872](#)
- When laser receiver power gets -inf , the telemetry value corresponding to -infinity should be equivalent to IEEE 754, which is a single-precision float and the 32-bit value should be 0xff800000. [PR1441015](#)
- New OID is added that calculates the buffer utilization where inactive memory is not considered as free memory. [PR1441680](#)
- Egress stream flush failure and traffic black hole might occur. [PR1441816](#)
- LINX:SNMP trap comes twice for FRU removal in MX10000, with one trap with FRU name as FPC: JNP10K-LC2101 and second with FRU name as FPC @ 1/*/* . [PR1441857](#)
- The packets originating from the IRB interface might get dropped in a VPLS scenario. [PR1442121](#)
- The chassisd is unable to power off a faulty FPC after Routing Engine switchover, leading to chassisd restart loop. [PR1442138](#)
- The operational status of the interface in hardware and software might be out of synchronization in EVPN setup with arp-proxy feature enabled. [PR1442310](#)
- In the enhanced-ip or enhanced-ethernet mode with DCU (destination-class-usage) accounting enabled, MS-DPC might drop all traffic that should egress through aggregated Ethernet interface. [PR1442527](#)
- EVENT UpDown interface logs are partially collected in syslog messages. [PR1442542](#)
- Different formats of the B4 addresses might be observed in the SERVICES_PORT_BLOCK_ALLOC/RELEASE/ACTIVE log messages. [PR1442552](#)
- A few Path Computation Element Protocol (PCEP) logs are marked as error even though they are not an error. The severity of those logs is now marked as INFO. [PR1442598](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- On MX Series platforms, the bbe-smgd might crash. [PR1443109](#)
- The BGP session fails to be establish when you use the firewall filter to de-capsulate BGP packets from the GRE tunnel. [PR1443238](#)
- The kmd process might crash and restart with a kmd core file created if IP of NAT mapping address for IPsec-VPN remote peer is changed. [PR1444183](#)
- MX204: GRE data packets with size greater than the MTU get dropped when sampling is enabled on the egress interface. [PR1444186](#)

- For eventd, you might observe high CPU utilization along with error logs. [PR1444462](#)
- Inline-keepalive might stop working for LNS subscribers if the routing-services statement is enabled. [PR1444696](#)
- MX:EAPoL: MACsec sessions are down with unicast EAPOL destination address. [PR1445052](#)
- Access route might be stuck in bbe-smgd and rpd might not be cleared. [PR1445155](#)
- The CPCDD process continuously generates core files and stops the process in ServicesMgr::ServicesManager::cpcddSmdlInterface::processInputMsg. [PR1445382](#)
- ECMP-FRR might not work for BGP multipath ECMP routes. [PR1445391](#)
- Detached LACP member link gets LACP state as enabled in Packet Forwarding Engine when switchover occurs because of device reboot. [PR1445428](#)
- The 1-Gbps interface on MX204 might stay down after the device reboots. [PR1445508](#)
- Junos OS Release 19.2 group level uses wildcard <*>. [PR1445651](#)
- The Layer 2 ald might crash when FPC restarts. [PR1445720](#)
- The mspmand process might crash if URL filtering is configured and one blacklisted domain name is a substring of another blacklisted domain name in URL filter database file. [PR1445751](#)
- On Ex3400, DOT1XD core file is found at **macsec_update_intf macsec_destroy_ca** directory. [PR1445764](#)
- The jdncpd process might crash after issuing the **show access-security router-advertisement-guard** command. [PR1446034](#)
- When you use a converged CPCD, MX Series router rewrites the HTTPS request with destination-port 80. [PR1446085](#)
- When switchover happens with MX Series router with service interface that has NAT and GR configuration, the static route for NAT never comes up. [PR1446267](#)
- The following rpd core file appears: **task_block_verify(task_io_hook_block, hook),jtask_jthr_endpoint_internal_sanity ,jtask_jthr_endpoint_sanity**. [PR1446320](#)
- Accurate statistics might not include the forwarded packets during the last 2 seconds before subscriber termination. [PR1446546](#)
- NAT service set in certain scale might fail to get programmed. [PR1446931](#)
- ISSU: Core-RMPC3.gz.core.0 and ISSU failure are seen for MPC5. [PR1446993](#)
- The J-Flow version 5 stops working after input rate values are changed. [PR1446996](#)
- Sonet option is enabled for the xe interface. [PR1447487](#)
- DT_BNG: bbe-smgd core file on backup Routing Engine in bbe_ifd_add_vlan (ifd=0x8c3e835, ifl=0xcaf59f18) at **../src/junos/usr.sbin/bbe-svcs/smd/infra/bbe_ifd.c:6374**. [PR1447493](#)
- On MX Series routers, when you use ps interface over redundant logical tunnel in Layer 2 circuit, the pseudowire traffic gets dropped or discarded if **no-control-word** is enabled. [PR1447917](#)

- The rpd process might crash if BGP is activated or deactivated multiple times. [PR1448325](#)
- PCEP: PCE-initiated SR LSP in the first PCE tears down when PCInitiate LSP is brought up and brought down in the second PCE. [PR1448665](#)
- DCD CPU spike is observed after Junos OS upgrade from Junos OS Release 14.2 to Release 16.1. [PR1448858](#)
- Unexpected behavior might occur when you use the **load override** command. [PR1448965](#)
- IPv6 packets might get dropped when vMX acts as a VRRPv3 gateway. [PR1449014](#)
- FPC reboots when PIC 0 is taken offline. [PR1449067](#)
- The DHCP relay feature might not work as expected with **helpers bootp** configured. [PR1449201](#)
- The packets might get dropped when the usage of CPU Core 0 on the host is high. [PR1449289](#)
- There might be an increase in the maximum value of **delegation-cleanup-timeout**. [PR1449468](#)
- Changing the hostname triggers LSP on-change notification and not the adjacency on-change notification. [PR1449837](#)
- The following error message is changed: **Failed to fetch JDM software version from <other_server_full_name>**. If authentication of peer server is not done yet, run **request server authenticate-peer-server from the earlier message: Failed to fetch software version from <other_server_full_name>** to make the error message meaningful. [PR1449871](#)
- On MX Series router running Junos OS enhanced subscriber management feature, no localhost logical interface for rtt 65535 is observed. [PR1450057](#)
- The power that supplies LED on the status panel remains green while one or more PEMs have FAULT LED turned on. [PR1450090](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)
- MoFRR: Issue with MLD plus IGMP scale. [PR1450803](#)
- On VLAN configuration changes with Layer 2 ald, restart might cause kernel synchronization issues and impact forwarding. [PR1450832](#)
- On MPC10E, dcd is unable to clean stale mt- logical interfaces while reloading rosen configuration on the DUT. [PR1450953](#)
- When you use the Standard_D5_v2, which has 16 vCPUs and 56 GB of memory, the deployment fails. [PR1450975](#)
- JNP10000-LC2101 FPC generates **Voltage Tolerance Exceeded** major alarm for EACHIP 2V5 sensors. [PR1451011](#)
- The burst size is not updated when the dynamic profile uses the static traffic control profile. [PR1451033](#)
- SNMP query for IPsec with decrypted or encrypted packets does not fetch the correct values. The following error is observed: **KMD_SNMP_FATAL_ERROR** [PR1451324](#)

- The VFP external static IP configuration is not persistent after rebooting the VFP instance. [PR1451709](#)
- RMPC core files are found after the configuration changes are done on the network for PTP or clock synchronization. [PR1451950](#)
- On MX Series, the dropped packets are seen on MQ/XM-based MPCs, although there is no traffic flowing through the system. [PR1451958](#)
- The mgd might crash when you use the **replace pattern** command. [PR1452136](#)
- On the MX10000 and PTX10000 lines of routers with Routing Engine redundancy configuration enabled, the firmware upgrade for PSU (JNP10000-AC2) and JNP10000-DC2) might fail due to lcmd being disabled by the firmware upgrade command. [PR1452324](#)
- PLL errors might be seen after FPC reboots or restarts. [PR1452604](#)
- On MX10003, MACsec framing errors are seen whenever the sequence number exceeds 2 power 32 with extended packet numbering (XPN). [PR1452851](#)
- Hide the **drop-flow** command under tcp-non-syn configuration. [PR1452902](#)
- On MPC10E, inconsistency between AFT and non-AFT line cards occurs while displaying ldp p2mp traffic-statistics on bud node. [PR1453130](#)
- The values displayed in the output of the **show snmp mib walk jnxTimingNotfnsMIB.3** command are not correct. This MIB table is responsible for timing feature defect or event notification. [PR1453436](#)
- PTP can go out of synchronization due to Layer 2 ald hwdb access failure. [PR1453531](#)
- On MX10003 platform, alarms are not sent to syslog. [PR1453533](#)
- Delay in freeing processed defragment buffers lead to prolonged flow control and might crash. [PR1453811](#)
- The ANCP interface-set QoS adjusts might not be processed. [PR1453826](#)
- The FPC might crash when the severity of error is modified. [PR1453871](#)
- Timestamp is not shown with count option after changing the match condition for the **show <> | mathc <> | count** command. [PR1454387](#)
- On the MX204 platform, the **radius-acct-interim** statistics are not populated for subscribers. Statistics are properly populated in the **radius-acct-stop** packets. [PR1454541](#)
- The 100-Gigabit Ethernet interfaces might not come up again after going down on MPC3E-NG. [PR1454595](#)
- The access request for Layer 2 BSA port up might not be retransmitted if the RADIUS server was unreachable. [PR1454975](#)
- JNS/GNF: CRAFTD syslog fatal errors along with junk characters appear upon startup and exits after four startup attempts. [PR1454985](#)
- JET/JSD RPC tag handling bug. [PR1455426](#)

- Device chooses incorrect source address for locally originated IPv6 packets in routing-instance when destination address is reachable through static route with next-table command. [PR1455893](#)
- Excessive fragmentations of IKE packets might cause failure in the tunnel establishment. [PR1455896](#)
- The BgpRouteInitialize API exits with error code 2. [PR1455967](#)
- The rpd crashes at `__mem_assert func=0x2266f3a "free_jemalloc"`, while adding and deleting the sensors. [PR1456049](#)
- High temperature from the **show chassis environment** output is observed after MPC4E is inserted to slot 5. [PR1456457](#)
- CLI command with invoke-on and display xml rpc results in unexpected multiple RPC commands. [PR1456578](#)
- All the IPsec tunnels might be cleared when the **clear** command is executed for only one IPsec tunnel with specified service-set name. [PR1456749](#)
- The bbe-statsd process might continuously crash if any parameter is set to 0 in the **mx_large.xml** file. [PR1457257](#)
- Default value of 2^{32} replay-window size results in framing errors at an average of one in 2^{32} frames received. [PR1457555](#)
- The chassisd process and all FPCs might restart after Routing Engine switchover. [PR1457657](#)
- The **show subscriber extensive** command incorrectly displays DNS address provided to the DHCP clients. [PR1457949](#)
- The subscriber routes are not cleared from the backup Routing Engine when the session is aborted. [PR1458369](#)
- Traffic black hole or MPC crash might be seen on MPC10E during firewall filter terms change. [PR1458499](#)
- If you use dynamic VoIP VLAN assignment, the correct VoIP VLAN information in LLDP-MED packets might not be sent after you commit. [PR1458559](#)
- The FPC X major errors alarm might be raised after committing the PTP configuration change. [PR1458581](#)
- When you perform delete operations, the gRPC updates on_change does not work. [PR1459038](#)
- After you set interface <> is disabled with QSA, the link still remains up. [PR1459093](#)
- The traffic might be stuck on MS-MPC or MS-MIC with sessions receiving a huge number of affinity packets. [PR1459306](#)
- The following error message might be seen after the chassisd restarts: **create_pseudos: unable to create interface device for pip0 (File exists)** [PR1459373](#)
- The **show ancp subscriber access-aggregation-circuit-id < access aggregation circuit ID>** command displays incomplete output. [PR1459386](#)
- Telemetry streaming of mandatory TLV 'ttl' learned from LLDP neighbor is missing. [PR1459441](#)

- The traffic might be silently dropped or discarded during link recovery in an open Ethernet access ring with ERPS configured. [PR1459446](#)
- In MC-LAG scenario, the traffic destined to VRRP-virtual MAC gets dropped. [PR1459692](#)
- After the DRD auto recovery, the traffic blackholing upon interface flaps. [PR1459698](#)
- CPCDD core file is found at
**ServicesMgr::ServicesManager::cpcddSmdInterface::processServiceNotifyMsg
 ,SmdInterface::cbStateSyncServiceNotifyMsgHandler ,statesync_consumer_poll_new_state_cb.**
[PR1459904](#)
- Initial synchronization for **OpenConfig** event sensors are streamed only from producers supporting event paths. [PR1459927](#)
- The PPTP does not work with destination NAT. [PR1460027](#)
- If **vlan-offload** is configured on the VMX platform, **input-vlan-map** might not work. [PR1460544](#)
- The bbe-smgd generates a core file when all RADIUS servers are unreachable. [PR1461340](#)
- When you receive IPv6 over IPv4 IBGP session, the IPv6 prefix is hidden. [PR1460786](#)
- The repd generates a core file during system startup. [PR1461796](#)
- During the BBE statistics collection and management process, issues with the **bbe-statsd** memory on backup Routing Engine occurs. [PR1461821](#)
- JET RIB API RouteRemove and RouteRemoveMatching RPCs do not work as the first RIB API call. [PR1461974](#)
- The rpd might crash after committing **dynamic-tunnel-anchor-pfe** command. [PR1461980](#)
- The **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed**" message appears when both DIP switches and power switch are turned off. [PR1462065](#)
- The flow stuck and flowd watchdog generate core files while trying to ping DNS server 8.8.8.8 on the internet through DUT configured with NAPT44. [PR1462277](#)
- On MX204, RADIUS interim accounting statistics are not populated. [PR1462325](#)
- The vty remote MAC addresses are not learned with correct age if vty is from a line card without Juniper Trio 5 silicon. [PR1463040](#)
- MAC-learning is broken for vlan-id all scenario. [PR1463078](#)
- The subscribers might not pass traffic after making some changes to the dynamic-profiles filter. [PR1463420](#)
- The MPC2E-NG or MPC3E-NG card with specific MIC might crash after a high rate of interface flaps. [PR1463859](#)
- RPC ALG causes MSPMAND core files when MX Series router is used as a stateful firewall with the MS-MIC or MS-MPC service cards. [PR1464020](#)

- The PPP IPv6CP might fail if the **routing-services** command is enabled. [PR1464415](#)
- The **show task memory detail** command shows incorrect cookie information. [PR1464659](#)
- The PPPoE session goes in to terminated state and the accounting stops for the session that is delayed. [PR1464804](#)
- MPC5E or MPC6E might crash due to internal thread hogging of the CPU. [PR1464820](#)
- DNS sinkhole server results in multiple core files. [PR1466567](#)
- Layer 2 wholesale does not forward all the client requests with stacked VLAN. [PR1467468](#)
- Crypto code might cause high CPU utilization. [PR1467874](#)
- The process rpd might crash after making several changes to the flow-spec routes. [PR1467838](#)

Infrastructure

- The duplex status of management interface might not be updated in the output of **show** command. [PR1427233](#)
- On all Junos OS VM based platforms, FPC might reboot if jlock hog occurs. [PR1439906](#)
- The operations on console might not work if the **system ports console log-out-on-disconnect** command is configured. [PR1433224](#)
- The Routing Engine might go to amnesiac mode an earlier version of Junos OS is installed on an upgraded device. [PR1445151](#)
- The scheduled tasks might not be executed if the **cron** daemon goes down without restarting automatically. [PR1463802](#)

Interfaces and Chassis

- Unrelated aggregated Ethernet interfaces might go down if the committing configuration changes. [PR1409535](#)
- The demux interfaces goes down after changing the MTU of the underlying et interface. [PR1424770](#)
- Mixed link-speed aggregated Ethernet bundle are not able to add new sub-interface successfully. [PR1437929](#)
- Targeted-distribution for static demux interface over aggregate Ethernet interface does not take correct LACP link status into consideration when choosing primary and backup links. [PR1439257](#)
- Mgd processes increases because the mgd processes are not closed properly. [PR1439440](#)
- The cfmd process might crash after a restart on Junos OS Release 17.1R1 and later. [PR1443353](#)
- Unrelated aggregated Ethernet interfaces might go down if changes in the configuration are committed. [PR1409535](#)
- Need enhancement to add or delete a single VLAN in vlan-id-list under interface family bridge. [PR1443536](#)

- ISSU might fail when you upgrade a device that has an aggregated Ethernet bundle with more than 64 logical interfaces. [PR1445040](#)
- The OAM CCM messages are sent with single-tagged VLAN even when configuring with two VLANs. [PR1445926](#)
- Not able to connect to newly installed Routing Engine from other Routing Engines in Routing Engines in MX Series Virtual Chassis. [PR1446418](#)
- Initiating a Routing Engine switchover on VRRP backup router through a CLI command (even protocols vrrp delegate-processing ae-irb) might cause VRRP state for aggregated Ethernet bundle interfaces transitions to the master state, then very shortly afterward to backup again. [PR1447028](#)
- The Layer 2 ald might fail to update composite next hop. [PR1447693](#)
- The ifinfo daemon might crash on the execution of the **show interface extensive** command. [PR1448090](#)
- Dual VRRP mastership might be seen after ungraceful Routing Engine switchover. [PR1450652](#)
- LACP daemon crashed continuously. [PR1450978](#)
- The severity level log might be flooded when the QSFP-100GE-DWDM2 is inserted. [PR1453919](#)
- In the CFM UP MEP over Layer 2 VPN or LAYer 2 circuit service, the CFM UP MEP session might get stuck in the failed state. [PR1454187](#)
- The VRRP traffic loss is longer than 1 second for some backup groups after performing GRES. [PR1454895](#)
- Mismatched MTU value causes the RLT interface to flap. [PR1457460](#)

J-Web

- Some error messages might be seen when you use J-Web. [PR1446081](#)

Layer 2 Features

- LSI interface might not be created, causing remote MACs not to be learned and display of the following error log: **RPD_KRT_Q_RETRIES: ifl iff add: Device busy**. [PR1295664](#)
- VPLS neighbors might stay in the down state after configuration changes in vlan-id. [PR1428862](#)
- Traffic drop might be seen when one MX Series Virtual Chassis member reboots and rejoins the Virtual Chassis. [PR1453430](#)

Layer 2 Ethernet Services

- DHCP request might get dropped in DHCP relay. [PR1435039](#)
- The jdhcpd process might go into infinite loop and cause CPU full utilization. [PR1442222](#)
- On MX10008 or MX10016 platforms, the **dhcp-relay** command might not work. [PR1447323](#)
- Some additional information can be provided in DHCPv6 option 17. This option can be in SOLICIT or REQUEST messages. BNG should relay the information from this option to RADIUS servers in ACCESS

REQUEST message in the attribute 26-207. Before the fix from the PR the information was not relayed. [PR1448100](#)

- PPPoE holding DHCPv6 prefix causes DHCPv6 binding failure due to duplicate prefix. [PR1453464](#)
- DHCP subscriber might not come online after the router reboots. [PR1458150](#)
- DHCP packet might not be processed correctly if DHCP option 82 is configured. [PR1459925](#)
- The ISSU might fail during subscriber in-flight login. [PR1465964](#)

MPLS

- The FPC might be stuck in the Ready state after making a change in the configuration that removes RSVP and triggers FPC restart. [PR1359087](#)
- Static MPLS LSP label might not get installed in MPLS.0 after the link flaps. [PR1457432](#)
- Traffic is silently discarded after the LSP protection link on Huawei transit router goes down. [PR1439251](#)
- Continuous rpd core files are generated at l2ckt_alloc_label, l2ckt_standby_assign_label, and l2ckt_intf_change_process in new backup during GRES in MX2010 box. [PR1427539](#)
- The LDP might withdraw a label for an FEC once the IGP route is inactive in inet.0. [PR1428843](#)
- Dynamic SPRING-TE tunnel creation to LDP (non SR) speaking nodes are is supported. [PR1432791](#)
- Root XML tag in the output is changed from **rsvp-pop-and-fwd-info** to **rsvp-pop-and-fwd-information** to be consistent with the XML tag convention. [PR1365940](#)
- SRLG entry shows unknown after removing it from configuration in **show mpls lsp extensive** or **show mpls srlg** output. [PR1433287](#)
- The P2MP LSP branch traffic might be dropped for a while when the sender PE device performs switchover. [PR1435014](#)
- The flow label is not pushed when **chained-composite-next-hop ingress l2ckt/l2vpn** is enabled. [PR1439453](#)
- LSI interface Layer 2 Virtual Chassis goes down on one router in VPLS domain through the MPLS path is still available in inet.3. Reason shows **as mpls label out of range**. [PR1442495](#)
- The backup LSP path messages are rejected if the bypass tunnel path is an inter-area LSP. [PR1442789](#)
- RSVP path message with long refresh interval is dropped between devices running Junos OS releases earlier than Release 16.1 and devices running versions later than Release 16.1. [PR1443811](#)
- TRUE POC: rpd core files are generated with SNMP polling. [PR1457681](#)
- P2MP LSP might get stuck in the down state after link flaps. [PR1444111](#)
- The rpd memory leak might be seen when the inter domain RSVP LSP is in the down state. [PR1445024](#)
- Traffic might be silently dropped or discarded if two consecutive PLRs along the LSP performs local repair simultaneously under certain misconfigured conditions. [PR1445994](#)

- The transit packets might be dropped if an LSP is added or changed on an MX Series or PTX Series device. [PR1447170](#)
- Traffic drop might be seen after traceoption configuration is committed in RSVP P2MP. [PR1447480](#)
- The rpd generated a core file at **ted_delete_abstract_hop (instance=0x75d33c0, hop_name=< optimized out>)** during abstract-hop testing. [PR1448769](#)
- The LDP route timer resets when committing unrelated configuration changes. [PR1451157](#)
- All LDP adjacencies flap after changing LDP preference. [PR1459301](#)
- The previously configured credibility preference is not considered by CSPF even though the configuration has been deleted or changed to prefer another protocol in the traffic engineering database. [PR1460283](#)
- High CPU usage and rpd core file might be observed if **ldp track-igp-metric** is configured and IGP metric is changed. [PR1460292](#)
- MPLS trace route does not trace the SRUDP tunnel ingress router. [PR1460516](#)

Network Address Translation (NAT)

- The nsd process might crash during SNMP query for deterministic NAT pool information. [PR1436775](#)

Network Management and Monitoring

- MX10000 reports the **jail socket** errors message. [PR1442176](#)
- The **Wrong Type** error message might be seen for the hrProcessorFrwID object. [PR1446675](#)

Platform and Infrastructure

- On all the EX9200 line of switches, MX Series routers, and T4000, LACP DDoS policer is incorrectly triggered by other protocols traffic. [PR1409626](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- Packet drops, replication failure, or ksyncd crash might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- The RPM udp-ping probe does not work in multiple routing instance scenario. [PR1442157](#)
- With CNH for 6PE, MPLS EXP rewrite rule for non-VPN IPv4 over MPLS traffic might not work. [PR1430878](#)
- Traffic from the same physical interface cannot be forwarded. [PR1434933](#)
- The BGP session might flap after Routing Engine switchover is done simultaneously on both boxes of BGP peer in scaled BGP session setup. [PR1437257](#)
- GRE traffic might get dropped if the terminating routing-instance name contains dots. [PR1437872](#)
- ARP resolution might fail after ARP HOLD net hops are added and deleted continuously. [PR1442815](#)
- Some duplicate flowtap filters are programmed after the restart of dynamic-flow-capture. [PR1442868](#)

- When host-bound packet is received in MAP-E BR router, service interface statistics counter shows incorrect number of bytes. [PR1443204](#)
- Packets drop due to missing destination MAC address in the Packet Forwarding Engine. [PR1445191](#)
- Python op scripts are executed as user nobody if started from NETCONF session, not as logged in user, resulting in failing PyEZ connection to the device. [PR1445917](#)
- On certain MPC line cards cm errors need to be reclassified. [PR1449427](#)
- Some hosts behind unnumbered interface are unreachable after the router or FPC restarts. [PR1449615](#)
- FPC might reboot with vmcore due to memory leak. [PR1449664](#)
- The DF flag BGP packets are dropped over MPLS LSP path. [PR1449929](#)
- REST API process becomes non-responsive when a number of requests come in at a high rate. [PR1449987](#)
- In EVPN-VXLAN scenario, sometimes host-generated packets gets dropped as hitting reject route in Packet Forwarding Engine. [PR1451559](#)
- The Routing Engine originated IPv6 packets might be dropped when interface-group rule is configured under IPv6 filter [PR1453649](#)
- Multicast traffic loss occurs in rare case in a seamless MPLS with MVPN configuration is observed. [PR1456905](#)
- Port mirroring does not occur with VPLS. [PR1458856](#)
- DDoS-protection does not stop logging when remote tracing is enabled. [PR1459605](#)
- Traceroute initiated from PE device does not show the tunnel endpoint hop in the output. [PR1461441](#)
- CLI configuration flag **version-03** must be optional. [PR1462186](#)

Routing Policy and Firewall Filters

- The rib-group might not process the exported route correctly. [PR1450123](#)
- Routes resolution might be inconsistent if any route resolves over the multipath route. [PR1453439](#)
- The rpd might crash after the Routing Engine switch overs when prefix-list is configured. [PR1451025](#)

Routing Protocols

- The rpd crashes in Junos OS Release 16.1 or later during BGP convergence. [PR1351639](#)
- The rpd process might crash with BGP multipath and damping configured. [PR1472671](#)
- Need to install all possible next hops for OSPF network LSAs. [PR1463535](#)
- The **other querier present interval** timer cannot be changed in a IGMP or MLD snooping scenario. [PR1461590](#)
- BGP peers might flap if the parameter of hold-time is set as small. [PR1466709](#)
- The rpd might crash under a rare condition if GR helper mode is triggered. [PR1382892](#)

- The rpd crashes in Junos OS 16.1 or later during BGP convergence. [PR1351639](#)
- BFD link failure detection of the broken path gets delayed when IGP link-state update is received from the same peer through an alternative path. [PR1410021](#)
- BGP might become stuck in the Idle state when the peer triggers a GR restart event. [PR1412538](#)
- BGP might get stuck in the Idle state when the peer triggers a GR restart event. [PR1412538](#)
- TI-LFA cannot find backup path when IS-IS overLoad bit is set on computing. node [PR1412923](#)
- Per-prefix LFA might not work as expected where the last hop needs to be protected on the penultimate node. [PR1432615](#)
- Unsupported configuration (EPE with dynamic-next-hop GRE tunnels) continuously rpd to generat core files. [PR1431536](#)
- The **show isis adjacency extensive** output does not contain the state transition details. [PR1432398](#)
- The next-hop of IPv6 route remains empty when a new IS-IS link comes up. [PR1430581](#)
- With SR enabled, 6PE next hop is not installed. [PR1435298](#)
- Clearing BGP neighbors takes longer time to delete routes. [PR1435466](#)
- Wrong next hop might be seen when BGP PIC edge is enabled. [PR1437108](#)
- The rpd might crash in case multipath is enabled, as BGP multipath teardown is called for secondary route even though secondary routes are considered for multipath. [PR1437837](#)
- The backup Routing Engine might go out of synchronization if you clear BGP sessions on the master Routing Engine. [PR1439620](#)
- Removing SSH Protocol version 1 from configuration. [PR1440476](#)
- RIP routes might be discarded by Juniper device over a /31 subnet interface. [PR1441452](#)
- The rpd might crash with SR-TE configuration change. [PR1442952](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- The rpd crash might be seen after configuring OSPF nssa area-range and summaries. [PR1444728](#)
- The rpd might crash in OSPF scenario due to invalid memory access. [PR1445078](#)
- The SSH login might fail if a user account exists in both the local database and RADIUS/TACACS+. [PR1454177](#)
- MoFRR with MLDP inband signaling is not working. [PR1454199](#)
- BRP: RPC call is not available for **show bgp output-scheduler**. [PR1445854](#)
- The BGP route prefixes are not being advertised to the peer. [PR1446383](#)
- The as-external route might not work in OSPF overload scenario for VRF instance. [PR1446437](#)
- The rpd uses full CPU utilization due to incorrect path selection. [PR1446861](#)

- The multicast traffic might be dropped in PIM with BGP PIC setup. [PR1447187](#)
- The rpd crashes and commit fails when trying to commit configuration changes. [PR1447595](#)
- On the MX2000 and PTX10000 lines of devices , Layer 3 VPN PE-CE link protection exhibits unexpected behavior. [PR1447601](#)
- Junos OS BFD sessions with authentication flaps occurs after sometime. [PR1448649](#)
- The connection between ppmd (Routing Engine) and ppman (FPC) might get lost due to session time out. [PR1448670](#)
- The BGP routes might fail to be installed in routing instance if the **from next-hop** policy match condition is used in the VRF import policy. [PR1449458](#)
- SPRING-LDP interoperability issues are observed with colocated SRMS+SR-client+LDP-stitching. [PR1452956](#)
- The rpd scheduler slip for BGP GR might be up to 120 second after the peer goes down. [PR1454198](#)
- The rpd memory might leak in a certain MSDP scenario. [PR1454244](#)
- Permanent rpd core files are seen with BGP configuration option **optimal-route-reflection** set. [PR1454803](#)
- Rpd might crash when multipath is in use. [PR1454951](#)
- The rpd might crash continuously due to memory corruption in IS-IS setup. [PR1455432](#)
- Prefix SID conflict might be observed in IS-IS. [PR1455994](#)
- Packet drop and CPU spike on Routing Engine might be seen in certain conditions if **labeled-unicast protection** is enabled for a CsC-VRF peer. [PR1456260](#)
- Rpd core file is seen at **rt_nhn_tree_stop,rt_table_tree_free_family,bgp_sync_free_tsp** after deactivating protocols. [PR1457358](#)
- The rpd might crash when OSPF router-id gets changed for NSSA with area-range configured. [PR1459080](#)
- The rpd memory leak might be observed on backup Routing Engine due to BGP flap. [PR1459384](#)
- Rpd scheduler slips might be seen on RPKI route validation enabled BGP peering router in a scaled setup. [PR1461602](#)
- Rpd core file is seen with BMP configured and BGP peer flapping. [PR1462441](#)
- IS-IS IPv6 multi-topology routes might flap every time when there is an unrelated commit under protocol stanza. [PR1463650](#)
- The rpd might crash if both BGP add-path and BGP multipath are enabled. [PR1463673](#)
- MX80 EVPN-VXLAN RT5 does not work properly and ip-prefix-routes are not reachable. [PR1466602](#)

Services Applications

- The kmd process might crash when DPD time outs for some IKEv2 SAs occurs. [PR1434521](#)
- On platforms running Junos OS Evolved, the **show ipsec security-associations** command throws an error. [PR1442161](#)
- Phase 1 SA is migrated to new remote IP because of the source-address translation for the static NAT tunnel. [PR1477181](#)
- Output of the **show subscriber user-name** command on LTS shows only one session instead of two. [PR1446572](#)
- The jl2tpd process might crash during the restart procedure. [PR1461335](#)
- BGP multipath does not work for MT on cRPD. [PR1467091](#)

Subscriber Access Management

- Subscriber filtering for general authentication services traceoptions could report debug messages for other users. [PR1431614](#)
- Subscriber deactivation might get stuck in the terminated state. [PR1437042](#)
- Test aaa ppp, output enhancement. [PR1444438](#)
- On MX Series platforms, there might be a false error for SAE policy activation or deactivation failure. [PR1447632](#)
- Subscribers login fails when PCRF server is unreachable. [PR1449064](#)
- The authd crashes on backup Routing Engine during execution of the slax script, running < get-jsrc-counters> RPC call. [PR1458185](#)
- DHCPv6 subscribers might be stuck in a state after the authd process crashes. [PR1460578](#)
- Problem with linked-pool-aggregation after attempting to delete a pool in the middle of the chain. [PR1465253](#)

User Interface and Configuration

- The **show chassis hardware satellite** command is not available in Junos OS Release 17.3. [PR1388252](#)
- Changing nested apply-groups does not occur. [PR1427962](#)
- In the Juno OS Fusion environment, the **show chassis hardware satellite** command is not available on Junos OS Release 17.3. [PR1388252](#)

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the Layer 2 circuit flaps repeatedly. [PR1282875](#)
- The rpd core file is seen at **rtbit_reset**, **rte_tgtexport_rth**. [PR1379621](#)
- The rpd crash might be seen if Layer 2 circuit or local switching connections flap continuously. [PR1418870](#)

- P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. [PR1432434](#)
- The resumed multicast traffic for certain groups might be stopped in overlapping MVPN scenario. [PR1441099](#)
- Result of the **show task replication** command shows MVPN as InProgress when the active master Routing Engine is forcibly removed and NSR are enabled. [PR1441292](#)
- Memory leak might happen if PIM messages are received over an MDT (mt- interface) in Draft-Rosen MVPN scenario. [PR1442054](#)
- The rpd process might crash due to memory leak in **MVPN RPF Src PE** block. [PR1460625](#)
- The Layer 2 circuit displays MM status which might cause traffic loss. [PR1462583](#)

SEE ALSO

[What's New | 87](#)

[What's Changed | 110](#)

[Known Limitations | 116](#)

[Open Issues | 119](#)

[Documentation Updates | 168](#)

[Migration, Upgrade, and Downgrade Instructions | 169](#)

Documentation Updates

IN THIS SECTION

- [Advanced Subscriber Management Provider | 169](#)
- [Feature Guides Are Renamed As User Guides | 169](#)

This section lists the errata and changes in Junos OS Release 19.4R2 documentation for MX Series.

Advanced Subscriber Management Provider

- The Broadband Subscriber Services User Guide incorrectly stated that for Routing Engine-based, converged HTTP redirect services, a CPCD service rule can include both a redirect term and a rewrite term. It also incorrectly stated that you can include separate rewrite and redirect rules in the same service profile.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

[What's New | 87](#)

[What's Changed | 110](#)

[Known Limitations | 116](#)

[Open Issues | 119](#)

[Resolved Issues | 135](#)

[Migration, Upgrade, and Downgrade Instructions | 169](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.4 | 170](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 170](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 173](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 175](#)
- [Upgrading a Router with Redundant Routing Engines | 175](#)
- [Downgrading from Release 19.4 | 176](#)

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

Starting in Junos OS 17.4R1 release, FreeBSD 11.x is the underlying OS for all Junos OS platforms which were previously running on FreeBSD 10.x based Junos OS. FreeBSD 11.x does not introduce any new Junos OS related modifications or features but is the latest version of FreeBSD.

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.x-based Junos OS	FreeBSD 11.x-based Junos OS
MX5,MX10, MX40,MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 19.4

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files) might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [Installation and Upgrade Guide](#).

For more information about the installation process, see [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).

Procedure to Upgrade to FreeBSD 11.x based Junos OS

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install FreeBSD 11.x based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.

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.

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-32-19.4R2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-19.4R2.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-19.4R2.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.4R2.9-limited.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname**

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.x) to Junos OS (FreeBSD 11.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 11.x, and Junos OS (FreeBSD 6.x) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

Use the **reboot** command to reboot the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 19.4R2, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following MX Series routers:
 - MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008

[See <https://kb.juniper.net/TSB17603>.]

NOTE: After you install a Junos OS Release 19.4 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. See the VM Host Software Administrative Commands in the [Installation and Upgrade Guide](#).

Procedure to Upgrade to FreeBSD 6.x based Junos OS

Products impacted: MX5, MX10, MX40, MX80, MX104.

To download and install FreeBSD 6.x based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.

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.

- All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/jinstall-ppc-19.4R2.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot  
source/jinstall-ppc-19.4R2.9-limited-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname**

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Use the **reboot** command to reboot the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 19.4 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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>.

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform the following Junos OS installation on each Routing Engine separately to avoid disrupting network operation:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine, and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.

3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Downgrading from Release 19.4

To downgrade from Release 19.4 to another supported release, follow the procedure for upgrading, but replace the 19.4 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the [Installation and Upgrade Guide](#).

SEE ALSO

What's New 87
What's Changed 110
Known Limitations 116
Resolved Issues 135
Open Issues 119
Documentation Updates 168

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 177](#)
- [What's Changed | 179](#)
- [Known Limitations | 180](#)
- [Open Issues | 181](#)
- [Resolved Issues | 183](#)

- Documentation Updates | 188
- Migration, Upgrade, and Downgrade Instructions | 188

These release notes accompany Junos OS Release 19.4R2 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os

What's New

IN THIS SECTION

- General routing | 178
- Hardware | 178
- Architecture | 178

Learn about new features introduced in the Junos OS main and maintenance releases for NFX Series devices.

General routing

- **Support for MAP-E customer edge encapsulation and decapsulation (NFX Series)**—Starting in Junos OS release 19.4R1, Mapping of Address and Port with Encapsulation (MAP-E) customer edge (CE) encapsulation and decapsulation are supported on NFX Series devices. MAP-E is an IPV6 transition technique that encapsulates an IPv4 packet in an IPv6 and carries the packet over IPv4-over-IPv6 tunnel from MAP-E CE devices to the MAP-E provider edge (PE) devices (also called as border relay [BR] devices) through an IPv6 routing topology, where the packet is de-tunneled for further processing.

MAP-E uses network address port translation (NAPT) features for restricting transport protocol ports, Internet Control Message Protocol (ICMP) identifiers, and fragment identifiers to the configured port sets. Existing NAPT feature is enhanced to add this capability.

[See [How to Configure the NFX150](#).]

[See [How to Configure the NFX250 NextGen](#).]

Hardware

- **NFX350 Platform**— With Junos OS Release 19.4R1, the NFX portfolio introduces the NFX350 Network Services Platform, which is a secure, automated, software-driven customer premises equipment (CPE) platform that delivers virtualized network and security services on demand. The NFX350 is part of the Juniper Cloud CPE solution, which leverages Network Functions Virtualization (NFV). The NFX350 platform completes the uCPE portfolio to provide end-to-end platforms for medium, large, and extra-large deployments. In addition to IPsec, Layer 2 features, and SD-WAN functionality, the NFX350 provides features such as LAN or WAN isolation, software and hardware resiliency, redundant power supply, and serial over LAN. The NFX350 device supports two external SSD and LTE expansion module.

The NFX350 devices are available in the following variants:

- **NFX350-S1**—Rack-mount model with 8-core Intel Skylake D-2146NT CPU, 100-GB SSD, 32-GB RAM, eight 1-Gigabit Ethernet RJ-45 LAN ports, and eight 10-Gigabit Ethernet SFP+ WAN ports.
- **NFX350-S2**—Rack-mount model with 12-core Intel Skylake D-2166NT CPU, 100-GB SSD, 64-GB RAM, eight 1-Gigabit Ethernet RJ-45 LAN ports, and eight 10-Gigabit Ethernet SFP+ WAN ports.
- **NFX350-S3**—Rack-mount model with 16-core Intel Skylake D-2187NT CPU, 100-GB SSD, 128-GB RAM, eight 1-Gigabit Ethernet RJ-45 LAN ports, and eight 10-Gigabit Ethernet SFP+ WAN ports.

[See [NFX350 Hardware Guide](#).]

[See [How to Configure the NFX350](#).]

Architecture

- **NFX350 Architecture**—The NFX350 architecture enables unified management of its components through the Junos Control Plane (JCP). It supports the following modes to effectively manage system resources:

- Throughput mode—Provides maximum resources (CPU and memory) for Junos software. The default mode is throughput mode.
- Hybrid mode—Provides a balanced distribution of resources between the Junos software and third-party VNFs.
- Compute mode—Provides minimal resources for Junos software and maximum resources for third-party VNFs

[See [NFX350 Hardware Guide](#).]

[See [How to Configure the NFX350](#).]

SEE ALSO

[What's Changed | 179](#)

[Known Limitations | 180](#)

[Open Issues | 181](#)

[Resolved Issues | 183](#)

[Documentation Updates | 188](#)

[Migration, Upgrade, and Downgrade Instructions | 188](#)

What's Changed

IN THIS SECTION

- [What's Changed in 19.4R2 Release | 179](#)

- [What's Changed in 19.4R1 Release | 180](#)

Learn about what changed in Junos OS main and maintenance releases for NFX Series devices.

What's Changed in 19.4R2 Release

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 19.4R2 for NFX Series devices.

What's Changed in 19.4R1 Release

System Logging

- **Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

[What's New | 177](#)[Known Limitations | 180](#)[Open Issues | 181](#)[Resolved Issues | 183](#)[Documentation Updates | 188](#)[Migration, Upgrade, and Downgrade Instructions | 188](#)

Known Limitations

IN THIS SECTION

- [Interfaces | 181](#)

- [Platform and Infrastructure | 181](#)

Learn about known limitations in this release for NFX Series devices. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Interfaces

- On an NFX150 device running Junos software release 19.4R1 or later, whenever there is no vmhost mapping configuration for a specific heth interface, then the **show interfaces extensive heth-<x>-<y>** command output displays the **Mapped to** status as **None** for **VF Number 0**. [PR1484388](#)
- On NFX150 devices, the link does not come up if a 1-Gigabit SFP transceiver is connected from heth-0-4 and heth-0-5 to a peer device. As a workaround, disable the auto-negotiation for the interface connected to the NFX150 on the remote device. [PR1428020](#)
- On NFX350 devices, when there is any change in vmhost interface mapping for an FPC, that FPC is restarted to activate the changed mapping. [PR1471694](#)

Platform and Infrastructure

- On NFX150 devices, random RPM probe losses are noticed if the probe packets are fragmented because the data-size is larger than the inet MTU. [PR1447082](#)

SEE ALSO

What's New 177
What's Changed 179
Open Issues 181
Resolved Issues 183
Documentation Updates 188
Migration, Upgrade, and Downgrade Instructions 188

Open Issues

IN THIS SECTION

- [Interfaces | 182](#)
- [Platform and Infrastructure | 182](#)
- [Virtual Network Functions \(VNFs\) | 183](#)

Learn about open issues in this release for NFX Series devices. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Interfaces

- When you issue a **show interface** command on NFX150 devices to check the interface details, the system will not check whether the interface name provided is valid or invalid. The system will not generate an error message if the interface name is invalid. [PR1306191](#)
- When a DHCP server assigns a conflicting IP address to the NFX device interfaces, the NFX device will not send a **DHCP DECLINE** message in response. [PR1398935](#)
- On NFX150 and NFX250 NextGen devices, when you add, modify, or delete a VNF interface that is mapped to an L2 or L3 data plane, kernel traces might be observed on the NFX Series device console. [PR1435361](#)
- Only ge-1/0/1 is mapped to OVS by default and 8 logical interfaces are created on this. [PR1452743](#)
- On NFX350 devices, sxe interfaces take more time to come online after fpc0 restart via CLI or after device reboot. [PR1465436](#)

Platform and Infrastructure

- If you plug an unsupported SFP-T transceiver into an NFX150 device and reboot the device, the FPC1 WAN port does not come online. [PR1411851](#)
- Jumbo frames are not supported through OVS on an NFX250 device. [PR1420630](#)
- On NFX250 devices, Virtual Port Peer (VPP) is not running on dual CPE and occasionally on single CPE. [PR1461238](#)
- On NFX350 devices, secure boot is not enabled after BIOS recovery using SSD and USB. [PR1480165](#)
- On NFX350 devices, an srpxfe core file is generated when VF mapping to srpxfe changes. When mapping the backplane's NIC changes for FPC1 to a VF, the srpxfe restarts. In NFX350 devices, the internal NICs are Intel NICs and the DPDK library in srpxfe is unable to handle the PF reset event generated during the remapping. This causes the srpxfe to crash just before the restart. There is no impact on functionality as a result of this issue; however, graceful restart doesn't happen, and instead the srpxfe generates a core file. [PR1469201](#)
- With SRX1500 device used as HUB and NFX350 device as SPOKE, **IPSEC replay-errors** are seen with HTTP traffic when the AppQoE passive probing is enabled. As a workaround, use SRX4200 as HUB. [PR1461068](#)
- When the AppQoe passive probing is enabled on NFX350 devices, packet drops are seen with the **First path drop: Policy check failed** message. [PR1473810](#)

- On NFX350 devices, the LCM peer connection un-stable alarm is raised during the device reboot, and the alarm clears after 2-3 minutes. [PR1473531](#)
- After you upgrade JDM image from Junos OS Release D497.1 to Junos OS Release 18.4R3-S2, tunnels are down in Gateway Router (GWR). [PR1507165](#)

Virtual Network Functions (VNFs)

- On NFX350 devices, management daemon (MGD) is broken when cross-connect is committed before VNF is created. As a workaround, create VNF before creating cross-connect. [PR1480740](#)

SEE ALSO

[What's New | 177](#)

[What's Changed | 179](#)

[Known Limitations | 180](#)

[Resolved Issues | 183](#)

[Documentation Updates | 188](#)

[Migration, Upgrade, and Downgrade Instructions | 188](#)

Resolved Issues

IN THIS SECTION

● [Resolved Issues: 19.4R2 | 184](#)

● [Resolved Issues: 19.4R1 | 185](#)

Learn which issues were resolved in Junos OS main release and the maintenance releases for NFX Series devices.

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

Resolved Issues: 19.4R2

Interfaces

- On NFX350 devices, if any xSTP protocol is enabled on all interfaces, it gets enabled on L3 interfaces, which are enabled with vlan-tagging or flexible-vlan-tagging. This results in blocking the SXE interface. [PR1475854](#)
- Management ports are not disabled with link disable command on NFX150-S1 devices. [PR1442064](#)
- On NFX Series devices, ping is not working between the cross-connected interfaces with interface deny-forwarding configuration. [PR1442173](#)
- On NFX150 and NFX250 NextGen devices, when two srxpfe interfaces are mapped to the same physical interface, data packets received on physical NIC from external device are still sent to the old VF mapping instead of new mapping. [PR1448595](#)
- On NFX150 devices, the heth-0-4 and heth-0-5 ports do not detect traffic when you try to activate the ports by plugging or unplugging the cable. [PR1449278](#)
- On NFX Series devices, the static MAC address is replaced by random MAC address. [PR1458554](#)
- On NFX150 devices running Junos OS Release 19.3R1, the **vmhost virtualization-options** command is not working as expected for heth to ge interface mapping. [PR1459885](#)
- On NFX250 NextGen devices, the **monitor interface traffic** command might not display the **pps** output for SXE and physical interfaces. [PR1464376](#)
- When traffic goes through vSRX3.0 platforms, core-dump files are generated and traffic is dropped. This issue might result in Packet Forwarding Engine being inactive and all interfaces being down. [PR1465132](#)
- On NFX Series devices, the GRE tunnel interface (gr-1/0/0) may not appear if the **clear-dont-fragment-bit** option is configured for the GRE interface. [PR1472029](#)

Mapping of Address and Port with Encapsulation (MAP-E)

- On NFX Series devices, IP identification (IP ID) is not changed after MAP-E NAT44 is performed on fragment packets when the packets reach the customer edge (CE) device. [PR1478037](#)

Platform and Infrastructure

- On NFX Series devices, if there are any conditional groups, the l2cpd process might crash and generate a core dump when interfaces are flapping and the lldp neighbors are available. It might cause the dot1x process to fail and all the ports have a short interruption at the time of process crash. As a workaround, delete the conditional group in the device. [PR1431355](#)
- On NFX350 devices, if you execute the **show vmhost mode** command multiple times, JDM may crash and cause the show commands to stop working. [PR1474220](#)
- After a power outage, JDMD is not responsive because the **/etc/hosts** file is being corrupted. [PR1477151](#)

- Coredumps on NFX250 while adding the second LAN subnet. [PR1490077](#)
- AppQoS is sending active probing packets for the deleted **active-probe-params** option. [PR1492208](#)
- On NFX250 NexGen devices, the **request vmhost power-off** command reboots the device instead of powering off. [PR1493062](#)
- Package files are lost after you upgrade the software image from Junos OS Release D497.1 to Junos OS release 18.4R3.3 on NFX250 devices. [PR1493711](#)

Virtual Network Function (VNF)

- No error is displayed for **native-vlan-id** option that is configured on an access VNF interface though the commit fails. [PR1438854](#)
- On NFX350 devices, VNF instantiation does not work properly. [PR1478456](#)

Resolved Issues: 19.4R1

Class of Service

- On NFX Series devices, when CoS rewrite rule is configured for st0 interface, the CoS value will not take effect on corresponding forwarding class. It causes the CoS not to work as expected. This issue has traffic impact. [PR1439401](#)

High Availability

- On an NFX150 high availability chassis cluster, the host logs updated in the system log messages might not show the correct time stamp. As a workaround, convert the UTC time stamp to local time zone. [PR1394778](#)

Interfaces

- When you transition NFX150 devices from PPPoE configuration to non-PPPoE configuration in a non-promiscuous mode, the interface hangs without any traffic flow. [PR1409475](#)
- The limit on maximum OVS interfaces is restored to the originally defined limit of 25 for backward compatibility. As a workaround, reduce the number of OVS interfaces in the configuration to 20 or fewer. [PR1439950](#)
- On NFX150 and NFX250 NextGen devices, cross-connect stays down even if all linked interfaces are up. [PR1443465](#)

Layer 2 Ethernet Services

- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case the subscriber may need more time to get IP address assigned. The subscriber may remain in this state until it's lease expires if it had previously bound with the address in the option 50. [PR1435039](#)

Platform and Infrastructure

- On NFX Series devices, the HTTP traffic flow is created with a different routing instance when an APBR profile is configured with category and application in the same profile. [PR1447757](#)
- When applying firewall filters on lo0.0 on an NFX250 NextGen device, FPC0 disappears. [PR1448246](#)
- On NFX150 devices, the **show security dynamic-address** command does not work for port 3. [PR1448594](#)
- Half duplex configuration on 1G ports is not supported when autonegotiation is disabled. [PR1453911](#)
- Informational log message, **LIBCOS_COS_RETRIEVE_FROM_PVIDB: feature cos_fc_defaults num elems 4 rc 0** is displayed on the console when you commit after you configure AppQoS rule set. [PR1457328](#)
- REST API process will get non responsive when a number of requests start coming at a high rate. [PR1449987](#)
- Packet drops, replication failure or ksyncd crashes might be seen on the logical system of a Junos OS device after Routing Engine switchover. [PR1427842](#)
- After upgrading the NFX Series devices to Junos OS Release 19.2R2-S1.4, the following commit warning is seen even though there is no configuration change under the **forwarding-options vxlan-overlay-load-balance** option:

```
# commit and-quit
re0:
[edit]
  'forwarding-options'
    warning: vxlan-overlay-load-balance configuration for forwarding options has
    been changed. A system reboot is mandatory. Please reboot *ALL* routing engines
    NOW. Continuing without a reboot might result in unexpected system behavior.
    configuration check succeeds
re1:
configuration check succeeds
commit complete
re0:
commit complete
Exiting configuration mode
```

[PR1459833](#)

Routing Protocols

- On NFX Series devices, changing the **other querier present interval** timer is not working on IGMP or MLD snooping device in the existing Bridge Domain (BD) or Listener Domain (LD). As a workaround, deactivate or activate the IGMP snooping via configuration or run the **restart multicast-snooping** command. [PR1461590](#)

SNMP

- On NFX150 devices, SNMP does not work for the following commands:

- **show snmp mib walk jnxIpSecTunMonOutEncryptedBytes**
- **show snmp mib walk jnxIpSecTunMonOutEncryptedPkts**
- **show snmp mib walk jnxIpSecTunMonInDecryptedBytes**
- **show snmp mib walk jnxIpSecTunMonInDecryptedPkts**
- **show snmp mib walk jnxIpSecTunMonLocalGwAddr**
- **show snmp mib walk jnxIpSecTunMonLocalGwAddrType**

[PR1386894](#)

- Version compare in phc may fail causing the phc to download the same image. [PR1453535](#)

Virtual Network Functions (VNFs)

- On NFX150 devices with VNFs configured, when the VNF interfaces are moved from default OVS bridge to custom OVS bridge, there will be duplicate VNF host entries in the /etc/hosts file on JDM. [PR1434679](#)
- On NFX150 devices, when you need to change the vmhost mappings of a particular NIC or NICs, you must delete the existing vmhost mapping and commit the configuration. Now you can configure the new mappings for the respective NICs. You cannot change the NIC vmhost mappings in the same commit to delete and add a new mapping to the heth NICs. [PR1450147](#)
- NFX250 devices do not allow *jdm* (case-insensitive) as a VNF name. You can use *jdm* as a part of the name. For example, *jdm123*, *abcJDM*, *abcJDM123* are valid VNF names, whereas, *jdm*, *JDM*, *Jdm*, *JDm* are not valid VNF names. [PR1463963](#)

SEE ALSO

[What's New | 177](#)

[What's Changed | 179](#)

[Known Limitations | 180](#)

[Open Issues | 181](#)

[Documentation Updates | 188](#)

[Migration, Upgrade, and Downgrade Instructions | 188](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 188](#)

This section lists the errata and changes in Junos OS Release 19.4R1 documentation for the NFX Series.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

[What's New | 177](#)

[What's Changed | 179](#)

[Known Limitations | 180](#)

[Open Issues | 181](#)

[Resolved Issues | 183](#)

[Migration, Upgrade, and Downgrade Instructions | 188](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 189](#)
- [Basic Procedure for Upgrading to Release 19.4 | 189](#)

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

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.

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

Basic Procedure for Upgrading to Release 19.4

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

NOTE: NFX150, NFX250 NextGen, and NFX350 devices run VMhost supported routing engine, and should follow the [VMhost Support on Routing Engines](#) upgrade procedure.

NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the device, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the device. For more information, see the [Software Installation and Upgrade Guide](#).

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.

To download and install Junos OS Release 19.4R2:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the **Software** tab.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the Download Software page.
5. In the Install Package section of the Software tab, select the software package for the release.
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. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the device or to your internal software distribution site.
10. Install the new package on the device.

SEE ALSO

[What's New | 177](#)

[What's Changed | 179](#)

[Known Limitations | 180](#)

[Open Issues | 181](#)

[Resolved Issues | 183](#)

[Documentation Updates | 188](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- What's New | [191](#)
- What's Changed | [199](#)
- Known Limitations | [202](#)
- Open Issues | [203](#)
- Resolved Issues | [206](#)
- Documentation Updates | [208](#)
- Migration, Upgrade, and Downgrade Instructions | [209](#)

These release notes accompany Junos OS Release 19.4R2 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- What's New in 19.4R2 | [192](#)
- What's New in 19.4R1 | [192](#)

Learn about new features introduced in this release for PTX Series routers.

What's New in 19.4R2

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 19.4R2.

What's New in 19.4R1

General Routing

- **Optimized BGP peer reestablishment (MX Series, PTX Series, and QFX Series)**—Starting with Junos OS Release 19.4R1, BGP peers in different groups can close in parallel. The connect/retry algorithm makes more frequent attempts to reestablish BGP peers, which reduces downtime. The connect/retry algorithm makes 16 attempts instead of 5 to reestablish BGP peers in the first 256 seconds after they go down. Peers can reestablish while cleanup of the Adj-RIB-In routes is in progress. If a peer comes back up before its route has been deleted from the routing table, that route is not deleted. The **DeletePending** flag in the **show route detail** and **show route extensive** command output indicates that a BGP route needs to be processed. **PurgePending**, **PurgeInProgress**, and **PurgeImpatient** flags in the **show bgp neighbor** command output show the status of the purge of routing table entries.

[See [Understanding External BGP Peering Sessions](#), [show bgp neighbor](#), [show route detail](#), and [show route extensive](#).]

Hardware

- **Support for 40-Gbps ports to operate at 10-Gbps speed (PTX10002-60C)**—You can use the Mellanox 10-Gbps pluggable adapter (model number: MAM1Q00A-QSA) to convert quad-lane-based ports to a single-lane-based SFP+ port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ module. Use the QSA adapter to convert a 40-Gbps port to a 10-Gbps port. You can plug a 10-Gbps SFP+ transceiver into the QSA adapter, which is then inserted into the QSFP or QSFP+ port of the PTX10002-60C router.

[See [Supported Transceivers on PTX10002-60C](#).]

High Availability (HA) and Resiliency

- **View ISSU status during an upgrade (MX240, MX480, MX960, MX2010, MX2020, PTX3000, and PTX5000)**—Starting in Junos OS Release 19.4R1, you can use the **request system software in-service-upgrade status** command to display the status of a unified ISSU. You will need to run this command on the Routing Engine where the unified ISSU was triggered to display the correct unified ISSU log file.

[See [request system software in-service-upgrade](#).]

Junos OS, XML, API, and Scripting

- **Python 3 support for commit, event, op, and SNMP scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, you can use Python 3 to execute commit, event, op, and SNMP scripts on devices running Junos OS. To use Python 3, configure

the **language python3** statement at the **[edit system scripts]** hierarchy level. When you configure the **language python3** statement, the device uses Python 3 to execute scripts that support this Python version and uses Python 2.7 to execute scripts that do not support Python 3 in the given release.

The Python 2.7 end-of-support date is January 1, 2020, and Python 2.7 will be EOL in 2020. The official upgrade path for Python 2.7 is to Python 3. As support for Python 3 is added to devices running Junos OS for the different types of onbox scripts, we recommend that you migrate supported script types from Python 2 to Python 3, because support for Python 2.7 might be removed from devices running Junos OS in the future.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

- **Automation script library upgrades (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, devices running Junos OS that support the Python extensions package include upgraded Python modules. Python scripts can leverage the upgraded versions of the following modules:

- **idna** (2.8)
- **jinja2** (2.10.1)
- **jnpr.junos** (Junos PyEZ) (2.2.0)
- **lxml** (4.3.3)
- **markupsafe** (1.1.1)
- **ncclient** (0.6.4)
- **packaging** (19.0)
- **paho.mqtt** (1.4.0)
- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

[See [Overview of Python Modules Available on Devices Running Junos OS.](#)]

Junos Telemetry Interface

- **Physical Ethernet interface sensor support on JTI (MX960, MX2020, PTX1000, and PTX5000)**—Starting in Junos OS Release 19.4R1, you can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) services or gRPC Network Management Interface (gNMI) services to export physical Ethernet interface statistics from MX960, MX2020, PTX1000, and PTX5000 routers to outside collectors. This feature supports OpenConfig model `openconfig-if-ethernet.yang` (physical interface level) version 2.6.2 (no configuration). Both streaming and ON-CHANGE statistics are supported using the following resource paths:
- `/interfaces/interface/ethernet/state/mac-address` (with ON_CHANGE support)
 - `/interfaces/interface/ethernet/state/auto-negotiate` (with ON_CHANGE support)
 - `/interfaces/interface/ethernet/state/duplex-mode` (with ON_CHANGE support)

- `/interfaces/interface/ethernet/state/port-speed` (with `ON_CHANGE` support)
- `/interfaces/interface/ethernet/state/enable-flow-control` (with `ON_CHANGE` support)
- `/interfaces/interface/ethernet/state/hw-mac-address` (with `ON_CHANGE` support)
- `/interfaces/interface/ethernet/state/negotiated-duplex-mode` (with `ON_CHANGE` support)
- `/interfaces/interface/ethernet/state/negotiated-port-speed` (with `ON_CHANGE` support)
- `/interfaces/interface/ethernet/state/counters/in-mac-control-frames`
- `/interfaces/interface/ethernet/state/counters/in-mac-pause-frames`
- `/interfaces/interface/ethernet/state/counters/in-oversize-frames`
- `/interfaces/interface/ethernet/state/counters/in-jabber-frames`
- `/interfaces/interface/ethernet/state/counters/in-fragment-frames`
- `/interfaces/interface/ethernet/state/counters/in-8021q-frames`
- `/interfaces/interface/ethernet/state/counters/in-crc-errors`
- `/interfaces/interface/ethernet/state/counters/in-block-errors`
- `/interfaces/interface/ethernet/state/counters/out-mac-control-frames`
- `/interfaces/interface/ethernet/state/counters/out-mac-pause-frames`

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

- **Transceiver sensor support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000)**—In Junos OS Release 19.4R1, you can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) or gRPC Network Management Interface (gNMI) services to export transceiver statistics from MX960, MX2010, MX2020, PTX1000 and PTX5000 routers to outside collectors. This feature supports OpenConfig transceiver model `openconfig-platform-transceiver.yang` 0.5.0.

Both streaming and ON-CHANGE statistics are supported using the following base path:

- `/components/components/transceiver/`

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#) and [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

- **Support for Segment Routing telemetry statistics and binding SIDs routes for uncolored Segment Routing Traffic Engineering policies (PTX1000, PTX3000, and PTX5000)**—Starting in Junos OS Release 19.4R1, Junos OS supports collection of traffic statistics for both ingress IP traffic and transit mpls traffic that take non-colored SR-TE paths on PTX Series routers. Binding SIDs for SRTE paths that have labels as first-hops in their segment lists are also now supported on PTX Series routers.

The `show spring-traffic-engineering lsp` command now has a `tunnel-source` filter, to display only the tunnels created from the specified sources by which the SRTE policy was provisioned. Also, the `show spring-traffic-engineering lsp detail` command now displays information on the source of the tunnel configuration and statistics. By default, traffic sensors and statistic collection are disabled for static SRTE

routes. To enable provisioning of JVISION traffic sensors in Junos OS data plane to stream out traffic statistics on SR policies and their Binding-SID routes, enable **statistics** under **telemetry** at the **[edit source-packet-routing telemetry]** hierarchy level, and sensors will be created for both the SRTE policy nexthop and Binding SID that are installed in the forwarding plane.

[See [source-packet-routing](#)]

MPLS

- **update-threshold statement modified to generate IGP update for lower bandwidth reservation (PTX Series)**—Starting in Junos OS Release 19.4R1, you can configure the threshold value of the **update-threshold** statement to accept:

- an integer or floating point values up to 3 significant digits after decimal point using the **threshold-percent** option
- an absolute value of bandwidth threshold which generates an IGP update using the **threshold-value** option

These options are mutually exclusive and can be used for generating an IGP update for lower bandwidth reservations.

[See [update-threshold](#).]

- **Distributed CSPF for segment routing LSPs (PTX Series)**—Starting in Junos OS Release 19.4R1, you can compute a segment routing LSP locally on the ingress device according to the constraints you have configured. With this feature, the LSPs are optimized based on the configured constraints and metric type. The LSPs are computed to utilize the available ECMP paths to the destination.

Prior to Junos OS Release 19.4R1, for traffic engineering of segment routing paths, you could either explicitly configure static paths, or use computed paths from an external controller.

[See [Enabling Distributed CSPF for Segment Routing LSPs](#).]

- **Color-based mapping of VPN services over SRTE (PTX Series)**—Starting in Junos OS Release 19.4R1, you can specify a color attribute along with an IP protocol next hop to resolve transport tunnels over static colored and BGP segment routing traffic-engineered (SRTE) label-switched paths (LSPs). This is called the color-IP protocol next hop resolution, where you are required to configure a resolution-map and apply it to the VPN services. Prior to this release, the VPN services were resolved over IP protocol next hops only.

With this feature, you can enable color-based traffic steering of Layer 2 and Layer 3 VPN services.

[See [Color-Based Mapping of VPN Services Overview](#).]

- **Support for segment routing features (PTX10002)**—Starting with Junos OS Release 19.4R1, PTX10002 router support the following segment routing features:
 - BGP link-state distribution with SPRING extensions
 - SRGB for SPRING in IS-IS domain
 - Anycast and prefix segments in SPRING for IS-IS protocols

- IS-IS SPRING and RSVP coexistence
- Segment routing policy for traffic engineering on BGP
- Static adjacency segment identifier for ISIS and OSPF
- Static adjacency segment identifier for aggregate Ethernet member links
- Interoperability of segment routing with LDP
- RSVP-TE pop-and-forward LSP tunnels
- BGP Labeled Unicast traffic statistics collection
- Static segment routing label switched path
- Interoperability of segment routing with LDP
- Topology Independent Loop-Free Alternate for IS-IS and OSPF
- MPLS ping and traceroute for segment routing
- Anycast and prefix segments in SPRING for OSPF protocols
- Configurable SRGBs used by SPRING in OSPF protocols

[See [Link-State Distribution Using BGP Overview](#), [Understanding Adjacency Segments](#), [Anycast Segments](#), [and Configurable SRGB in SPRING](#), [BGP Egress Traffic Engineering](#), [Static Adjacency Segment Identifier for ISIS](#), [Static Adjacency Segment Identifier for OSPF](#), [IS-IS User Guide](#), [OSPF User Guide](#).]

Routing Protocols

- **Support for BGP PIC Edge with BGP labeled unicast (MX Series and PTX Series)**—Starting with Junos OS Release 19.4R1, MX Series and PTX Series routers support BGP PIC Edge with BGP labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

[See [Load Balancing for a BGP Session](#).]

- **Unnumbered interface support for IS-IS and OSPFv2 with topology-independent loop-free alternate (ACX Series, MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, you can enable IPv4 processing on a point-to-point interface without assigning it an explicit IPv4 address. The router borrows the IPv4 address of another Ethernet or loopback interface already configured on the router and assigns it to the unnumbered interface to conserve IPv4 addresses.

To enable IPv4 processing for unnumbered interfaces include **unnumbered-address source** at the **[edit interfaces [name] unit [name] family inet]** hierarchy level.

[See [Configuring an Unnumbered Interface](#).]

- **Support for flexible algorithm in IS-IS for segment routing–traffic engineering (MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, you can thin slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For

example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path based on SPF calculation type to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

To define a flexible algorithm, include **flex-algorithm** statement at the **[edit routing-options]** hierarchy level.

To configure participation in a flexible algorithm include the **flex-algorithm** statement at the **[edit protocols isis segment routing]** hierarchy level.

[See [Understanding IS-IS Flexible Algorithm for Segment Routing](#).]

- **Decouple RSVP for IGP-TE (MX Series and PTX Series)**—Starting in Junos OS Release 19.4R1, a device can advertise selective **traffic-engineering** attributes such as **admin-color** and **maximum-bandwidth**, without enabling RSVP, for segment routing and interior gateway protocol (IGP) deployments.

Services Applications

- **Inline J-Flow scale enhancement (PTX10002)**—Starting in Junos OS Release 19.4R1, 100,000 flows per Packet Forwarding Engine are supported.

[See [Understanding Inline Active Flow Monitoring](#).]

- **Support for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP inline flow monitoring (PTX10002-60C)**—Starting in Junos OS Release 19.4R1, you can perform inline flow monitoring for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP traffic. Both IPFIX and version 9 templates are supported.

[See [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers](#).]

- **MPLS-over-UDP inner payload flow monitoring with IPFIX and version 9 formats (PTX10002-60C)**—Starting in Junos OS Release 19.4R1, on the PTX10002-60C router, you can perform flow monitoring for MPLS-over-UDP traffic to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. This feature supports MPLS IPv4 and IPv6 payloads and both IPFIX and version 9 templates. Only ingress sampling is supported.

[See [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers](#).]

Software-Defined Networking

- **Tunnel templates for PCE-initiated segment routing LSPs (PTX Series)**—Starting in Junos OS Release 19.4R1, you can configure a tunnel template for Path Computation Element (PCE)-initiated segment routing LSPs and apply it through policy configuration. These templates enable dynamic creation of segment routing tunnels with two additional parameters – Bidirectional forwarding detection (BFD) and LDP tunneling.

With the support for tunnel configuration, the LSPs that you would configure statically can now be automatically created from the PCE, thereby providing the benefit of reduced configuration on the device.

[See [Understanding Static Segment Routing LSP in MPLS Networks](#).]

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:

- Additional information is now logged for MGD-FFP intermodule communication.
- Commit errors that previously were only shown onscreen are now logged.

We provide a new operational command, **request debug information**, to speed up the initial information-gathering phase of debugging.

[See [request debug information](#).]

SEE ALSO

[What's Changed | 199](#)

[Known Limitations | 202](#)

[Open Issues | 203](#)

[Resolved Issues | 206](#)

[Documentation Updates | 208](#)

[Migration, Upgrade, and Downgrade Instructions | 209](#)

What's Changed

IN THIS SECTION

- [What's Changed in 19.4R2-S4 | 199](#)
- [What's Changed in 19.4R2 | 199](#)
- [What's Changed in 19.4R1 | 200](#)

Learn about what changed in Junos OS main and maintenance releases for PTX Series routers.

What's Changed in 19.4R2-S4

MPLS

- **Change in auto-bandwidth adjustment (PTX5000)**—If auto-bandwidth adjustment fails because of bandwidth unavailable error, the router tries to bring up the LSP with the same bandwidth during the subsequent reoptimization. In earlier releases, when the auto-bandwidth adjustment fails, the current bandwidth is reset to the bandwidth that was already active.

See [rsvp-error-hold-time](#).

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the `cspf-backoff-time` statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

See [cspf-backoff-time](#).

What's Changed in 19.4R2

General Routing

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:
 - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.
 - ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Routing Protocols

- **Advertising 32 secondary loopback addresses to Traffic Engineering Database (TED) as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—In Junos OS Release, multiple loopback addresses export into `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier Junos OS releases, multiple secondary loopback addresses in TED were added into `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router-id.

Software Installation and Upgrade

- **Install or activate the RIFT package to include the `request rift package activate-as-top-of-fabric` option (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Install or activate the RIFT package to include the `request rift package activate-as-top-of-fabric` option. This option is same as the `activate` option but it adds additional configuration to act as a **top-of-fabric** node.

What's Changed in 19.4R1

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, MX Series, PTX Series, and SRX Series)**—Starting with Junos OS Release 19.4R1, the `persist-group-inheritance` option at the `[edit system commit]` hierarchy is enabled by default. To disable this option, use `no-persist-groups-inheritance`.

[See [commit \(System\)](#).]

Interfaces and Chassis

- **Updates to the `show interfaces` and `show policer` commands (PTX Series)**—Starting in Junos OS release 19.4R1, on PTX Series routers, when you issue the `show interfaces` command or the `show policer` command, the output does not display the default arp policer (`_default_arp_policer_`). In earlier releases, when you issue the `show interfaces` command or the `show policer` command, the output displays the default arp policer (`_default_arp_policer_`) though the default arp policer is not supported on PTX series routers.
- **Change in Fabric Error Handling Behavior (PTX10008, PTX10016, PTX5000 routers (with FPC3-PTX-U2, FPC3-PTX-U3 FPCs), QFX10008, QFX10016, and QFX10002 switches)**—Starting in Junos OS release 19.4R1, when the PFE encounters ECC errors or parity errors related to fabric which are fatal, major, or correctable minor errors, the interfaces on the PFE are disabled. You must reboot the FPC manually to recover from the error. If you still face an issue after rebooting the FPC, contact our Customer Service. In earlier releases, when the PFE encounters any error (fatal, major, minor_correctable, minor_transient, and info), the errors were incorrectly classified as info and as a result, ignored.

Routing Protocols

- **XML RPC equivalent included for the show bgp output-scheduler | display xml rpc CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.4R1, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler](#).]

Software-Defined Networking

- **Increase in the maximum value of delegation-cleanup-timeout (PTX Series)**—Starting in Junos OS Release 19.4R1, you can configure a maximum of 2,147,483,647 seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

With the increase in maximum value of **delegation-cleanup-timeout** from 600 to 2,147,483,647 seconds, you can benefit during a Path Computation Element (PCE) failover or other network issues that might disrupt the PCEP session with the main active stateful PCE.

[See [delegation-cleanup-timeout](#).]

System Logging

- **Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

[What's New | 191](#)

[Known Limitations | 202](#)

[Open Issues | 203](#)

[Resolved Issues | 206](#)

[Documentation Updates | 208](#)

[Migration, Upgrade, and Downgrade Instructions | 209](#)

Known Limitations

IN THIS SECTION

- [General Routing](#) | 202

Learn about known limitations in Junos OS Release 19.4R2 for PTX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- On the PTX Platform with FPC Model FPC-PTX-P1-A or FPC2-PTX-P1A, you might encounter single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message gets reported: **Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002** The Junos OS Chassis Management Error handling detects such a condition, raises an Alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper support representative if the issue persists even after the FPC restarts. [PR1254415](#)
- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error.** [PR1268678](#)
- In the specific case of semigraceful RCB reboot initiated by the internal shell command **vhclient init 0**, GRES takes longer than 3 minutes to complete as opposed to 21 seconds. As a workaround, the CLI command **request vmhost reboot** (graceful) and a jack-out-jack-in of the Routing Engine (ungraceful) do not exhibit this delay. [PR1312065](#)

- Traffic loss for more than 15 seconds is seen when 50 percent of the aggregated Ethernet links are brought down by restarting multiple FPCs. [PR1412578](#)
- Because of an issue in the BIOS:QFXS_SFP_00.32_02.01 version, when the watchdog is disabled, the device does not reboot. [PR1441963](#)
- Call trace is observed during image upgrade from WRL6 to WRL9. [PR1442017](#)
- When a filter is attached in the outbound direction, GRE encapsulated headers are applied after the filter block in the egress direction. So in this case, it is possible that the filter is evaluated on an old header content (and not on the new GRE encapsulated header) and hence filter evaluation turns true and the new GRE encapsulated gets recirculated for another GRE encapsulation. This issue is difficult to fix as filter block evaluation happens before the new header is attached. [PR1465837](#)
- For scaled MACs as per the current design, the learn rate is expected. [PR1473334](#)

SEE ALSO

[What's New | 191](#)

[What's Changed | 199](#)

[Open Issues | 203](#)

[Resolved Issues | 206](#)

[Documentation Updates | 208](#)

[Migration, Upgrade, and Downgrade Instructions | 209](#)

Open Issues

IN THIS SECTION

- [General Routing | 204](#)
- [Multiprotocol Label Switching \(MPLS\) | 205](#)
- [Routing Protocols | 205](#)

Learn about open issues in Junos OS Release 19.4R2 for PTX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- Uneven load balancing of traffic might occur if the traffic stream changes only in 0-15 bits of the Layer 3 destination IPv6 address. This limitation might not be visible if the other parameters affecting the load balance change along with L3_DST, such as Layer 3 source IP address, Layer 4 source/destination ports, and so on. [PR1065515](#)
- Power budget values for a PTX5000 chassis, FPC, and PICs have been revised. For routers operating on limited power, this can change the point where alarms for power-over-budget or insufficient power are raised or cleared. [PR1216404](#)
- On a PTX Series router with a third-generation FPC, an error message is displayed when the FPC goes online or offline. [PR1322491](#)
- Alarm action does not work for minor errors after the threshold is changed to 1. [PR1345154](#)
- You might not be able to stop the ZTP bootstrap process when a PTX10016 or PTX10008 router with many line cards is powered on with the factory-default configuration. [PR1369959](#)
- VTY command **show filter index <number> counter** shows values as zero at 28-02-HOSTBOUND_NDP_DISCARD_TERM on PTX5000 platform. Basically, the counter does not increase for NDP packets. The issue is only with **show filter index**, which is a debug tool in VTY. This issue has no impact on NDP functionality for user traffic. [PR1420057](#)
- The firewall counter for lo0 interface might not increase. As a workaround, set the lo0 filter family inet and family inet6 counters instead of filter family any. [PR1420560](#)
- The em2 interface configuration causes FPC to crash during initialization and FPC does not come online. After deleting the em2 configuration and restarting the router, FPC comes online. [PR1429212](#)
- Interface statistics does not get updated with port-mirroring. [PR1431607](#)
- There is an increase in traffic loss after a unified ISSU with InterAS Layer 3 VPN OptionB configuration. [PR1435578](#)
- Memory leaks are observed while running the longevity check. [PR1438358](#)
- The full-resolve tunnel uses chain composite next-hop to program tunnel composite next-hop. Since the chain composite next-hop is created from the resolver, it has delay to delete logic to save the IPC call to the kernel or Packet Forwarding Engine. If the full-resolve dynamic tunnel (in this case, IPoIP tunnel) is deleted and created again within 10 seconds, it reclaims the old tunnel. This is expected behavior. Consequently, you can see the old statistics of the tunnel. [PR1444081](#)
- The XML output for colored routes displays **<c>** instead of **<c>** for colored routes. It does not have any impact on functionality. It is just a display issue. [PR1447156](#)
- The **show route** command does not give all the next-hop information in the case of multipath routes. It is just a display issue. [PR1458000](#)
- Statistic comparison between CLI and Junos telemetry interface for queue fails as the buffers show incorrect values. [PR1460246](#)

- Traffic fails with **gcm-aes-xpn-128** cipher when you perform an event. [PR1460254](#)
- Optics-options syslog and link-down do not work as expected on PTX5000 with FPC3. For low warning breach event, when the configured action is syslog only, link goes down. [PR1461404](#)
- When you configure the **inline-jflow** statement for sampling the MPLS traffic, multi-svsc process runs on the FPC might crash because of MPLS resend flow referencing to an unavailable memory location. The issue is related only to MPLS traffic profile and has no relation with any other J-Flow (that is, IPv4 and IPv6). [PR1477445](#)
- On PTX1000 routers, port mirroring is not working in Junos OS Release 19.4R2. [PR1491789](#)

Multiprotocol Label Switching (MPLS)

- In an RSVP LSP with loose or undefined path scenario, the LSP might stay in down state due to loop detection after the link in the path flaps. [PR1384929](#)
- On all Junos OS platforms with distributed CSPF under SR-TE scenario, if you execute some operations like deactivate or activate SR protocols, restart routing, and so on, the rpd crash might be observed. [PR1493721](#)

Routing Protocols

- With BFD configured on an aggregated Ethernet interface, if you disable or enable the aggregated Ethernet interface, then the aggregated Ethernet interface and the BFD session might not come up. [PR1354409](#)
- The **show dynamic-tunnels database** command does not show the current value of traffic statistics. It shows the cached value of traffic statistics which might not be equal to the current value. [PR1445705](#)
- Post IGP convergence backup IPoIP tunnel remains up. As a workaround, you must deactivate or activate dynamic tunnel. [PR1447153](#)
- NSR with larger scales of 4000 and more IPv6 peers is not qualified or supported. [PR1461436](#)

SEE ALSO

[What's New | 191](#)

[What's Changed | 199](#)

[Known Limitations | 202](#)

[Resolved Issues | 206](#)

[Documentation Updates | 208](#)

[Migration, Upgrade, and Downgrade Instructions | 209](#)

Resolved Issues

IN THIS SECTION

- General Routing | 206
- Infrastructure | 207
- Layer 2 Ethernet Services | 207
- Multiprotocol Label Switching (MPLS) | 207
- Routing Protocols | 207

This section lists the issues fixed in Junos OS Release 19.4R2 for the PTX 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

- On PTX3000 and PTX5000 platforms, interface might stay down after maintenance. [PR1412126](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- Upgrading fails due to communication failure between Junos VM and host OS. [PR1438219](#)
- FPC reboot might in the event of J-Lock hog for more than 5 seconds. [PR1439929](#)
- The 100-Gigabit Ethernet interface might not come up after flapping on PTX5000. [PR1453217](#)
- The local-loopback test fails with together-options. [PR1458814](#)
- On PTX1000 and PTX10002 devices traffic might get silently dropped and discarded after transient SIB and FPC voltage alarms. [PR1460406](#)
- The sample/syslog/log action in output firewall filter with packet of size less than 128 bytes might cause ASIC wedge (all packet loss) on PTX Series platforms. [PR1462634](#)
- PIC might restart if the temperature of QSFP optics is overheated on PTX3000 and PTX5000. [PR1462987](#)
- FPC might restart during runtime on PTX10000 and QFX10000 platforms. [PR1464119](#)
- EBUF parity interrupt is not seen on PTX Series platforms. [PR1466532](#)
- Packet Forwarding Engine error logs **prds_packet_classify_notification: Failed to find fwd nh for flabel 48** might be reported when IGMP packets got sampled on PTX5000 platform. [PR1466995](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)

- The `aftd-expr` core file generated at `JexprDfwTlvIffBindPoint`, `JexprHandleFilterBindPoint`, `EmExprFilterBindPoint`, `EmEngineCoreGroupManager`, `EmEngine::performAction`. [PR1468483](#)
- Incorrect counter value is seen for **arrival rate** and **peak rate** for DDoS command. [PR1470385](#)
- A PTX5000 SIB3 might fail to come up in slot 0 and/or slot 8 when RE1 is master. [PR1471178](#)
- The **input-vlan-map** or **output-vlan-map** might not work properly in Layer 2 circuit local-switching scenario. [PR1474876](#)
- `mib2d` generates core files while deleting channelized interfaces. [PR1479642](#)
- Adding or deleting multicast routes might cause adjacency and LSPs to go down. [PR1479789](#)
- FPC might crash when dealing with invalid next hops. [PR1484255](#)
- BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)

Infrastructure

- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later releases. [PR1462986](#)

Layer 2 Ethernet Services

- Member links state might be asynchronized on a connection between PE device and CE device in an EVPN A/A scenario. [PR1463791](#)

Multiprotocol Label Switching (MPLS)

- Kernel crashes and the device might restart. [PR1478806](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect use of the TCP-MSS. [PR1493431](#)

Routing Protocols

- SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. [PR1454177](#)
- On PTX3000 and PTX5000, `ppmd` crashes and a core file is generated after configuring `sbfd responder` on RE-DUO-2600. [PR1477525](#)
- The `rpd` process might crash with BGP multipath and routes withdraw occasionally. [PR1481589](#)

SEE ALSO

What's New	 191
What's Changed	 199
Known Limitations	 202
Open Issues	 203
Documentation Updates	 208
Migration, Upgrade, and Downgrade Instructions	 209

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed as User Guides](#) | 208

This section lists the errata and changes in Junos OS Release 19.4R2 documentation for the PTX Series.

Feature Guides Are Renamed as User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

What's New	 191
What's Changed	 199
Known Limitations	 202
Open Issues	 203
Resolved Issues	 206
Migration, Upgrade, and Downgrade Instructions	 209

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.4 | 209](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 212](#)
- [Upgrading a Router with Redundant Routing Engines | 212](#)

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

Basic Procedure for Upgrading to Release 19.4

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

```
user@host> request system snapshot
```

NOTE: The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the router, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and SSH files), might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [Installation and Upgrade Guide](#).

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.

To download and install Junos OS Release 19.4R2:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Click the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the router.

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.

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot  
source/junos-install-ptx-x86-64-19.4R2.9.tgz
```

Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-19.4R2.9-limited.tgz
```

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

- **/pathname**— For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname**

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

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

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 19.4R2, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following PTX Series routers:
 - PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

[See <https://kb.juniper.net/TSB17603>.]

NOTE: After you install a Junos OS Release 19.2j**install** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the [Installation and Upgrade Guide](#).

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 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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://support.juniper.net/support/eol/software/junos/>.

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation as follows:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine.
3. After making sure that the new software version is running correctly on the backup Routing Engine, switch over to the backup Routing Engine to activate the new software.
4. Install the new software on the original master Routing Engine that is now active as the backup Routing Engine.

For the detailed procedure, see the [Installation and Upgrade Guide](#).

SEE ALSO

What's New	191
What's Changed	199
Known Limitations	202
Open Issues	203
Resolved Issues	206
Migration, Upgrade, and Downgrade Instructions	209

Junos OS Release Notes for the QFX Series

IN THIS SECTION

●	What's New	214
●	What's Changed	226
●	Known Limitations	230
●	Open Issues	232
●	Resolved Issues	236
●	Documentation Updates	249
●	Migration, Upgrade, and Downgrade Instructions	250

These release notes accompany Junos OS Release 19.4R2 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- [What's New in Release 19.4R2 | 214](#)
- [What's New in Release 19.4R1 | 214](#)

Learn about new features introduced in the Junos OS main and maintenance releases for QFX Series switches.

NOTE: The following QFX Series platforms are supported in Release 19.4R2: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5200-32CD, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.

Junos on White Box runs on Accton Edgecore AS7816-64X switches in this release. The software is based on Junos OS running on QFX5210 switches, so release-notes items that apply to QFX5210 switches also apply to Junos on White Box.

What's New in Release 19.4R2

Software Defined Networking

- **VMware NSX Data Center for vSphere 6.4.5 and 6.4.6 certification (QFX5100 Virtual Chassis)**—Starting with Junos OS Release 19.4R2, Juniper Networks certifies QFX5100 Virtual Chassis as a hardware Virtual Extensible LAN (VXLAN) gateway in an Open vSwitch Database (OVSDb) and VXLAN network with a VMware NSX Data Center for vSphere 6.4.5 or 6.4.6 controller.

[See [OVSDb-VXLAN User Guide for QFX Series Switches \(VMware NSX\)](#).]

What's New in Release 19.4R1

EVPN

- **Layer 2 and Layer 3 VXLAN gateways (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, you can deploy QFX5120-32C switches as follows:
 - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
 - As a Layer 2 VXLAN gateway in an Open vSwitch Database (OVSDb) overlay network

VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See [Understanding VXLANs](#).]

- **EVPN pure type-5 route support (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, you can configure pure type-5 routing in an EVPN-VXLAN environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which carries the MAC address of the sending switch and provides next-hop reachability for the prefix. To configure pure type-5 routing, include the **ip-prefix-routes advertise direct-nexthop** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. To enable two-level equal-cost multipath (ECMP) next hops in an EVPN-VXLAN overlay network, you must also include the **overlay-ecmp** statement at the **[edit forwarding-options vxlan-routing]** hierarchy level.

[See [ip-prefix-routes](#).]

- **EVPN control plane and VXLAN data plane support (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, QFX5120-32C switches support EVPN-VXLAN. By using a Layer 3 IP-based underlay network coupled with an EVPN-VXLAN overlay network, you can place endpoints anywhere in the network and remain connected to the same logical Layer 2 network.

EVPN-VXLAN is commonly deployed over the following physical underlay architectures:

- A two-layer IP fabric that includes spine devices (Layer 3 VXLAN gateways) and leaf devices (Layer 2 VXLAN gateways). You can deploy EX4650 and QFX5120 switches as spine or leaf devices in this fabric.
- A one-layer IP fabric that includes leaf devices that function as both Layer 2 and Layer 3 VXLAN gateways. You can deploy EX4650 and QFX5120 switches as leaf nodes in this fabric.

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

- **Dynamic load balancing in an EVPN-VXLAN overlay network (QFX5200 and QFX5210)**—In Junos OS Releases before Release 19.4R1, QFX5200 and QFX5210 switches support a static load-balancing scheme based on destination MAC addresses. This scheme distributes traffic on a round-robin basis among virtual tunnel endpoints (VTEPs) in an EVPN-VXLAN overlay network.

Starting in Junos OS Release 19.4R1, QFX5200 and QFX5210 switches that function as leaf or spine devices in an EVPN-VXLAN overlay network (centrally-routed and edge-routed bridging overlays) support dynamic load balancing among different equal-cost VTEPs. When enabled, the dynamic load-balancing feature supersedes the static load-balancing feature. With the dynamic feature, traffic is hashed among equal-cost paths based on packet fields. We support this feature in the following use cases:

- A leaf device is multihomed to multiple spine devices.
- A host is multihomed to multiple leaf devices.

In both use cases, each multihomed physical, aggregated Ethernet, or logical interface is configured with an Ethernet segment identifier (ESI). Dynamic load balancing supports a maximum of 255 ESIs. If you exceed this maximum (for example, you configure 256 ESIs), traffic destined for the 256th ESI is flooded to the VLAN associated with the ESI.

The hashing takes place before a packet undergoes VXLAN encapsulation. We use these fields to load-balance traffic:

- Packets with an IP header:
 - IP header fields:
 - Source IP address
 - Destination IP address
 - Protocol
 - VLAN ID
 - Layer 4 (TCP and UDP) source and destination ports
- Packets with an MPLS/IP header:
 - Up to three top labels
 - IP header fields:
 - Source IP address
 - Destination IP address
 - Layer 4 (TCP and UDP) source and destination ports
- Packets with a Layer 2 header only:
 - Source MAC address
 - Destination MAC address
 - VLAN ID

To enable dynamic load balancing, include the **vxlان-overlay-load-balance** configuration statement at the **[edit forwarding-options]** hierarchy level and restart your switch.

To further control the hashing input used by this feature, include the **enhanced-hash-key** configuration statement at the **[edit forwarding-options]** hierarchy level.

- **Assisted replication in data centers with EVPN-VXLAN overlay networks (QFX Series switches)**—Starting in Junos OS Release 19.4R1, QFX Series switches support assisted replication (AR) in data centers with EVPN-VXLAN networks to optimize replication of BUM traffic forwarded into the EVPN core. Instead of flooding BUM traffic using ingress replication to multiple remote virtual tunnel endpoints (VTEPs) for a VLAN or virtual network identifier (VNI), devices configured as AR leaf devices (also called AR clients) forward the traffic to an AR replicator device that can better handle the replication load. The AR replicator

then replicates and forwards the traffic to the VXLAN overlay tunnels. For further optimization, you can configure AR with IGMP snooping.

Switches in the QFX10000 line can be AR replicators. Any QFX Series switches that support EVPN-VXLAN can be AR leaf devices.

[See [Assisted Replication Multicast Optimization in EVPN Networks](#).]

- **Support for EVPN routing policies (ACX5448, EX4600, EX4650, EX9200, MX Series, QFX Series, and vMX)**—Starting in Junos OS Release 19.4R1, Junos OS has expanded routing policy support to include the creation and application of policy filters specific to EVPN routes. You can create policies and apply policy filters to import and export EVPN routes at the routing-instance level or at the BGP level. Junos OS supports the following matching criteria for EVPN routes:
 - Route distinguisher ID
 - NLRI route type
 - EVPN Ethernet tag
 - BGP path attributes
 - Ethernet Segment Identifier
 - MAC Address on EVPN route type 2 routes
 - IP address on EVPN route type 2 and EVPN route type 5 routes
 - Extended community

[See [Routing policies for EVPN](#).]

- **Features supported on EX4650 and QFX5120 switches**—Starting with Junos OS Release 19.4R1, the following Junos OS features are supported on EX4650 and QFX5120 switches:
 - Automatically generated Ethernet segment identifiers (ESIs) in EVPN-VXLAN and EVPN-MPLS networks.
[See [Understanding Automatically Generated and Assigned ESIs in EVPN Networks](#).]
 - Firewall filtering and policing on EVPN-VXLAN traffic.
[See [Understanding VXLANs](#) and [Overview of Firewall Filters](#).]
 - Graceful restart on EVPN-VXLAN.
[See [Graceful Restart in EVPN](#).]
 - IGMPv2 snooping for EVPN-VXLAN in a multihomed environment.
[See [Overview of IGMP Snooping in an EVPN-VXLAN Environment](#).]
 - IPv6 data traffic support through an EVPN-VXLAN overlay network.
[See [Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay](#).]
 - Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface.
[See [Understanding Flexible Ethernet Services Support with EVPN-VXLAN](#).]

- MAC limiting, storm control, and port mirroring support in EVPN-VXLAN overlay networks.
[See [MAC Limiting, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.](#)]
- Multihomed proxy advertisement.
[See [EVPN Multihoming Implementation.](#)]
- Selective multicast forwarding and SMET route support in EVPN-VXLAN.
[See [Overview of Selective Multicast Forwarding.](#)]
- Standard class-of-service (CoS) features—classifiers, rewrite rules, and schedulers—are supported on VXLAN interfaces.
[See [Understanding CoS on OVSD-Managed VXLAN Interfaces.](#)]
- VMTO for ingress traffic.
[See [Ingress Virtual Machine Traffic Optimization.](#)]

General Routing

- **Optimized BGP peer reestablishment (MX Series, PTX Series, and QFX Series)**—Starting with Junos OS Release 19.4R1, BGP peers in different groups can close in parallel. The connect/retry algorithm makes more frequent attempts to reestablish BGP peers, which reduces downtime. The connect/retry algorithm makes 16 attempts instead of 5 to reestablish BGP peers in the first 256 seconds after they go down. Peers can reestablish while cleanup of the Adj-RIB-In routes is in progress. If a peer comes back up before its route has been deleted from the routing table, that route is not deleted. The **DeletePending** flag in the **show route detail** and **show route extensive** command output indicates that a BGP route needs to be processed. **PurgePending**, **PurgeInProgress**, and **PurgeImpatient** flags in the **show bgp neighbor** command output show the status of the purge of routing table entries.

[See [Understanding External BGP Peering Sessions](#), [show bgp neighbor](#), [show route detail](#), and [show route extensive](#).]

Interfaces and Chassis

- **QFX5110 supports JNP-SFPP-10GE-T**—Starting in Junos OS Release 19.4R1, QFX5110 switches support the new copper 10GBASE-T SFP+ transceiver (JNP-SFPP-10GE-T), which provides a speed of 10 Gbps. Use the existing **show** commands such as **show chassis pic** and **show chassis hardware** to view the details of the transceivers.

NOTE: In case a device with a different interface speed (that is, 1 Gbps or 100 Mbps) is connected on the other side of the wire, the interface on the Juniper device does not come up.

[See [show chassis pic](#) and [show chassis hardware](#).]

- **Support for dynamic load balancing (QFX5120-32C and QFX5120-48Y)**—Starting in Junos OS Release 19.4R1, QFX5120-32C and QFX5120-48Y switches support dynamic load balancing (DLB) for ECMP and LAG. DLB is an enhancement to static load balancing. DLB considers member bandwidth utilization along with packet content for member selection.

You can use the following DLB modes to load-balance traffic:

- Flowlet
- Assigned flow
- Per-packet

To configure DLB for ECMP, include the **ecmp-dlb** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy level.

To configure DLB for LAG, include the **dlb** statement at the **[edit interfaces aex aggregated-ether-options]** hierarchy level.

NOTE: You cannot configure both DLB and resilient hashing at the same time. Otherwise, commit error will be thrown.

[See [Understanding Dynamic Load Balancing](#) and [Configuring Dynamic Load Balancing](#).]

- **Support for 10-Gbps speed using JNP-SFP-25G-DAC (QFX5120-48Y)**—Starting in Junos OS Release 19.4R1, you can use any of the following JNP-SFP-25G-DAC cables to set 10-Gbps speed on the SFP28 ports of a QFX5120-48Y switch:

- JNP-SFP-25G-DAC-1M
- JNP-SFP-25G-DAC-3M
- JNP-SFP-25G-DAC-5M

If you've plugged a JNP-SFP-25G-DAC cable into a QFX5120-48Y switch, then the SFP28 ports come up with 10-Gbps speed by default. To configure the SFP28 ports to operate at 25-Gbps speed, you must explicitly configure the speed of the first port in the port group using the **set chassis fpc 0 pic 0 port *port-num* speed 25g** command.

[See [Channelizing Interfaces on QFX5120-48Y Switches](#).]

- **Support for 10-Gbps speed on JNP-SFPP-10GE-T transceiver (QFX5100-48S)**—Starting in Junos OS Release 19.4R1, QFX5100-48S switches support JNP-SFPP-10GE-T transceiver. This transceiver supports 10-Gbps speed by default.

Junos OS, XML, API, and Scripting

- **Automation script library upgrades (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, devices running Junos OS that support the Python extensions package include upgraded Python modules. Python scripts can leverage the upgraded versions of the following modules:

- **idna** (2.8)
- **jinja2** (2.10.1)
- **jnpr.junos** (Junos PyEZ) (2.2.0)
- **lxml** (4.3.3)
- **markupsafe** (1.1.1)
- **ncclient** (0.6.4)
- **packaging** (19.0)
- **paho.mqtt** (1.4.0)
- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

[See [Overview of Python Modules Available on Devices Running Junos OS.](#)]

- **Python 3 support for commit, event, op, and SNMP scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, you can use Python 3 to execute commit, event, op, and SNMP scripts on devices running Junos OS. To use Python 3, configure the **language python3** statement at the **[edit system scripts]** hierarchy level. When you configure the **language python3** statement, the device uses Python 3 to execute scripts that support this Python version and uses Python 2.7 to execute scripts that do not support Python 3 in the given release.

The Python 2.7 end-of-support date is January 1, 2020, and Python 2.7 will be EOL in 2020. The official upgrade path for Python 2.7 is to Python 3. As support for Python 3 is added to devices running Junos OS for the different types of onbox scripts, we recommend that you migrate supported script types from Python 2 to Python 3, because support for Python 2.7 might be removed from devices running Junos OS in the future.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

Junos Telemetry Interface

- **JTI and OpenConfig support for VLAN sensors (EX4650, QFX5120)**—Junos OS Release 19.4R1 supports the export of VLAN statistics using either Junos telemetry interface (JTI) services or remote procedure call (gRPC) services. You can export statistics at configurable intervals to an outside collector.

This feature includes OpenConfig support for the data model **openconfig-vlan.yang** for VLAN configuration version 1.0.2.

Use the following resource paths in a gRPC or gNMI subscription:

- **/vlans/**
- **/vlans/vlan/state/name**
- **/vlans/vlan/state/vlan-id**
- **/vlans/vlan/state/status**
- **/vlans/vlan/members/**
- **/vlans/vlan/members/member/interface-ref/state/interface/**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/interface-mode**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/native-vlan**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/access-vlan**
- **/vlans/vlan/members/member/interface-ref/state/interface/switched-vlan/state/trunk-vlan**
- **/vlans/vlan/members/member/interface-ref/state/interface/vlan/state/vlan-id**

Streaming telemetry data through gRPC or gNMI also requires the OpenConfig for Junos OS module.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Layer 2 Features

- **Ethernet ring protection switching (ERPS)(EX4650 and QFX5120)**—Starting in Junos OS Release 19.4R1, the EX4650 and QFX5120 support ERPS to reliably achieve carrier-class network requirements for Ethernet topologies forming a closed loop. The ITU-T Recommendation is G.8032 version 1.

ERPS version 1 comprises the following features:

- Revertive mode of operation of the Ethernet ring
- Multiple ring instances on the same interfaces
- Multiple ring instances on different interfaces
- Interworking with Spanning Tree Protocol, Multiple Spanning Tree Protocol, and redundant trunk groups

[See [Ethernet Ring Protection Switching Overview](#).]

- **Redundant Trunk Group support (EX4650 and QFX5120)**—Starting with Junos OS Release 19.4R1, EX4650 and QFX5120 switches support redundant trunk group (RTG) links.

[See [Redundant Trunk Groups](#).]

MPLS

- **MPLS scaling enhancements (EX4650 and QFX5120)**—Starting in Junos OS Release 19.4R1, MPLS scaling is enhanced on EX4650 and QFX5120 switches. For instance, you can increase the scale from its default 1024 to 8192 on QFX5120 switches. This enhancement optimizes and increases the ingress tunnel scale to address the current needs of data center networks either in IP-CLOS or IP over MPLS application spaces.

[See [Supported MPLS Scaling Values](#).]

Routing Protocols

- **Integrating RIFT protocol into Junos OS (MX240, MX480, MX960, QFX5100, QFX5110, QFX5120-32C, QFX5120-48Y, QFX5120-48YM, QFX5200, QFX5210, QFX10008, and VMX virtual routers)**—Starting in Junos OS Release 19.4R1, you can integrate a new IGP protocol, Routing in Fat Tree (RIFT), into Junos OS to route packets in variants of CLOS-based and fat tree network topologies (also called the spine and leaf model).

The RIFT protocol is capable of automatic construction of fat-tree topologies, providing you the benefit of having a close to zero necessary configuration. RIFT makes networks resilient, extensively traceable, and simpler to manage, thereby overcoming the deployment limitations of evolving IP fabrics.

[See [RIFT Overview and Set Up](#).]

Software Defined Networking (SDN)

- **OVSDB support with VMware NSX for vSphere (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, the Open vSwitch Database (OVSDB) management protocol provides a control plane through which an NSX controller can provision QFX5120-32C switches. In an environment in which NSX Release 6.4.5 or later is deployed, an NSX controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and the reverse.

The physical underlay network over which OVSDB-VXLAN is commonly deployed is a two-layer IP fabric that includes spine and leaf devices. The spine devices function as Layer 3 VXLAN gateways, and the leaf devices function as Layer 2 VXLAN gateways. You can deploy QFX5120 switches as leaf devices in this fabric.

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

- **Layer 2 and Layer 3 VXLAN gateways (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, you can deploy QFX5120-32C switches as follows:
 - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
 - As a Layer 2 VXLAN gateway in an OVSDB overlay network

VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not

need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See [Understanding VXLANs](#).]

- **Map PCE-initiated P2MP LSPs to MVPN (QFX Series)**—Starting in Junos OS Release 19.4R1, you can associate a single or range of MVPN multicast flows (S,G) to a dynamically created PCE-initiated point-to-multipoint label-switched path (LSP). You can specify only selective types of flows, which include a route distinguisher (RD), (S,G) address, and LSP name. When the incoming traffic matches the specified flows, it is mapped to the point-to-multipoint PCE-initiated LSP.

With this feature, you can benefit from reduced configuration as the PCE-initiated point-to-multipoint LSPs are dynamically mapped, thereby eliminating the need to statically enable MVPN and point-to-multipoint LSPs.

[See [Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs](#).]

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:

- Additional information is now logged for MGD-FFP intermodule communication.
- Commit errors that previously were only shown onscreen are now logged.

We provide a new operational command, **request debug information**, to speed up the initial information-gathering phase of debugging.

[See [request debug information](#).]

System Management

- **Precision Time Protocol (PTP) transparent clock (QFX5120 and QFX5210)**—Starting in Junos OS Release 19.4R1, you can use a transparent clock to update the PTP packets with the residence time as the packets pass through the switch. There is no master/slave designation. The switches support end-to-end transparent clocks, which include only the residence time. The transparent clock can update the residence time in a one-step process, which means it sends the timestamps in one packet.

To use a transparent clock, enable the **e2e-transparent** statement at the **[edit protocols ptp]**.

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

- **Additional support for Bidirectional Forwarding Detection (QFX5110, QFX5120, QFX5200, and QFX5210)**—Starting in Junos OS Release 19.4R2, Bidirectional Forwarding Detection (BFD) can support sessions of less than 1-second intervals. The inline BFD feature needs to be enabled with the **set routing-options ppm inline-processing-enable** command in order to support sub second BFD timers. Performance might vary depending on the configuration load within the system.

NOTE: IPv4 and standalone BFD sessions, as well as inline single-hop sessions are supported. Micro BFD implementation and logical router support are not supported.

NOTE: You must clear all BFD sessions after applying **set routing-options ppm inline-processing-enable**.

[See [Understanding Bidirectional Forwarding Detection \(BFD\)](#).]

VLAN Infrastructure

- **Support for multiple Q-in-Q tags (QFX10000 switches)**—Starting in Junos OS Release 19.4R1, the QFX10000 line of switches support the third and fourth Q-in-Q tags as payload (also known as pass-through tag) along with the existing two tags (for VLAN matching and operations). The QFX10000 switches support multiple Q-in-Q tags for both layer 2 bridging and EVPN-VXLAN cases. The Layer 2 access interfaces accept packets with three or four tags (all tags with the TPID value 0x8100). All the tags beyond the fourth tag (that is, from the fifth tag onward) are considered part of the Layer 3 payload and are forwarded transparently.

NOTE: In a one or two tagged packet, the tags (tag 1 and tag 2) can carry any TPID values (0x8100, 0x88a8, 0x9100, and 0x9200).

[See [Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation](#).]

SEE ALSO

[What's Changed | 226](#)

[Known Limitations | 230](#)

[Open Issues | 232](#)

[Resolved Issues | 236](#)

[Documentation Updates | 249](#)

[Migration, Upgrade, and Downgrade Instructions | 250](#)

What's Changed

IN THIS SECTION

- [What's Changed in 19.4R2-S4 | 226](#)
- [What's Changed in 19.4R2 | 226](#)
- [What's Changed in 19.4R1 | 228](#)

Learn about what changed in Junos OS main and maintenance releases for QFX Series.

What's Changed in 19.4R2-S4

MPLS

- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the `cspf-backoff-time` statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

See [cspf-backoff-time](#).

What's Changed in 19.4R2

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (EX Series and QFX Series)**—Starting with Junos OS Release 19.4R2, the `persist-groups-inheritance` option at the `[edit system commit]` hierarchy level is enabled by default. To disable this option, use `no-persist-groups-inheritance`.

[See [commit \(System\)](#).]

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:
 - When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.
 - ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered.

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#).]

Multicast

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 19.4R2, EX4600, EX4650, and the QFX5000 line of switches provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... n packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). In earlier releases, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** did not provide valid statistics for multicast traffic at Layer 2.

[See [show multicast snooping route](#).]

Software Installation and Upgrade

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.

What's Changed in 19.4R1

General Routing

- **Automatic installation of YANG-based CLI for RIFT protocol (MX Series, QFX Series, and vMX with 64-bit and x86-based servers)**—In RIFT 1.2 Release, installation of the CLI for RIFT protocol occurs automatically along with the installation of the junos-rift package. In the pre-1.0 releases of the junos-rift package, the RIFT CLI had to be installed separately using **request system yang** command after installation of the junos-rift package.

Interfaces and Chassis

- **Logical Interface is created along with physical interface by default (EX Series, MX Series, and QFX Series)**—Starting in Junos OS Release 19.4R1, logical interfaces are created on ge, et, and xe interfaces along with the physical interface, by default. In earlier Junos OS releases, by default, only physical interfaces are created.

For example, for ge interfaces, previously when you viewed the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

Management

- **entPhysicalTable fetched on QFX10002**—In Junos OS Release 19.4R1, the MIB data for entPhysicalTable will be fetched on a QFX10002-72Q or QFX10002-36Q switch.

[See [SNMP Explorer](#).]

Routing Protocols

- **XML RPC equivalent included for the `show bgp output-scheduler | display xml rpc` CLI command (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, we have included an XML RPC equivalent for the `show bgp output-scheduler | display xml rpc` CLI command. In Junos OS releases before Release 19.4R1, the `show bgp output-scheduler | display xml rpc` CLI command does not have an XML RPC equivalent.

[See [show bgp output-scheduler](#).]

Software Defined Networking (SDN)

- **Increase in the maximum value of `delegation-cleanup-timeout` (QFX Series)**—You can now configure a maximum of 2147483647 seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

With the increase in maximum value of **`delegation-cleanup-timeout`** from 600 to 2147483647 seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

[See [delegation-cleanup-timeout](#).]

System Logging

- **Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

SEE ALSO

[What's New | 214](#)

[Known Limitations | 230](#)

[Open Issues | 232](#)

[Resolved Issues | 236](#)

[Documentation Updates | 249](#)

[Migration, Upgrade, and Downgrade Instructions | 250](#)

Known Limitations

IN THIS SECTION

- [Layer 2 Ethernet Services | 230](#)
- [Layer 2 Features | 230](#)
- [Network Management and Monitoring | 230](#)
- [Platform and Infrastructure | 231](#)
- [Routing Protocols | 231](#)

Learn about known limitations in Junos OS Release 19.4R2 for QFX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Layer 2 Ethernet Services

- LACP force-up and EVPN core isolation features are not supported together. [PR1461581](#)

Layer 2 Features

- The **Targeted-broadcast forward-only** command does not broadcast the traffic. [PR1359031](#)
- On the QFX5000 line of switches, the following error message is reported in the log: **fpc0 Pools exhausted for Table:EGR_DVP_ATTRIBUTE_1**. [PR1479826](#)

Network Management and Monitoring

- The number of possible output interfaces in remote port mirroring varies among the various switches in the QFX5000 line of switches:
 - QFX5110, QFX5120, and QFX5210—Support a maximum of 4 output interfaces.

- QFX5100 and QFX5200—Support a maximum of 3 output interfaces.

Platform and Infrastructure

- The IPv6 transit statistics counter does not work. [PR1327811](#)
- VLAN does not get deleted in the hardware on IRB when disabled earlier, which causes ARP to refresh even though IRB is disabled. [PR1421382](#)
- On the QFX5110-32Q line of switches, when ports 0 to 23 are channelized as 4x10-Gigabit Ethernet in the default system-mode, all the ports do not work. [PR1438319](#)
- On the QFX5100 Virtual Chassis, the following error message is observed while issuing commit: **could not open configuration database (juniper.data+)**. [PR1452154](#)
- One of the ports of the VCP throughput test result for most of the frame sizes is not close to 100 percent. [PR1453709](#)
- Observed 100 percent Layer 2 MAC scaling traffic loss in the QFX10002-60C switch after loading the EVPN-VXLAN collapsed profile configurations. [PR1489753](#)

Routing Protocols

- On the QFX5000 line of switches, the targeted broadcast functionality with VXLAN is not supported. [PR1397086](#)

SEE ALSO

What's New 214
What's Changed 226
Open Issues 232
Resolved Issues 236
Documentation Updates 249
Migration, Upgrade, and Downgrade Instructions 250

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 232](#)
- [EVPN | 232](#)
- [Interfaces and Chassis | 232](#)
- [Junos Fusion for Provider Edge | 233](#)
- [Layer 2 Features | 233](#)
- [Platform and Infrastructure | 233](#)
- [Routing Protocols | 235](#)
- [User Interface and Configuration | 236](#)
- [Virtual Chassis | 236](#)

Learn about open issues in Junos OS Release 19.4R2 for QFX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- PFC feature is not supported with QFX5120 or EX4650 2-member Virtual Chassis due to BCM limitation. [PR1431895](#)

EVPN

- OVSDB-managed QFX5100 or QFX5110 line of switches encapsulates the VXLAN traffic and sends to the incorrect destination MAC address when multiple remote VTEPs are in the same subnet and reach through the IRB interface in a stretched VLAN. [PR1424698](#)
- The l2ald memory leak might be observed in an EVPN scenario. [PR1498023](#)
- EVPN-VXLAN core isolation does not work upon system reboot or routing restart. [PR1461795](#)

Interfaces and Chassis

- On the QFX5110 MC-LAG, flooding of the multicast packets for around 16 to 20 seconds is observed after disabling and enabling a member link of ICL after reboot. [PR1422473](#)

- ARP reply unicast packets might be flooded to all the interfaces in VLAN. [PR1454764](#)

Junos Fusion for Provider Edge

- IGMP membership does not get learned by the AD fully even when the IGMP queries are being sent out. [PR1419265](#)

Layer 2 Features

- On the QFX5000 Virtual Chassis, multicast traffic gets flooded even when the IGMP report times out. [PR1431893](#)
- The ARP packets are not counted against the storm control on the EVPN configured interface. [PR1469837](#)

Platform and Infrastructure

- On the QFX Series with a third-generation FPC, an error message is displayed when the FPC goes online or offline. [PR1322491](#)
- The PFE gets stuck in the **Init** state for one or two member with the TCP connection timeout error. [PR1503655](#)
- Vmcore is observed after the Virtual Chassis split event. [PR1498485](#)
- On the QFX5000 line of switches, misleading ISSU logs are printed during the NSSU process even when the box does not perform ISSU. [PR1451375](#)
- Inter and Intra VNI or VRF traffics are dropped between the CE devices when the interfaces connected between the TOR and multi-homes PE devices are disabled. [PR1498863](#)
- On the QFX5100 line of switches, the VTEP traffic are not seen as expected with the VXLAN filter policier. [PR1503122](#)
- The dcpfe process might generate core file with the non-oversubscribed mode. [PR1485854](#)
- Virtual Chassis is not stable with 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces. [PR1497563](#)
- On the QFX5120 line of switches, one of tport VCP throughput test result for most of the frame sizes is not close to 100% [PR1453709](#)
- On the QFX10000 line of switches, the Aruba wireless access point (AP) heartbeat packets get dropped. As a result, the Aruba wireless AP cannot work. [PR1352805](#)
- When a QFX10016 or QFX10008 switch with may line cards is powered on with the factory-default configuration, the switch is not able to stop the ZTP bootstrap. [PR1369959](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)

- With MLD-snooping enabled and when we have two receivers in the same VLAN interested in the same group address but from a different source, traffic will be received on only one receiver that sent the latest MLD report. This is because we do not install S, G routes in hardware when MLD snooping is enabled. [PR1386440](#)
- On Junos OS Release 18.4R1, intermittent traffic loss is observed with the RTG streams while flapping the RTG primary interface. [PR1388082](#)
- The subscriber bindings might not be successful. [PR1396470](#)
- Configurations with IRB interfaces might produce Packet Forwarding Engine error messages during the interface deletions. [PR1054798](#)
- Some interfaces of aggregated Ethernet bundle might go to the **Detached** state after the bulk configurations change. [PR1406691](#)
- On the QFX5110 and QFX5120 line of switches, the uRPF check in strict mode might not work properly. [PR1417546](#)
- On the QFX10000 line of switches, the analyzer does not mirror after adding the child member to an aggregated Ethernet interface. [PR1417694](#)
- On the QFX Series Virtual Chassis during shutdown, vmcore is observed if there is an interrupt. [PR1421250](#)
- LLDP frames received on a QFX5210 management em0 port might not be displayed in the **show LLDP operational** queries. Other non-em0 interfaces display statistics. [PR1426753](#)
- The l2ald memory leakage is observed when the VTEP or overlay BGP session flaps. [PR1435561](#)
- On the QFX5200 line of switches, the ISSU might fail. [PR1438690](#)
- The RIB installation or deletion time consumption is reduced. [PR1441737](#)
- On the QFX10000 line of switches, removal of the EVPN-VXLAN Layer 3 gateway on the IRB interface from the spine switches might cause traffic to be silently discarded. [PR1446291](#)
- Interface still sends mirrored traffic out even after it is removed from the output VLAN. [PR1452459](#)
- On the QFX5120 line of switches, one of the VCP ports of the throughput test result for most of the frame sizes is not close to 100 percent. [PR1453709](#)
- On the QFX5200-32C-32Q line of switches, vmcore occurs after upgrading from Junos OS Release 18.3 throttle image to Junos OS Release 19.3R1 at
...../amd/svl-engdata1vs1/occamdev/build/freebsd/stable_11/20190614.234225
__ci_fbsd_builder_stable_11.0.269d466/src/sys/kern/kern_shutdown.c:313. [PR1455851](#)
- On the QFX5000 and QFX10000 line of switches, 9.51 percent of degradation with commit time and 12 percent of degradation with VLAN commit convergence are observed while comparing 19.4DCB with 19.3DCB. [PR1457939](#)
- The **show dynamic-tunnels** database does not show IPv6 mapped next-hop flag for the 6 PE device routes that have labels. It is just a display issue. [PR1458634](#)

- On the QFX5110 line of switches, the VXLAN VNI (mcast) scaling causes traffic issue. [PR1462548](#)
- On the QFX5110 line of switches, the `libvirtMib_suba` might generate a core file after an image upgrades from Junos OS Release 17.3R3-S5.2 to Junos OS Release 17.3R3-S6.3. [PR1462725](#)
- The `show interfaces xe-a/b/c` statement on a disabled or enabled configuration change displays fiber intermittently. [PR1467509](#)
- On the QFX5100 line of switches, traffic loss might be seen with framing errors or runs if MACsec is configured. [PR1469663](#)
- If the system has 1000 BGP-V4 VRF (120,000 routes) + 700 OSPF-V2 VRF (70,000 routes) + 300 P2P IS-IS V4 VRF (30,000 routes) (nearly 220000 hardware routes), then the deletion or reading of the VRF configurations might lead to all BGP sessions going down. [PR1469881](#)
- ERP might not come up properly when the MSTP and ERP are enabled on the same interface. [PR1473610](#)
- Interfaces are not detected on some of the ports when the 25-Gigabit Ethernet SFP is swapped and 10-Gigabit Ethernet SFP is inserted. [PR1475574](#)
- The `dcpe` process might generate core file with the non-oversubscribed mode after SDK upgrade. [PR1485854](#)

Routing Protocols

- On the QFX-5100 VC/VCF line of switches, the following error is observed in the hardware with the mini-PDT base configurations: `BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:I3 nh 6594 unintsall failed`. [PR1407175](#)
- On the QFX5110-32Q switch, inter pod traffic loss is observed after loading the non collapsed VXLAN configurations. [PR1507471](#)
- On the QFX 5100-48T-6Q Virtual Chassis or Virtual Chassis fan, the following error message is observed while copying image to the Virtual Chassis fan member and trying to downgrade the image: `rcp for member 14, failed`. [PR1486632](#)
- On the QFX5100 Virtual Chassis, when you run the `show bgp neighbors` command, change in the x-path output for the value `input-updates` is observed. [PR1504399](#)
- Traffic drop might be observed after modifying the FBF firewall filter. [PR1499918](#)
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, traffic loss on multiple traffic streams is observed after reboot and the interfaces of the Virtual Chassis node flaps. [PR1500508](#)
- Traffic loss occurs in the BGP streams while triggering GRES and reboot with the base configuration. [PR1508133](#)
- The `rpdp` process generates core file at `rt_nh_resolve_add_gen` in `../../../../../../../../src/junos/usr/sbin/rpd/lib/rt/rt_resolve_ind.c`: with the evpn-dhcp configurations. [PR1494005](#)

- BGP session might go into the **Down** state once the traffic flow starts. [PR1431259](#)
- On the QFX5110 MC-LAG, the following error message is observed after the node reboot:
L2_L3_INTF_OPS_ERROR. [PR1435314](#)
- The VRF traffic loss on the mini-PDT based configurations occurs when you perform ISSU. [PR1449581](#)
- The IGMP snooping on the QFX5110 or QFX5129 leaf devices might cause multicast packets to loop in the multihomed scenario. [PR1457725](#)
- On the QFX10002-60C line of switches, the Packet Forwarding Engine installation or deletion, and link flap convergence time are reduced in Junos OS Release 19.4 compared to Junos OS Releases 19.3R1 and 19.2R1. [PR1464572](#)

User Interface and Configuration

- The version information under the configuration changes from Junos OS release 19.1. [PR1457602](#)

Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. [PR1490552](#)

SEE ALSO

What's New 214
What's Changed 226
Known Limitations 230
Resolved Issues 236
Documentation Updates 249
Migration, Upgrade, and Downgrade Instructions 250

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.4R2 | 237](#)
- [Resolved Issues: 19.4R1 | 242](#)

Learn which issues were resolved in Junos OS main and maintenance releases for QFX Series.

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

Resolved Issues: 19.4R2

Class of Service (CoS)

- Shaping does not work after the reboot if **shaping-rate** is configured. [PR1432078](#)
- The traffic is placed in the network-control queue on an extended port even if the traffic comes in with a different DSCP marking. [PR1433252](#)
- On the QFX5120-32C line of switches, when you move the unicast traffic to the multicast queue through the MF classifier, the **show interface queue <>** does not display any statistics. [PR1459281](#)

EVPN

- The ESI of the IRB interfaces does not get updated after the autonomous-system number change if the interface is down. [PR1482790](#)

Forwarding and Sampling

- The type 1 ESI/AD route might not be generated locally on the EVPN PE devices in the all-active mode. [PR1464778](#)

High Availability (HA) and Resiliency

- On the QFX5000 line of switches, ISSU is not supported for some versions. [PR1472183](#)

Interfaces and Chassis

- Executing commit might not work due to dcd process being nonresponsive. [PR1470622](#)
- Commit error is not thrown when the member link is added to the multiple aggregation group with different interface specific options. [PR1475634](#)
- Error message does not get generated while verifying the GRE limitation. [PR1495543](#)

Junos Fusion Enterprise

- Loop detection might not work on the extended ports. [PR1460209](#)

Layer 2 Features

- The LLDP function might fail when a device running Junos OS connects to a device that does not run Junos OS. [PR1462171](#)
- The FPC goes down when 100-Gigabits Ethernet link comes up on the QFX5110-48S switch. [PR1499422](#)
- On the QFX5000 line of switches, few MAC addresses might be missing from the MAC table in the software. [PR1467466](#)

- Ingress traffic might silently discard traffic if the underlying interfaces flap in the EVPN-VXLAN scenario. [PR1469596](#)
- Traffic might get affected if the composite next-hop is enabled. [PR1474142](#)

Layer 2 Ethernet Services

- The **relay-source** command is now applicable for the forward-only subscribers. [PR1455076](#)
- Member links state might be asynchronized on the connection between the PE devices and the CE devices in the EVPN A/A scenario. [PR1463791](#)
- When the wired DHCPv6 clients changes VLAN and an existing DHCPv6 relay binding exists on another VLAN, the DHCPv6 confirm packets from the client might not get processed correctly on the relay resulting in connectivity issues. [PR1496220](#)

MPLS

- On the QFX10002 line of switches, the **show mpls static-lsp | display xml** command produces invalid XML. [PR1469378](#)
- Traffic might be silently discarded on the PE devices when the CE devices send traffic to the PE devices and the destination is resolved with two LSPs through one upstream interface. [PR1475395](#)
- The traffic might get lost over the QFX5100 switch acting as a transit PHP node in the MPLS network. [PR1477301](#)
- The BGP session might keep flapping between two directly connected BGP peers because of the wrong TCP-MSS in use. [PR1493431](#)

Platform and Infrastructures

- On the QFX5100 Virtual Chassis, the following error message is observed: **MacDrainTimeOut and bcm_port_update failed: Internal error.** [PR1284590](#)
- The system might stop new MAC learning and impact the Layer 2 traffic forwarding. [PR1475005](#)
- On the QFX5120 switch under the MC-LAG scenario, traffic loss are observed. [PR1494507](#)
- FPC major error is observed after system boots up or FPC restarts. [PR1475851](#)
- The following error message is generated while booting: **CMQFX: Error requesting SET BOOLEAN, illegal setting 66.** [PR1385954](#)
- On the QFX5000 line of switches, the **show chassis errors active detail** command is not supported. [PR1386255](#)
- The 10-Gigabit Ethernet fiber interfaces might flap frequently when connected to a switch from another vendor. [PR1409448](#)
- The **show interface** command output indicates **Media type: Fiber** on the QFX5100-48T switch running the QFX 5E Series image. [PR1419732](#)

- The CoS rewrite rules applied under an aggregated Ethernet interface might not take effect after NSSU. [PR1430173](#)
- The l2cpd process might crash and generate a core file when the interfaces flap. [PR1431355](#)
- The default logical interface on the channelized IFD might not get created after ISSU or ISSR. [PR1439358](#)
- CRC error might be seen on the VCPs of the QFX5100 Virtual Chassis. [PR1449406](#)
- The em0 route might be rejected after the em0 interface is disabled and then enabled. [PR1449897](#)
- FPC does not restart immediately after the system reboot that causes packet loss. [PR1449977](#)
- On the QFX10000 line of switches, the CoS classification does not work. [PR1450265](#)
- On the QFX5000 line of switches, when a dual VLAN tag feature is configured on the physical interface that is not supported, no warning or error message is displayed. [PR1450455](#)
- Members might stay disconnected from the QFX5120-32C/QFX5120-48T Virtual Chassis after a full-stack reboot. [PR1453399](#)
- Changing the VLAN name associated with the access ports might prevent the MAC addresses from being learned under the EVPN-VXLAN scenario. [PR1454095](#)
- Master FPC might come up in the **Master** state again after the reboot instead of backup. [PR1454343](#)
- On the QFX5000 line of switches, the dcpfe process crashes due to the usage of data that is not null getting terminated. [PR1454527](#)
- In a 16+ member QFX5100 Virtual Chassis fan, the **FROM** column under the output of the **show system users** output reports feb0, feb1, feb2, or feb3 for FPC16, FPC17, FPC18, or FPC19, respectively. [PR1455201](#)
- On the QFX10000 line of switches, the PFC feature does not work. [PR1455309](#)
- The cosd process might crash if **forwarding-class-set** is directly applied on the child interface of an aggregated Ethernet interface. [PR1455357](#)
- The telemetry traffic might not be sent out when the telemetry server is reachable through a different routing-instance. [PR1456282](#)
- Link up delay and traffic drop might be observed on the mixed SP Layer 2 or Layer 3, and EP Layer 2 type configurations. [PR1456336](#)
- The QFX5110 QSFP-100GBASE-SR4 switch made by Avago cannot link up. [PR1457266](#)
- On the QFX10000 line of switches, the Packet Forwarding Engine might crash after the Routing Engine switchover. [PR1457414](#)
- In the QFX5110 line of switches, port 51 has one LED blinking amber. [PR1457516](#)
- On the QFX5000 line of switches, the **accept-source-mac** feature with VXLAN does not work. [PR1460885](#)
- On the QFX10000 line of switches, **show forwarding-options enhanced-hash-key** does not work. [PR1462519](#)

- On the QFX5000 line of switches, the fxpc process might generate a core file when you change MTU in a VXLAN scenario with firewall filters applied. [PR1462594](#)
- On the QFX10000 line of switches, the FPC process might restart during run-time. [PR1464119](#)
- A 10-Gigabit Ethernet interface might not come up on the QFX5100-48T switches or negotiate at speed 1-Gigabit Ethernet when connected with Broadcom 57800-T daughter card. [PR1465196](#)
- Starting from Junos OS Release 19.2R2, the PSU numbering on AS7816-64X platform has changed from 0->1 and 1->0 and the fan numbering has changed from 0,1,2,3->3,2,1,0. [PR1465327](#)
- On the QFX10000 line of switches, the EBUF parity interrupt is not observed. [PR1466532](#)
- On the QFX10000 line of switches, the EPR iCRC errors might cause protocols to go down. [PR1466810](#)
- The DHCPvX acknowledgment message does not receive a response to the Broadcast inform packets in Junos OS Release 19.1R2.3. [PR1467182](#)
- Ingress drops must be included at the CLI command from the interface statistics and added to the InDiscards. [PR1468033](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- MAC address might not be learned on a new extended port after VMotion in Junos fusion for data center environment. [PR1468732](#)
- If continuous interface flap occur at ingress or egress of the PE devices, the IP routed packets might get looped on the MPLS PHP node. [PR1469998](#)
- Incorrect counter value is observed for the arrival rate and the peak rate for the DDoS commands. [PR1470385](#)
- The speed 10m might not be configured on the ge- interface. [PR1471216](#)
- When the VTEP source interface is configured in the multiple routing instances, there might be traffic loss. [PR1471465](#)
- On the QFX5000 line of switches, the egress ACL filter entries is only 512 in Junos OS Release 19.4R1. [PR1472206](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- The detached interface in LAG might process the xSTP BPDUs. [PR1473313](#)
- The l2ald crash might be observed when around 16,000 VLAN-IDs share the same VXLAN tunnel and the Packet Forwarding Engine is rebooted. [PR1473521](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might get dropped. [PR1473685](#)
- On the QFX5000 line of switches in the EVPN-VXLAN scenario, continuous error log messages might be raised. [PR1474545](#)
- On the QFX5000 line of switches, the Layer 2 circuit might fail to communicate through VLAN 2. [PR1474935](#)

- On the QFX5200 line of switches, the DAC cables are not being properly detected in the Packet Forwarding Engine in Junos OS Release 18.4R2-S2.4. [PR1475249](#)
- There might be traffic drop on the QFX5110 or QFX5120 line of switches acting as a leaf switch in a multicast environment with VXLAN. [PR1475430](#)
- The QFX Series platforms exhibit invalid Packet Forwarding Engine PG counter pairs to copy src 0xfffff80, dst 0. [PR1476829](#)
- On the QFX10002-36Q/72Q switches, the following continuous error messages are logged on the device on getting adoption valid bit[8] asserted: **prds_ptc_wait_adoption_status: PECHIP[1] PTC[1]: timeout.** [PR1477192](#)
- The default Virtual Chassis MAC persistence timer is incorrectly set to 20 seconds instead of 20 minutes. [PR1478905](#)
- The remaining interface might be still in the **Down** state even though the number of channelized interfaces is no more than 5. [PR1480480](#)
- The ARP request packets for an unknown host might get dropped in the remote PE device in the EVPN-VXLAN scenario. [PR1480776](#)
- VLAN creation failure might be observed with the scaled VLAN and Layer 3 configuration. [PR1484964](#)
- On the QFX5120 Virtual Chassis, the output of the **show chassis alarm** command displays incorrect PEM status after multiple GRES events. [PR1486736](#)
- After ISSU or ISSR, a port using SR4 or LR4 optics might not come up. [PR1490799](#)
- The BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)
- The SLAX script might get lost after upgrading the software. [PR1479803](#)

Routing Protocols

- The OSPF VRF sessions take a long time to come up when the host table is full and the host routes are in the LPM table. [PR1358289](#)
- On the QFX5110 Virtual Chassis peers, invalid VRRP mastership election is observed. [PR1367439](#)
- On the QFX5100 line of switches, the BGP IPv4 or IPv6 convergence and the RIB installation or deletion time is reduced in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- PIM (S,G) joins might cause MSDP to incorrectly announce source active messages in some cases. [PR1443713](#)
- On the QFX5100 Virtual Chassis, the CRC errors might be observed. [PR1444845](#)
- Core files might get generated during the addition or removal of the EVPN type-5 routing instance. [PR1455547](#)
- The following error message is observed after the ISSU with Mini-PDT base configurations: **Error brcm_ipmc_route_counter_delete:3900Multicast stat destroy failed (-10:Operation still running).** [PR1460791](#)

- The mcsnood crash might be seen if one of the BD/VLANs is configured as a part of EVPN and has static or dynamic multicast router interfaces. [PR1468737](#)
- Flows do not fall back to a single link when the inactivity-interval is set higher than the IFG. [PR1471729](#)
- Traffic might not be forwarded over the ECMP link in the EVPN-VXLAN scenario. [PR1475819](#)
- ARP packets are always sent to CPU regardless of whether the storm-control is activated. [PR1476708](#)
- GRE transit traffic does not forward in the VRRP scenario. [PR1477073](#)
- The MUX state in the LACP interface does not go to the **Collecting and Distributing** state and remains in the **Attached** state after enabling the aggregated Ethernet interface. [PR1484523](#)
- The FPC process goes to the **NotPrsnt** state after upgrading the QFX5100 VC/VCF switch. [PR1485612](#)

Resolved Issues: 19.4R1

Class of Service (CoS)

- QFX10008: FPC0 generated core files after running the Packet Forwarding Engine command **show cos sched-usage**. [PR1449645](#)
- The **show cos scheds-per-pfe**, **show cos pfe-scheduler-ifds**, and **pfe** commands will restart forwarding planes on QFX10008 switches. [PR1452013](#)

EVPN

- Asynchronous result between ARP table and Ethernet switching table happens if EVPN ESI link flaps multiple times. [PR1435306](#)
- When using **no-arp-suppression**, an ARP request might not be sent out when an ARP entry aged out. [PR1441464](#)
- ARP and IPv6 neighbor entries cannot be cleared when they are learned from EVPN multihomed ESI. [PR1446957](#)
- EVPN-VXLAN NON-COLLAPSED: ARP will get resolved on QFX5100 for VXLAN having vlan-id of 2. [PR1453865](#)
- ARP request/NS might be sent back to the local segment by DF router. [PR1459830](#)

Forwarding and Sampling

- Commit error and dfwd core files might be observed when applying a firewall filter with action **then traffic-class** or **then dscp**. [PR1452435](#)

Interfaces and Chassis

- VRRPv6 state is flapping with init and idle states after configuring vlan-tagging. [PR1445370](#)
- On QFX10000 ARP entries might not be synchronized between MC-LAG devices. [PR1449806](#)

- The traffic might be forwarded to the incorrect interfaces in MC-LAG scenario. [PR1465077](#)
- Vrrpv3mibs are not working on QFX Series platform to poll VRRPv6 related objects. [PR1467649](#)

Layer 2 Features

- Storm control configuration might be disabled for the interface. [PR1354889](#)
- Packet loss might be seen when one of the spine switches fails or reboots. [PR1421672](#)
- Ethernet ring protection switching (ERPS) nodes might not converge to IDLE state after failure recovery or reboot. [PR1431262](#)
- EVPN-VXLAN NON-COLLAPSED: JTASK and multimove depth failed errors are seen after HALT. [PR1434687](#)
- The MAC/ARP learning might not work for copper base SFP-T on QFX5100/QFX5110/EX4600. [PR1437577](#)
- The traffic leaving QFX5000 and EX46000 switches might not be properly load-balanced over ae interfaces. [PR1448488](#)
- Unequal LAG hashing might happen on QFX devices. [PR1455161](#)
- The fxpc.core file might be seen when committing the configuration all together, for example, after the reboot. [PR1467763](#)

MPLS

- The l2circuit traffic might be silently dropped at EVPN SPINE/MPLS LSP TRANSIT device if VXLAN access interface flaps on remote PE node (QFX5110). [PR1435504](#)
- Packet loss might occur when ECMP resilient-hash is enabled on QFX5000 platforms. [PR1442033](#)

Platform and Infrastructure

- QFX5100-VC MacDrainTimeout and bcm_port_update failed: Internal error. [PR1284590](#)
- On QFX5100 platforms, LR4 QSFP can take up to 15 minutes to come up after Virtual Chassis reboot. [PR1337340](#)
- When powering off an individual FPC, the other FPC Packet Forwarding Engine might go offline too. [PR1344395](#)
- Mib2d core file in mib2d_write_snmpidx at snmpidx_sync.c on both ADs while bringing up base traffic profile. [PR1354452](#)
- Need new CLI command to enable copying of Open vSwitch Database (OVSDB) to RAM on Virtual Chassis backup Routing Engine instead of SSD. [PR1382522](#)
- FEC error counts are not updating for QFX5110. [PR1382803](#)
- QSFP-100GBASE-SR4/LR4 might take a long time to come up after disabling interface or reboot. [PR1402127](#)

- Ping over loopback might not work over type 5 tunnel on QFX10000 platforms. [PR1405786](#)
- QFX5200/5100 might not be able to send out control plane traffic to the peering device. [PR1406242](#)
- No inner VLAN tag is added even with **input-vlan-map push** configured on QFX10000 platforms. [PR1407347](#)
- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. [PR1411015](#)
- QFX5120 : Route table full for IPv6 routes in some scenarios. [PR1412873](#)
- Intermittently chassis alarms might not be raised after power-cycle of the device. [PR1413981](#)
- IPv6 multicast traffic received on one Virtual Chassis member might be dropped when egressing on other Virtual Chassis member if MLD snooping is enabled. [PR1423310](#)
- Ports might get incorrectly channelized if they are 10-Gigabit Ethernet already and they are channelized to 10-Gigabit Ethernet again. [PR1423496](#)
- On QFX5000 or QFX10000 switches, packet drops might be seen for the traffic that has to go over type-5 overlay tunnel. [PR1423928](#)
- The dcpfe/Packet Forwarding Engine might not start on AS7816-64X and QFX5000 TVP platform devices. [PR1426737](#)
- QFX5210: Received LLDP frames on em0 not displaying in LLDP neighbor output. [PR1426753](#)
- QFX5100-VCF - rollback for uncommitted configuration takes 1 hour. [PR1427632](#)
- Packet drops, replication failure, or ksynd crashes might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- The dcpfe process might crash and restart in MC-LAG scenario when the ARP/NDP next hop is changed. [PR1427994](#)
- The **global-mac-limit** and **global-mac-ip-limit** might allow more entries than the configured values. [PR1428572](#)
- [QFX10008] After Routing Engine switchover, LED status is not set for missing fan tray. [PR1429309](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- The dcpfe might crash on all line cards on QFX10000 in a scaled setup. [PR1431735](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- Outer VLAN tag might not be pushed in the egress VXLAN traffic toward the host for Q-in-Q scenario. [PR1432703](#)
- Line card might crash due to plug in unsupported SFP-T module. [PR1432809](#)
- Traffic loss might be seen on QFX10000/PTX10000 platforms using line card LC1105. [PR1433300](#)
- Layer 3 filters applied to PVLAN IRB interface might not work after unified ISSU. [PR1434941](#)

- QFX5100-Virtual Chassis : NSSU: there might be approximate 1 minute traffic loss during NSSU with LACP link protection configuration. [PR1435519](#)
- The mc-ae interface might get stuck in waiting state in dual mc-ae scenario. [PR1435874](#)
- QFX5200 NSSU: dcpfe core file is seen after NSSU upgrade of backup followed by reboot. [PR1435963](#)
- DHCP discover packets sent to IP addresses in the same subnet as IRB interface cause the QFX5110 to send bogus traffic out of DHCP-snooping enabled interfaces. [PR1436436](#)
- Unknown SNMP traps (1.3.6.1.4.1.2636.3.69.1.0.0.1) are sent on QFX5110 restart. [PR1436968](#)
- The FPC might crash if both the ae bundle flapping on the local device and the configuration change on peer device occur at the same time. [PR1437295](#)
- BGP neighborship might not come up if the MACsec feature is configured. [PR1438143](#)
- The DHCP snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it. [PR1438351](#)
- Port LED turns red when cable is connected on QFX5210. [PR1438359](#)
- Interfaces configured with **flexible-vlan-tagging** might loss connectivity. [PR1439073](#)
- The xSTP recognizes 1G SFP-T optic interface as LAN type resulting, in slow STP convergence. [PR1439095](#)
- LACP MUX state stuck in "Attached" after disabling peer active members when link protection is enabled on local along with force-up. [PR1439268](#)
- DHCPv6 relay binding is not up while verifying the DHCP Snooping along with DHCPv6 relay. [PR1439844](#)
- EX4600 Virtual Chassis does not comes up after replacing Virtual Chassis port from fiber connection to DAC cable. [PR1440062](#)
- MAC addresses learned on RTG might not be aged out after a Virtual Chassis member is rebooted. [PR1440574](#)
- QFX10002 MCLAG PDT: Layer 2, Layer 3 Traffic drop is seen on disabling/enabling MC-LAG. [PR1440732](#)
- The Layer 3 communication might break on an interface that is configured with **flexible-ethernet-services**. [PR1441690](#)
- The operational status of the interface in hardware and software might be out of synchronization in EVPN setup with **arp-proxy** feature enabled. [PR1442310](#)
- Flow control does not work as expected on 100-Gigabit Ethernet interface of QFX5110. [PR1442522](#)
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB. [PR1442587](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- When a line card is rebooted, the MC-LAG might not get programmed after the line card comes back online. [PR1444100](#)
- QFX5200: Observing **DCBCM[bcore_init]: ioctl call failed ret:0** failure message when changing UFT profile in FPC logs. [PR1445855](#)

- On QFX10008, traffic impact might be seen when the JSRV interface is used. [PR1445939](#)
- CoS classifier might not work as expected. [PR1445960](#)
- IPinIP: QFX - CoS rewrite happens to both inner and outer header. [PR1446128](#)
- IPinIP: ptx/qfx - Upon steering of underlay dynamic tunnel PNHs to a different set of ECMP next hops, unrelated IPv6 based tunnel traffic is tagged with the incorrect VLAN. [PR1446132](#)
- Traffic discarded for only specified VLAN in IPACL_VXLAN filters. [PR1446489](#)
- Long IPv6 address are not displayed fully on IPv6 neighbor table. [PR1447115](#)
- Unicast ARP requests are not replied with **no-arp-trap** option. [PR1448071](#)
- Rebooting QFX5120-48Y using **request system reboot** doesn't take physical links offline immediately. [PR1448102](#)
- QFX10000 -- QSFP28 100G AOC / 740-065632 & QSFP+ 40G / 740-043308 transceiver -- port LED remains lit green after disconnecting one end. [PR1448121](#)
- QFX5100-48t's in a mixed Virtual Chassis with QFX5110 switches are experiencing rx crc errors on vc-ports 53 and 52. [PR1449406](#)
- Except one AE member link, the other links do not send out sFlow sample packets for ingress traffic. [PR1449568](#)
- REST API process will get non-responsive when a number of request coming with a high rate. [PR1449987](#)
- RMPC core files are found after configuration changes are done on the network for PTP/Clock Synchronization. [PR1451950](#)
- Vgd core files might be generated when tunnel gets deleted twice. [PR1452149](#)
- DHCP offer packet with unicast flag set gets dropped by QFX10000 in a VXLAN multi-homed setup using anycast IP. [PR1452870](#)
- Configuration change in VLAN all option might affect the per-VLAN configuration. [PR1453505](#)
- The classifier configuration doesn't get applied to the interface in an EVPN/VXLAN environment. [PR1453512](#)
- The **show chassis led** shows incorrect status. [PR1453821](#)
- On QFX5100-VC VGD process hogs the CPU without **switch-options vtep-source-interface lo0.0** configuration. [PR1454014](#)
- Master FPC might come up in master state again after reboot instead of backup. [PR1454343](#)
- QFX10002-60c: EVPN-VXLAN: MAC+IP Count is shown as Zero. [PR1454603](#)
- QFX5120 : Untagged hosts ARP/NS connected on **encapsulation ethernet-bridge** interface are not being resolved. [PR1454804](#)
- The PFC feature doesn't work on QFX10000 platforms. [PR1455309](#)

- The laser from the 10G SFP+ interface is still on when the interface is disabled or the device is rebooted. [PR1456742](#)
- Over temperature SNMP trap messages are shown after update even though the temperatures are within the system thresholds. [PR1457456](#)
- Dual tag Q-in-Q is not working with EVPN-VXLAN. [PR1458206](#)
- QFX5210 : LED does not light on port 64 and 65 after upgraded to Junos OS Release 19.2R1. [PR1458514](#)
- The BPDU packet might be looped between leaf DF switch and non-DF switch and cause traffic blocking. [PR1458929](#)
- The dhcpv6 LDRA relay bounded count is not as expected after dhcp is configured. [PR1459499](#)
- The fxpc process might crash due to BGP IPV6 session flaps. [PR1459759](#)
- The **forwarding** option is missed in routing-instance type. [PR1460181](#)
- The 'entPhysicalTable' MIB is not fetching expected data on QFX10002-72Q / 36Q platforms. [PR1462582](#)
- The firewall filter does not get hit for traceroute packets when destination MAC address is VRRP virtual MAC. [PR1463425](#)
- On QFX5100 Virtual Chassis, the error **BRCM-VIRTUAL,brcm_vxlan_walk_svp(),6916:Failed to find L2-iff for ifl:** might appear during cleanup of EVPN-VXLAN configurations. These messages are harmless. [PR1463939](#)
- A few of the interfaces stay down and keep flapping for QFX ULC-3DWDM-MACsec line cards on reboot. [PR1464650](#)
- QFX5100-24Q: Not able to apply DSCP rewrite to firewall filter to a Layer 3 subinterface (for example, xe-0/0/0.100). [PR1464883](#)
- PEM is not present spontaneously on QFX5210. [PR1465183](#)
- The 10-Gigabit Ethernet port on QFX5100-48T negotiates with speed 1 GB with BRCM 10G/GbE 2+2P 57800-t rNDC. [PR1465196](#)
- The QSFP-100G-PSM4 could not be correctly identified on QFX5200 or QFX5110 platforms. [PR1465214](#)
- When BGP open messages with specific types of BGP optional capabilities are sent during BGP session establishment, incorrectly coded messages are later sent to the BMP Collector. [PR1466477](#)
- Slow packet drops might be seen on QFX5000 platforms. [PR1466770](#)
- Ingress drops to be included at CLI from interface statistics and added to InDiscards. [PR1468033](#)
- QFX5120 is looping the IP routed packet through IS-IS or MPLS. [PR1469998](#)
- l2ald core is seen (**l2ald_mem_free, l2ald_update_comp_vmenh**) after restarting dc-pfe in Virtual Chassis devices. [PR1473521](#)

Routing Protocols

- Host-directed packets with filter log action might not reach to Routing Engine if log/syslog is enabled. [PR1379718](#)
- The IRB transit traffic might not be counted for EVPN-VXLAN traffic. [PR1383680](#)
- QFX5100 : BGP IPv4 and IPv6 convergence and RIB installation and deletion time are degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The fxpc core file might be seen during the reboot of device on QFX5100/EX4600 switches. [PR1432023](#)
- The IPv4 fragmented packets might be broken if PTP transparent clock is configured. [PR1437943](#)
- Traffic might be dropped after the Q-in-Q enabled interface is flapped or a change is made to the vlan-id-list. [PR1441402](#)
- QFX5210: firewall Filter DSCP action modifier does not work when firewall filter is mapped to IRB. [PR1441444](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- PIM (S,G) joins can cause MSDP to incorrectly announce source active messages in some cases. [PR1443713](#)
- The QFX5120 might drop the tunnel encapsulated packets if it acts as a transit device. [PR1447128](#)
- Loopback address exported into other VRF instances might not work on ACX Series, EX Series, and QFX Series platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even the LDP neighbor's MAC changes. [PR1451217](#)
- A few seconds of traffic drop might be seen on the existing receivers when another receiver joins/leaves. [PR1457228](#)
- The egress interface in Packet Forwarding Engine for some end-hosts might not be correct on the Layer 3 gateway switch after it is rebooted. [PR1460688](#)
- The "other querier present interval" timer cannot be changed in IGMP/MLD snooping scenario. [PR1461590](#)
- When deleting IRB on the Layer 3 gateway, IRB does not get removed from Packet Forwarding Engine and will silently drop traffic to IRB MAC address. [PR1463092](#)

User Interface and Configuration

- EX4600 and QFX5100 were unable to commit baseline configuration after being returned to zero. [PR1426341](#)

SEE ALSO

| [What's New | 214](#)

What's Changed 226
Known Limitations 230
Open Issues 232
Documentation Updates 249
Migration, Upgrade, and Downgrade Instructions 250

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 249](#)

This section lists the errata and changes in Junos OS Release 19.4R2 for the QFX Series switches documentation.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [link](#).

SEE ALSO

What's New 214
What's Changed 226
Known Limitations 230
Open Issues 232
Resolved Issues 236
Migration, Upgrade, and Downgrade Instructions 250

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 250
- Installing the Software on QFX10002-60C Switches | 253
- Installing the Software on QFX10002 Switches | 253
- Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches | 254
- Installing the Software on QFX10008 and QFX10016 Switches | 256
- Performing a Unified ISSU | 260
- Preparing the Switch for Software Installation | 261
- Upgrading the Software Using Unified ISSU | 261
- Upgrade and Downgrade Support Policy for Junos OS Releases | 263

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

Upgrading Software on QFX Series 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 [Installation and Upgrade Guide](#) and [Junos OS Basics](#) in the QFX Series documentation.

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 **19.4** 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 19.4 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-host-qfx-5-x86-64-19.4-R1.n-secure-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 19.4 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.

NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

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

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-19.4R2.9.tgz
```

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

For example:

```
user@switch> request vmhost software add
ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-19.4R2.9.tgz
```

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

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

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R1.

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

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

For example:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-f-x86-64-19.4R2.n-secure-signed.tgz reboot
```

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

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-19.4R2.n-secure-signed.tgz reboot
```

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

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

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

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

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

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

To install the software on re0:

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

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

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

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

Installing the Software on QFX10008 and QFX10016 Switches

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

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



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

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

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

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

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

4. Disable nonstop-bridging:

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

5. Save the configuration change on both Routing Engines:

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

6. Exit the CLI configuration mode:

```
user@switch# exit
```

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

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

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

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

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

```
user@switch> request system software add validate  
/var/tmp/jinstall-host-qfx-10-f-x86-64-19.4R2.n-secure-signed.tgz
```

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

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

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

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

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

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

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

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

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

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

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

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

12. Transfer routing control to the backup Routing Engine:

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

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

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

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

```
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
Routing Engine status:
Slot 1:
  Current state           Master
  Election priority       Backup (default)
```

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

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-19.4R2.n-secure-signed.tgz
```

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

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

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

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

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

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

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

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

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

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

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

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

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

```
Routing Engine status:
  Slot 0:
    Current state           Master
    Election priority       Master (default)
Routing Engine status:
  Slot 1:
    Current state           Backup
    Election priority       Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

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

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 261](#)
- [Upgrading the Software Using Unified ISSU on page 261](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified 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 Unified ISSU

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

To upgrade the switch using unified 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 Series 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-host-qfx-10-f-x86-64-19.4R2.n-secure-signed.tgz*.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

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

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz 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: A unified 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 unified 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
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

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 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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>.

SEE ALSO

[What's New | 214](#)

[What's Changed | 226](#)

[Known Limitations | 230](#)

[Open Issues | 232](#)

[Resolved Issues | 236](#)

[Documentation Updates | 249](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 265](#)
- [What's Changed | 273](#)
- [Known Limitations | 279](#)
- [Open Issues | 282](#)
- [Resolved Issues | 284](#)
- [Documentation Updates | 295](#)
- [Migration, Upgrade, and Downgrade Instructions | 296](#)

These release notes accompany Junos OS Release 19.4R2 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

What's New

IN THIS SECTION

- [Release 19.4R2 New and Changed Features | 265](#)
- [Release 19.4R1 New and Changed Features | 265](#)

Learn about new features introduced in the Junos OS main and maintenance releases for SRX Series devices.

Release 19.4R2 New and Changed Features

There are no new features in Junos OS Release 19.4R2 for the SRX Series devices.

Release 19.4R1 New and Changed Features

Application Security

- **Selectively disable midstream APBR (SRX Series and vSRX)**—Starting in Junos OS Release 19.4R1, you can selectively turn-off midstream routing for a specific APBR rule while retaining the global APBR functionality for the remaining sessions.

When you disable the midstream routing for a specific APBR rule, the system does not apply midstream APBR for corresponding application traffic, and routes the traffic through a non-APBR route

[See [Advanced Policy-Based Routing](#).]

- **DSCP support for AppQoS (SRX Series and vSRX)**—Starting in Junos OS Release 19.4R1, AppQoS supports SLA-based path selection for an incoming traffic based on Differentiated Services Code Point (DSCP) value.

AppQoS depends on AppID and APBR to select the best possible link for the application traffic to meet the performance requirements specified in SLA. Junos OS Release 19.3R1 introduced APBR functionality for DSCP-tagged traffic. Using this enhancement, AppQoS selects the best possible link for the application

traffic based on the application signature, or DSCP value, or a combination of both application signature and DSCP value.

With this enhancement, now you can apply AppQoS for the encrypted traffic based on the DSCP value.

[See [Application Quality of Experience](#)].

- **Support for server certificates with key size 4096 bits (SRX300 and SRX320)**—Starting in Junos OS Release 19.4R1, SRX300 and SRX320 devices support RSA certificates with key size 4096 bits. You must explicitly configure the SSL proxy profile on these devices to use the server certificate with key size 4096 bits.

The RSA certificates with key size 4096 bits support is available only when the SRX300 and SRX320 devices are operating in standalone mode.

[See [Managing Certificates and Keys for SSL Proxy](#).]

Chassis Clustering

- **Increase in the maximum number of child links (SRX4600)**—Starting in Junos OS Release 19.4R1, you can configure up to eight child links in a redundant Ethernet bundle on each node of the chassis cluster. See [[Configuring Chassis Cluster Redundant Ethernet Interfaces on SRX4600](#)].

Flow-Based and Packet-Based Processing

- **Express Path for Flow Processing (SRX4600)**—Starting from Junos OS 19.4R1, Express Path is enabled by default on SRX4600 devices. You must configure Express Path only in policies. There is no need to configure Express Path on Flexible PIC Concentrator (FPC) or on Physical Interface Cards (PIC).

See [[Express Path](#)].

- **Support of IPFIX formatting for SRX J-Flow functionality (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vSRX, and vSRX3.0)** —Starting with Junos OS Release 19.4R1, you can use IPFIX flow templates to define a flow record for IPv4 traffic or IPv6 traffic. IPFIX is an enhanced version of J-flow version 9 template. Using IPFIX, you can collect a set of sampled flows and send the record to a specified host.

See [[Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches](#)].

- **Symmetric Fat Tunnel (SRX5400, SRX5600, and SRX5800 devices with SPC3 card, and vSRX)**—Starting from Junos OS 19.4R1, fat tunnel technology is introduced to improve the single IPsec tunnel throughput value to 10 times of current value.

To enable this feature, a new CLI command **fat-core** is introduced at the **set security distribution-profile** hierarchy level.

See [[Understanding Symmetric Fat IPsec Tunnel](#)].

General Packet Radio Switching (GPRS)

- **Increase in GTP scale for IoT and roaming firewall applications (SRX5400, SRX5600, SRX5800, and SRX4600)**—Starting in Junos OS Release 19.4R1, to enable Internet of Things (IoT) and roaming firewall use cases, the GTP tunnel scale per SPU is increased for the following SRX devices:
 - SRX5000 (SRX5400, SRX5600, SRX5800) SPC3: 1.2M to 12M
 - SRX5000 (SRX5400, SRX5600, SRX5800) SPC2: 600K to 3M
 - SRX4600: 400K to 4M

[See [Understanding Policy-Based GTP](#).]

Hardware

- **Wi-Fi mini-physical interface module (SRX320, SRX340, SRX345, and SRX550M)**—The Wi-Fi mini-physical interface module (mini-PIM) provides an integrated wireless LAN access point solution for branch SRX Series Services Gateways. The Mini-PIM supports the 802.11ac Wave 2 wireless standards and is backward-compatible with 802.11a, 802.11b, 802.11g, and 802.11n.

The Mini-PIM supports the following key features:

- 2x2 MU-MIMO
- Dual radios, which provide concurrent dual bands of 2.4 GHz and 5 GHz
- Eight virtual access points (VAPs) per radio
- Configurable transmit power
- 128 concurrent users

The Wi-Fi Mini-PIM is available in three models based on the regional wireless standards:

- SRX-MP-WLAN-US (United States)
- SRX-MP-WLAN-IL (Israel)
- SRX-MP-WLAN-WW (other countries)

[See [Wi-Fi Mini-Physical Interface Module](#).]

- **SRX5K-SPC3 LTC firmware version check and upgrade**—Starting in Junos OS Release 19.4R1, you can check the current LTC Firmware version on an SRX5K-SPC3 card and upgrade the firmware version manually.

The LEDs on the front panel of the services gateway chassis indicate a major alarm when the chassis detects that a newer version of LTC firmware is available and the firmware on the SRX5K-SPC3 card is outdated. The CLI commands:

- **show chassis alarm**—displays the alarm description
- **show system firmware**—displays the current version, available version, and the Status of the LTC firmware
- **request system firmware upgrade pic fpc-slot 0 pic-slot 0 tag 0**—updates the LTC firmware version.

[See [Chassis Component Alarm Conditions on SRX5400, SRX5600, and SRX5800 Services Gateways](#).]

Interfaces and Chassis

- **Wi-Fi Mini-Physical Interface Module (SRX320, SRX340, SRX345, and SRX550M)**—In Junos OS Release 19.4R1, we introduce the Wi-Fi Mini-Physical Interface Module (Mini-PIM). For retail and small offices, the Wi-Fi Mini-PIM provides secure wireless LAN connectivity to endpoint devices. The Wi-Fi Mini-PIM supports 802.11ac wave 2 wireless standards.

[See [Wi-Fi Mini-Physical Interface Module Overview](#).]

- **LTE Support in HA deployments (SRX300, SRX320, SRX340, SRX345, and SRX550HM)**—Starting in Junos OS Release 19.4R1, you can provide a backup WAN connection by configuring LTE modules on a pair of SRX devices operating in cluster mode.

[See [Configure LTE Mini-PIM in HA cluster mode](#).]

Intrusion Detection and Prevention (IDP)

- **IDP utility to read packet capture and generate protocol contexts (SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM)**—Starting from Junos OS Release 19.4R1, on SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM devices, to improve the IDP validation process, a CLI command is introduced to display and clear the contexts and the associated data only for the packet capture (PCAP) traffic. You can run the packet capture utility in either inet mode or transparent mode to generate attack contexts.

See [\[IDP Utility for PCAP\]](#).

- **Signature Language Constructs (SRX Series)**—Starting from Junos OS 19.4R1, signature language constructs are supported in the IDP engine code to write more efficient signatures that helps in reducing false positives.

The following constructs are supported:

- Depth
- Offset
- Within
- Distance
- Ipopts

See [\[IDP Signature Language Enhancements\]](#).

Junos OS XML API and Scripting

- **Python 3 support for commit, event, op, and SNMP scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, you can use Python 3 to execute commit, event, op, and SNMP scripts on devices running Junos OS. To use Python 3, configure the `language python3` statement at the `[edit system scripts]` hierarchy level. When you configure the

language python3 statement, the device uses Python 3 to execute scripts that support this Python version and uses Python 2.7 to execute scripts that do not support Python 3 in the given release.

The Python 2.7 end-of-support date is January 1, 2020, and Python 2.7 will be EOL in 2020. The official upgrade path for Python 2.7 is to Python 3. As support for Python 3 is added to devices running Junos OS for the different types of onbox scripts, we recommend that you migrate supported script types from Python 2 to Python 3, because support for Python 2.7 might be removed from devices running Junos OS in the future.

[See [Understanding Python Automation Scripts for Devices Running Junos OS.](#)]

- **Automation script library upgrades (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, devices running Junos OS that support the Python extensions package include upgraded Python modules. Python scripts can leverage the upgraded versions of the following modules:

- **idna** (2.8)
- **jinja2** (2.10.1)
- **jnpr.junos** (Junos PyEZ) (2.2.0)
- **lxml** (4.3.3)
- **markupsafe** (1.1.1)
- **ncclient** (0.6.4)
- **packaging** (19.0)
- **paho.mqtt** (1.4.0)
- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

[See [Overview of Python Modules Available on Devices Running Junos OS.](#)]

J-Web

- **Threat Assessment report enhancement (SRX Series)**—Starting in Junos OS Release 19.4R1, the Threat Assessment report displays a new *Filename* column in the *Malware downloaded by User* table. This column helps you to identify the malware filename.

[See [About Reports Page.](#)]

- **UTM enhancement (SRX Series)**—Starting in Junos OS Release 19.4R1, the following UTM pages (Configure > Security Services > UTM) are refreshed for a seamless experience:
 - Antivirus
 - Content Filtering
 - Policy

[See [About the Antivirus Page](#), [About the Content Filtering Page](#), and [About the Policy Page](#).]

- **Support for Wi-Fi Mini-PIM (SRX320, SRX340, SRX345, and SRX550M devices)**—Starting in Junos OS Release 19.4R1, J-Web supports the Wi-Fi Mini-Physical Interface Module (Mini-PIM). The physical interface for the Wi-Fi Mini-PIM uses the name `wl-x/0/0`, where x identifies the slot on the services gateway where the Mini-PIM is installed.

You can monitor and configure the wireless LAN settings using the J-Web interface.

[See [Dashboard Overview](#), [Monitor Ports](#), [About the Ports Page](#), [Monitor Wireless LAN](#), and [About the Settings Page](#).]

Logical Systems and Tenant Systems

- **Flow trace support at logical system and tenant system level (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.4R1, you can trace the packet flow at the logical system level and tenant system level. Traceoptions enables you to monitor traffic flow into and out of an SRX Series device.

When you trace traffic flow, you can generate and save the trace logs to the respective logical system and tenant system log files.

Flow trace at the level of logical system and tenant system helps you avoid generating large log files from the root level.

[See [Flow Trace Support for Logical Systems](#) and [Flow Trace Support for Tenant Systems](#).]

- **AppID statistics at tenant system level (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.4R1, you can view or clear the application identification (AppID) statistics, counters, and application system cache at the tenant system level.

[See [Application Security for Tenant Systems](#).]

Network Management and Monitoring

- **SNMP support for Wi-Fi Mini-Physical Interface Module (Mini-PIM) monitoring (SRX320, SRX340, SRX345, and SRX550M)**—Starting in Junos OS Release 19.4R1, you can monitor the Wi-Fi Mini-PIM status from remote network using SNMP. Use the `show snmp mib walk ascii jnxWlanWAPStatusTable` and `show snmp mib walk jnxWlanWAPClientTable` commands to monitor the Wi-Fi Mini-PIM status and client information.

[See [SNMP MIB Explorer](#) and `show snmp mib`.]

- **SNMP support for IPsec VPN flow monitoring (SRX5000 Series devices with SRX5K-SPC3 card)**—Starting in Junos OS Release 19.4R1, we have enhanced the existing IPsec VPN flow monitor MIB `jnxIpSecFlowMonMIB` to support the global IKE statistics for tunnels using IKEv2. Use the `show security ike stats` command to display the global statistics of tunnels such as in-progress, established, and expired negotiations using IKEv2.

[See [Enterprise-Specific SNMP MIBs Supported by Junos OS](#) and `show security ike stats`.]

- **Improved query performance in on-box reporting (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 19.4R1, we've upgraded the on-box logging database to improve query performance. For example, if you expect fewer traffic logs, you can use the default configuration with a start time and a stop time. If you expect a large number of traffic logs and greater time intervals for which the logs will be generated, we recommend you enable table dense mode.

[See [Understanding On-Box Logging and Reporting](#).]

- **Enhanced support for the non-default management instance (SRX Series)**—Starting in Junos OS 19.4R1, you can access information related to all routing instances and logical system networks and not specific to ingress routing instance by configuring the SNMPv3 management interface in a required management instance. Configuring the SNMPv3 management interface in a required management instance enables all the SNMPv3 requests coming from non-default routing instance is treated as if the requests are coming from default routing instance. You can configure the management instance configuration statement at the `[edit SNMP v3]` hierarchy level.

[See [SNMPv3 Management Routing Instance](#).]

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:
 - Additional information is now logged for MGD-FFP intermodule communication.
 - Commit errors that previously were only shown onscreen are now logged.

We provide a new operational command, **request debug information**, to speed up the initial information-gathering phase of debugging.

[See [request debug information](#).]

Unified Threat Management (UTM)

- **UTM support for active/active chassis cluster (SRX Series devices)**—Starting in Junos OS Release 19.4R1, you can configure all Unified Threat Management (UTM) features when the device is in active/active chassis cluster mode. The UTM features supported on an active/active chassis cluster include, Antispam Filtering, Content Filtering, Sophos Antivirus Scanning, Enhanced Web Filtering, Local Web Filtering, and Websense Redirect Web Filtering, and On-box/AviraAV. Enhanced Web Filtering and Sophos Antivirus Scanning remain active on both the primary node and the secondary node.

[See [Understanding UTM Support for Active/Active Chassis Cluster](#).]

- **UTM support for SMTPS, IMAPS, POP3S, and FTPS (SRX Series devices)** —Starting in Junos OS Release 19.4R1, UTM supports implicit and explicit SMTPS, IMAPS, and POP3S protocol and explicit passive-mode FTPS. SMTPS, IMAPS, POP3S, and FTPS are methods for securing SMTP, IMAP, POP3, FTP protocols

using Transport Layer Security (TLS). Antivirus and content filtering feature supports SMTPS, IMAPS, POP3S and FTPS protocol. Antispam feature only supports SMTPS protocol.

[See [Antispam Filtering Overview](#) and [Understanding Content Filtering Protocol Support](#).]

VPNs

- **Extended Sequence Number (SRX5400, SRX5600, and SRX5800 devices using SPC3)**—Starting from Junos OS Release 19.4R1, Extended Sequence Number (ESN) is introduced in IPsec VPN using IKE version 2 (IKEv2).
IPSec uses a 32-bit sequence number by default for the sequence number. When all sequence numbers are consumed, a rekey must be issued. By enabling ESN this 32-bit sequence numbering is increased to 64-bit.

You can enable ESN using the **set extended-sequence-number** command at the **edit security ipsec proposal proposal-name** level.

[See [Understanding Extended Sequence Number \(ESN\)](#).]

- **VPN support for inserting Services Processing Cards in Chassis Cluster (SRX5400, SRX5600, SRX5800)**—Starting in Junos OS Release 19.4R1, on all SRX5000 Series devices chassis cluster, you can insert a new SRX5K-SPC3 (SPC3) on the device without affecting or disrupting the traffic on the existing IKE or IPsec VPN tunnels. When you insert the new SPC3 in each chassis of the cluster, the existing tunnels are not affected and traffic continues to flow without disruption. You must reboot the node after you insert the SPC3 to activate the card. After the node reboot is complete, IPsec tunnels are distributed to the cards. After you reboot the secondary node where the new spc3 card is inserted, the IPsec sessions remain active on the other active node, without disruption to those sessions except during the failover time.

[See [Understanding VPN Support for Inserting Services Processing Cards](#).]

- **IPsec Encapsulating Security Payload authentication-only mode in PowerMode IPsec (SRX5000 Series devices with SRX5K-SPC3 card, and vSRX)**—Starting in Junos OS Release 19.4R1, you can enable the IPsec Encapsulating Security Payload (ESP) authentication-only mode in the PowerMode IPsec (PMI). The ESP authentication-only mode provides authentication, integrity checking, and replay protection in the PMI.

[See [Improving IPsec Performance with PowerMode IPsec](#).]

SEE ALSO

<i>What's Changed</i>
Known Limitations 279
Open Issues 282
Resolved Issues 284

[Documentation Updates | 295](#)

[Migration, Upgrade, and Downgrade Instructions | 296](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.4R2-S4 | 274](#)
- [What's Changed in Release 19.4R2 | 274](#)
- [What's Changed in Release 19.4R1 | 275](#)

Learn about what changed in Junos OS main and maintenance releases for SRX Series.

What's Changed in Release 19.4R2-S4

Flow-based and Packet-based Processing

- On SRX Series devices in earlier releases, when the session table was full there was no alarm set to indicate this. Starting from this release, when the percent of flow session table utilization is 95% on FPC and PIC, an alarm message ? Flow session table is almost full on FPC <number> PIC <number>? is set. Similarly, when the percent of DCP session table utilization is 95% on FPC and PIC, an alarm message ? DCP session table is almost full on FPC <number> PIC <number>? is set.

General Routing

- **Support for fully qualified domain name (FQDN) for log server (SRX Series)**—Starting in Junos OS Release, you can configure TTL value for a DNS server cache with hostname or IP address.

See [Configuring the TTL Value for DNS Server Caching](#).

What's Changed in Release 19.4R2

Authentication and Access Control

- **Enhanced user firewall support**—In Junos OS Release 19.4R2, for SRX300 devices with eUSB (SRX300, SRX320, SRX340, and SRX345), the SRX Series user firewall (UserFW) module tries to synchronize user entries from the domain controller or Juniper Identity Management Service (JIMS) after booting up. If the historical login events expired on the domain controller, then the SRX Series UserFW module is unable to retrieve those user entries after the UserFW module boots up.

[See [User Authentication Entries in the ClearPass Authentication Table](#).]

VPNs

- **Modified output of the show security pki local-certificate command (SRX Series)**—In Junos OS Release 19.4R2, the output of the **show security pki local-certificate** command is modified to include the logical systems name along with the existing output.

[See [show security pki local-certificate \(View\)](#).]

- **Change in encryption algorithm display (SRX Series)**—Starting in Junos OS Release 19.4R2, when you configure **aes-128-gcm** or **aes-256-gcm** as an encryption algorithm at the **[edit security ipsec proposal proposal-name]** hierarchy level, the authentication algorithm field of the **show security ike security-associations detail** and **show security ipsec security-associations detail** commands displays the same configured encryption algorithm.

Before Junos OS Release 19.4R2, the commands were displaying the authentication algorithm field as **UNKNOWN AUTH ALG**.

[See [show security ike security-associations](#) and [show security ipsec security-associations](#).]

What's Changed in Release 19.4R1

Application Security

- Starting in Junos OS Release 19.4R1, you have the flexibility to limit the application identification inspection as follows:

- Inspection Limit for TCP and UDP Sessions**

You can set the byte limit and the packet limit for application identification (AppID) in a UDP or in a TCP session. AppID concludes the classification based on the configured inspection limit. On exceeding the limit, AppID terminates the application classification.

If AppID does not conclude the final classification within the configured limits, and a pre-matched application is available, AppID concludes the application as the pre-matched application. Otherwise, the application is concluded as `junos:UNKNOWN` provided the global AppID cache is enabled. The global AppID cache is enabled by default.

To configure the byte limit and the packet limit, use the following configuration statements from the **[edit]** hierarchy:

- ```
user@host# set services application-identification inspection-limit tcp byte-limit byte-limit-number
packet-limit packet-limit-number
```
- ```
user@host# set services application-identification inspection-limit udp byte-limit byte-limit-number
packet-limit packet-limit-number
```

[Table 3 on page 275](#) provides the range and default value for configuring the byte limit and the packet limit for TCP and UDP sessions.

Table 3: Maximum Byte Limit and Packet Byte Limit for TCP and UDP Sessions

Session	Limit	Range	Default Value
TCP	Byte limit	0 through 4294967295	<ul style="list-style-type: none"> 6000 For Junos OS Release 15.1X49-D200, the default value is 10000.
	Packet limit	0 through 4294967295	Zero
UDP	Byte limit	0 through 4294967295	Zero
	Packet limit	0 through 4294967295	<ul style="list-style-type: none"> 10 For Junos OS Release 15.1X49-D200, the default value is 20.

The byte limit excludes the IP header and the TCP/UDP header lengths.

If you set both the **byte-limit** and the **packet-limit** options, AppID inspects the session until both the limits are reached.

You can disable the TCP or UDP inspection limit by configuring the corresponding **byte-limit** and the **packet-limit** values to zero.

- **Global Offload Byte Limit (Other Sessions)**

You can set the byte limit for the AppID to conclude the classification and identify the application in a session. On exceeding the limit, AppID terminates the application classification.

If AppID does not conclude the final classification within the configured limits, or the session is not offloaded due to tunneling behavior of some applications, and a pre-matched application is available, AppID concludes the application as the pre-matched application. Otherwise, the application is concluded as `junos:UNKNOWN` provided the global AppID cache is enabled (the global AppID cache is enabled by default).

To configure the byte limit, use the following configuration statement from the **[edit]** hierarchy:

```
set services application-identification global-offload-byte-limit byte-limit-number
```

The default value for the **global-offload-byte-limit** option is 10000 and the range is 0 through 4294967295.

You can disable the global offload byte limit by configuring the **global-offload-byte-limit** value to zero.

The byte limit excludes the IP header and the TCP/UDP header lengths.

- Starting in Junos OS Release 19.4R1, the maximum packet threshold for DPI performance mode option **set services application-identification enable-performance-mode max-packet-threshold *value*** is deprecated—rather than immediately removed—to provide backward compatibility and an opportunity to bring your configuration into compliance with the new configuration. This option was used for setting the maximum packet threshold for the DPI performance mode.

If your configuration includes enabled performance mode option with **max-packet-threshold** in Junos OS releases 15.1X49-D200 and 19.4R1, AppID concludes the application classification on reaching the lowest value configured in the TCP or UDP inspection limit or in the global offload byte limit, or in the maximum packet threshold for DPI performance mode option.

[See [Application Identification Inspection Limit](#) and [application-identification](#)]

- Starting in Junos OS Release 19.4R1, the **apbr-rule-type** field in the system log message displays the value as **none** if no rule is applied when you have disabled midstream for the application. Updated syslog message sample is as following:

```
<14>1 2019-07-11T03:06:27.276-07:00 pavna RT_FLOW - APPTRACK_SESSION_CLOSE
[junos@2636.1.1.1.2.140 reason="TCP FIN" source-address="4.0.0.1"
source-port="33810" destination-address="5.0.0.1" destination-port="80"]
```

```

service-name="junos-http" application="HTTP" nested-application="YAHOO"
nat-source-address="4.0.0.1" nat-source-port="33810"
nat-destination-address="5.0.0.1" nat-destination-port="80" src-nat-rule-name="N/A"
dst-nat-rule-name="N/A" protocol-id="6" policy-name="1" source-zone-name="trust"
destination-zone-name="untrust" session-id-32="370" packets-from-client="969"
bytes-from-client="50518" packets-from-server="1107" bytes-from-server="1057897"
elapsed-time="253" username="N/A" roles="N/A" encrypted="No"
profile-name="profile1" rule-name="N/A" routing-instance="default"
destination-interface-name="xe-1/1/4.0" uplink-incoming-interface-name="xe-1/1/4.0"
uplink-tx-bytes="50518" uplink-rx-bytes="1057897" category="Web"
sub-category="miscellaneous" apbr-policy-name="sla1" multipath-rule-name="N/A"
src-vrf-grp="N/A" dst-vrf-grp="N/A" dscp-value="N/A" apbr-rule-type="none" ]

```

- Starting in Junos OS Releases 19.4R1, security policy does not support using following applications as dynamic-applications match criteria:
 - junos:HTTPS
 - junos:POP3S
 - junos:IMAPS
 - junos:SMTPS

Software upgrade to the Junos OS Releases 19.4R1 fails during the validation if any of the security policies are configured with junos:HTTPS, junos:POP3S, junos:IMAPS, junos:SMTPS as dynamic-applications as match criteria. We recommend that you remove any configuration that includes these dynamic-applications as match criteria in security policies.

We recommend that you use the **request system software validate package-name** option before upgrading to the above mentioned releases.

Authentication and Access Control

- **Enabling and disabling SSH login password or challenge-response authentication (SRX Series)**—Starting in Junos OS Release 19.4R1, you can disable either the SSH login password or the challenge-response authentication at the **[edit system services ssh]** hierarchy level.

In Junos OS releases earlier than Release 19.4R1, you can enable and disable both SSH login password and the challenge-response authentication simultaneously at the **[edit system services ssh]** hierarchy level.

[See [Configuring SSH Service for Remote Access to the Router or Switch](#).]

Network Management and Monitoring

- **SSHD process authentication logs timestamp (SRX Series)**—Starting in Junos OS Release 19.4R1, the SSHD process authentication logs use only the time zone that is defined in the system time zone. In

Junos OS releases earlier than Release 19.4R1, the SSHD process authentication logs sometimes use the system time zone and the UTC time zone.

[See [Overview of Junos OS System Log Messages](#).]

- **Change in On-box reporting factory-default configuration (SRX1500, SRX4100, SRX4200, SRX4600 and vSRX)**—Starting in Junos OS Release 19.4R1, the factory-default configuration does not include on-box reporting configuration to increase the solid-state drive (SSD) lifetime. You can enable the on-box reporting by configuring the **set security log report** CLI command at **[edit security log]** hierarchy.

[See [Understanding On-Box Logging and Reporting](#).]

- **Change in jnxJsFlowMIB statistics display (SRX Series)**—Starting in Junos OS Release 19.4R1, in a chassis cluster, you can see the statistics on all SPUs of both nodes using the **show snmp mib walk jnxJsFlowMIB** command. In the earlier releases, you can see the statistics only on local SPUs.

[See [SNMP MIB Explorer](#).]

Port Security

- **Configuring source mac filters (SRX300 and SRX550 Services Gateway)**—In this release of Junos OS, fixed an issue that prevented source mac filters from being configured on an interface. The error effected both the **accept-source-mac** and **source-address-filter** statements and resulted in one of the following error messages: **accept-source-mac not allowed in switching mode** and **source mac filters not allowed in switching mode**.

System Logging

- **Preventing system instability during core file generation (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Release 19.4R1 onward, Junos OS checks for available storage space on the Routing Engine before generating core files either on request or because of an assertion condition. This check ensures that your device does not become unstable because of shortage of storage space on the Routing Engine. If the available space is not sufficient, core files are not generated. Instead, Junos OS either displays the **Insufficient Disk space !!! Core generation skipped** message as an output or issues the syslog message **core generation is skipped due to disk full**.

VPNs

- **IKE gateway dynamic distinguished name attributes (SRX Series devices)**—Starting in Junos OS Release 19.4R1, you can now configure only one dynamic distinguished name (DN) attribute among **container-string** and **wildcard-string** at **[edit security ike gateway gateway_name dynamic distinguished-name]** hierarchy. If you try configuring the second attribute after you configure the first attribute, the first attribute is replaced with the second attribute. Before you upgrade your device, you must remove one of the attributes if you have configured both the attributes.

[See [distinguished-name \(Security\)](#) and [Understanding IKE Identity Configuration](#).]

- **CoS Forward Class name (SRX Series devices)**—Starting in Junos OS Release 19.4R1, we have deprecated the CLI option **fc-name** (CoS Forward Class name) in the new **iked** process that displays security associations (SAs) under show command **show security ipsec sa**.

[See [show security ipsec security-associations](#).]

SEE ALSO

[What's New | 265](#)

[Known Limitations | 279](#)

[Open Issues | 282](#)

[Resolved Issues | 284](#)

[Documentation Updates | 295](#)

[Migration, Upgrade, and Downgrade Instructions | 296](#)

Known Limitations

Learn about known limitations in Junos OS Release 19.4R2 for SRX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Application Layer Gateways (ALGs)

- Due to an SSL-FP limitation, the active mode of TLS-based FTP is not supported in Junos OS Release 19.4R1. [PR1450924](#)

DHCP

- On all Junos OS devices, when the DHCP relay is configured with forward-only option, and the DHCP client is terminated on logical tunnel interface which includes multiple IFLs (logical interfaces) and use same VLAN on multiple IFLs of the same It interface, the DHCP relay might fail to send OFFER messages.

Flow-Based and Packet-Based Processing

- For any WiFi configuration change, the access point restarts to make the configuration active. [PR1436587](#)
- The SSID in different WLANs uses the same IP address as the source IP address of the RADIUS packet. [PR1445276](#)
- TKIP is not supported in acn mode. [PR1459160](#)

- Use 512 antireplay window size for IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). Hence there are no out-of-order packets with 512 antireplay window size. [PR1470637](#)
- Use **show security macsec statistics detail** command to see MACsec statistics for control interfaces. The interface specifier does not work. [PR1475371](#)

J-Web

- The CA profile group imported using J-Web is not populated in the Certificate Authority Group initial landing page grid, but all the CA profiles of a group are populated on the Trusted Certificate Authorities landing page. [PR1426682](#)
- When a dynamic application is created for an edited policy rule, the list of services will be blank when the services tab is clicked and then the policy grid will be autorefreshed. As a workaround, create a dynamic application as the last action while modifying the policy rule and click the Save button to avoid loss of configuration changes made to the policy rule. [PR1460214](#)

Platform and Infrastructure

- On SRX5000 line of devices with SPC3 card, if security datapath configuration is applied on tunnel transit traffic, ESP traffic is not captured. [PR1442132](#)

Routing Policy and Firewall Filters

- On SRX5400, SRX5600, and SRX5800 devices, on reth interfaces that are configured as DHCP clients, after a reboot of the device the interface might not get an IP address when you use the default number of DHCP retransmission attempts. When the number of retransmission attempts is increased to 5 or higher, it works fine. [PR1458490](#)

Switching

- SRX300, SRX320, SRX340, SRX345, and SRX550HM devices do not support CoS features such as classification, scheduling, shaping, policing, PCP, and DSCP rewrite in Ethernet switching mode. [PR1476310](#)

VPNs

- When using the operational mode **request security ike debug-enable** command for IKE debugging after using IKE traceoptions with a filename specified in the configuration, the debugs are written to the same filename. [PR1381328](#)
- In SRX Series HA environment with thousands of IPsec tunnels, if rebooting a secondary node, some of the IPsec/IKE tunnels might go down due to the iked getting stuck in cold sync for a long time and not responding to peer DPD messages. Because of above problem, the traffic on down tunnels would be affected until either the SRX Series devices or peer brings up the tunnel back, which typically happens in seconds. [PR1471243](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with 60,000 tunnels up, when RGO failover happens while an IPsec and/or IKE rekey is in progress, those rekeying tunnels might go down and traffic loss might be seen until the tunnel is reestablished. [PR1471499](#)
- In SPC2 and SPC3 mixed mode HA deployments, tunnel per second (TPS) is getting affected while dead peer detection (DCD) is being served on existing tunnels. This limitation is due to a large chunk of CPU being occupied by infrastructure (gencfg) used by IKED to synchronize its DPD state to the backup nodes. [PR1473482](#)

SEE ALSO

[What's New | 265](#)[What's Changed](#)[Open Issues | 282](#)[Resolved Issues | 284](#)[Documentation Updates | 295](#)[Migration, Upgrade, and Downgrade Instructions | 296](#)

Open Issues

IN THIS SECTION

- [Intrusion Detection and Prevention \(IDP\) | 283](#)
- [J-Web | 283](#)
- [Routing Policy and Firewall Filters | 283](#)
- [VPNs | 283](#)

Learn about open issues in Junos OS Release 19.4R2 for SRX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Intrusion Detection and Prevention (IDP)

- When intelligent inspection status changes, syslog is not getting generated on SRX300 and SRX500 line of devices. [PR1448365](#)

J-Web

- On the SRX5000 line of devices, J-Web might not be responsive sometimes when you commit configuration changes after adding a new dynamic application while creating a new firewall rule. J-Web displays a warning while validating the configuration due to dynamic application or any other configuration changes. As a workaround, refresh the J-Web page. [PR1460001](#)

Routing Policy and Firewall Filters

- The SSL reverse proxy feature must be used instead of the SSL inspection feature. SSL inspection on IDP level is being deprecated in favor of SSL reverse proxy. [PR1450900](#)
- If a large number of policies are configured on SRX Series devices and some policies are changed, the traffic that matches the changed policies might be dropped. [PR1454907](#)

VPNs

- On SRX Series devices, in case multiple traffic selectors are configured for a peer with IKEv2 reauthentication, only one traffic selector rekeys at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors are cleared without immediate rekeying. New negotiation of those traffic selectors might be triggered through other mechanisms such as traffic or peer. [PR1287168](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, when SRX Series device is configured in IKEv1 and NAT traversal is active, after a successful IPsec rekey, the IPsec tunnel index might change. In such a scenario, there might be some traffic loss for a few seconds. [PR1409855](#)
- On SRX5400, SRX5600, and SRX5800 devices, during in-service software upgrade (ISSU), the IPsec tunnels flap, causing a disruption of traffic. The IPsec tunnels recover automatically after the ISSU process is completed. [PR1416334](#)
- On the SRX5000 line of devices with SPC3 cards, sometimes IKE SA is not seen on the device when st0 binding on VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)
- With the NCP remote access solution, where IPsec traffic has to be encapsulated as TCP packets, TCP encapsulation for transit traffic is failing. [PR1442145](#)
- Tunnel debugging configuration is not synchronized to backup node. It needs to be configured again after RGO failover. [PR1450393](#)

- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with 60,000 tunnels up, when RGO failover happens while an IPsec and/or IKE rekey is in progress, those rekeying tunnels might go down and traffic loss might be seen until the tunnel is reestablished. [PR1471499](#)
- In SPC2 and SPC3 mixed mode HA deployments, tunnel per second (TPS) is getting affected while dead peer detection (DCD) is being served on existing tunnels. This limitation is due to a large chunk of CPU being occupied by infrastructure (gencfg) used by IKED to synchronize its DPD state to the backup nodes. [PR1473482](#)
- On the SRX5000 line of devices with SPC3 and SPC2 mixed mode, with a very large amount of IKE peers (60,000) with dead peer detection (DPD) enabled, IPsec tunnels might flap in some cases when IKE and IPsec rekeys are happening at the same time. [PR1473523](#)
- On SRX Series devices with SPC3, when overlapping traffic-selectors are configured, multiple IPsec SAs get negotiated with peer device. [PR1482446](#)

SEE ALSO

[What's New | 265](#)

[What's Changed](#)

[Known Limitations | 279](#)

[Resolved Issues | 284](#)

[Documentation Updates | 295](#)

[Migration, Upgrade, and Downgrade Instructions | 296](#)

Resolved Issues

Learn which issues were resolved in the Junos OS main and maintenance releases for SRX Series devices.

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

Resolved Issues: 19.4R2

Application Layer Gateways (ALGs)

- The H323 call with NAT64 cannot be established on the SRX5000 line of devices. [PR1462984](#)
- The flowd or srxpfe process might stop when an ALG creates a gate with an incorrect protocol value. [PR1474942](#)

- SIP messages that need to be fragmented might be dropped by the SIP ALG. [PR1475031](#)
- FTPS traffic might get dropped on SRX Series and MX Series devices if FTP ALG is used. [PR1483834](#)

Authentication and Access Control

- SRX Series: Unified Access Control (UAC) bypass vulnerability (CVE-2020-1637). [PR1475435](#)

Chassis Clustering

- IP monitoring might fail on the secondary node. [PR1468441](#)
- An unhealthy node might become primary in an SRX4600 chassis cluster scenario. [PR1474233](#)
- The **show chassis temperature-thresholds** command displays extensive FPC 0 output. [PR1485224](#)
- If a cluster id of 16 or multiples of 16 is used, the chassis cluster might not come up. [PR1487951](#)

Flow-Based and Packet-Based Processing

- The trusted-ca and root-ca names or IDs should not be the same within an SSL proxy configuration. [PR1420859](#)
- The SPC card might stop on SRX5000 line of devices. [PR1439744](#)
- SRX Series devices upgrades to Junos OS Release 19.4R1 and later versions fail when certain SSL based dynamic applications are used. [PR1444767](#)
- On E1 interface, BERT bits count is not within the range. [PR1445041](#)
- Introduction of default inspection limits to application identification to optimize CPU usage and improve resistance to evasive applications. [PR1454180](#)
- LTE dual CPE support with mPIMs when modem receives disconnect event from ISP; need to increase wait timer. [PR1460102](#)
- A core file is generated when you perform an ISSU on SRX Series devices. [PR1463159](#)
- The pkid process keeps leaking memory on SRX Series devices. [PR1465614](#)
- Tail drop on all ports is observed when any switch-side egress port gets congested. [PR1468430](#)
- RPM test probe failure due to exceeded round-trip time is not working. [PR1471606](#)
- Lookup failure for expected e-mail address in DUT. [PR1472748](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- The dfs-off function is enabled. [PR1475294](#)
- The flowd or srxpfe process might stop when deleting user firewall local authentication table entry. [PR1477627](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- The nsd process pause might be seen during device reboots if dynamic application groups are configured in policy. [PR1478608](#)

- The **show mape rule statistics command** might display negative values. [PR1479165](#)
- On Web proxy, memory leak in association hash table and DNS hash table. [PR1480760](#)
- IMAP curl sessions stuck in the active state if AAMW IMAP block mode is configured. [PR1484692](#)
- Sometimes multiple flowd core files are generated on both nodes of chassis cluster at the same time when changing media MTU. [PR1489494](#)

Installation and Upgrade

- CPU board inlet increases after OS upgrade from Junos OS Release 15.1X49 to Junos OS Release 18.x. [PR1488203](#)
- Has the risk of service interruption on all SRX Series devices with a dual stacked CA server. [PR1489249](#)

Interfaces and Chassis

- Static route through dlo.0 interface is not active. [PR1465199](#)
- All interfaces remain in the down status after the SRX300 line of devices power up or reboot. [PR1488348](#)

Intrusion Detection and Prevention (IDP)

- SNMP queries might cause **commit** or **show** command to fail due to IDP. [PR1444043](#)
- Rogue .gz files in **/var/tmp/sec-download/** might cause an offline security package update to fail. [PR1466283](#)

J-Web

- The Interconnect ports page cannot be used from J-Web because the Type list does not contain any values. [PR1478333](#)

Layer 2 Ethernet Services

- Member links state might be asynchronized on a connection between PE and CE devices in EVPN Active/Active scenario. [PR1463791](#)

Multiprotocol Label Switching (MPLS)

- BGP session might keep flapping between two directly connected BGP peers because of the wrong TCP MSS (maximum segment size) in use. [PR1493431](#)

Network Address Translation (NAT)

- The flowd and srpxfe process might stop when traffic is processed by both ALGs and NAT. [PR1471932](#)
- Issuing the **show security nat source paired-address** command might return an error. [PR1479824](#)

Network Management and Monitoring

- The flowd and srpxfe process might stop immediately after you commit the J-Flowv9 configuration or after you upgrade Junos OS to affected releases. [PR1471524](#)

- SNMP trap coldStart agent-address becomes 0.0.0.0. [PR1473288](#)

Platform and Infrastructure

- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- Physically disconnecting the cable from the fxp0 interface causes hardware monitor failure. [PR1467376](#)
- On the SRX300 line of devices, you might encounter slow loading of Authentication-Table while using user identification. [PR1462922](#)
- On SRX Series devices, Packet Forwarding Engine memory might be used up if the security intelligence feature is configured. [PR1472926](#)
- Supports LLDP on reth interfaces. [PR1473456](#)
- Certificate error while validating configuration during Junos OS upgrade. [PR1474225](#)
- Packet drop might be observed on the SRX300 line of devices when an interface is added to or removed from MACsec. [PR1474674](#)
- The flowd process core files might be seen when there is mixed NAT-T traffic or non-NAT-T traffic with PMI enabled. [PR1478812](#)
- When SRX5K-SPC3s or MX-SPC3s are installed in slots 0 or 1 in SRX5800 or MX960 devices, EMI radiated emissions are observed to be higher than regulatory compliance requirements. [PR1479001](#)
- The RGx might fail over after RG0 failover in a rare case. [PR1479255](#)
- The wl-interface stays in ready status after you execute **request chassis fpc restart** command in Layer 2 mode. [PR1479396](#)
- Recent changes to JDPI's classification mechanism caused a considerable performance regression (more than 30 percent). [PR1479684](#)
- The flowd or srxpfe process might stop when advanced anti-malware service is used. [PR1480005](#)
- Commit does not work after the installation through boot loader. [PR1487831](#)

Routing Policy and Firewall Filters

- Security policies cannot be synchronized between the Routing Engine and the Packet Forwarding Engine on SRX Series devices. [PR1453852](#)
- Some domains are not resolved by the SRX Series devices when using the DNS address book. [PR1471408](#)
- The policy detail does not display the policy statistics counter, even when policy count is enabled. [PR1471621](#)
- Support for dynamic tunnels on SRX Series devices was mistakenly removed. [PR1476530](#)
- **Request security policies check** output shows **policies out of syn** command. [PR1482200](#)
- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. [PR1492436](#)

Routing Protocols

- SSH login might fail if a user account exists in both local database and RADIUS/TACACS+. [PR1454177](#)
- The rpd might stop when both instance-import and instance-export policies contain the as-path-prepend action. [PR1471968](#)

Unified Threat Management (UTM)

- Increase the scale number of a UTM profile or policy for the SRX1500 device, and the SRX4000 and SRX5000 lines of devices. [PR1455321](#)
- The UTMD process pauses after you deactivate UTM configuration with predefined category upgrading used. [PR1478825](#)
- UTM Websense redirect support IPv6 message. [PR1481290](#)

VPNs

- The established tunnels might remain unchanged when an IKE gateway is changed from AutoVPN to Site-to-Site VPN. [PR1413619](#)
- The **show security ipsec statistics** command output displays buffer overflow and wraps around 4,---,---,--- count. [PR1424558](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with IKEv1 enabled IKE, the daemon might generate a core file, when IKE SA is expired and IPsec tunnel associated with the expired IKE SA exists in case of an RGO failover. Daemon recovers eventually. [PR1463501](#)
- IPsec tunnels might lose connectivity after an SRX Series chassis cluster failover when using AutoVPN point-to-multipoint mode. [PR1469172](#)
- IPsec tunnels might flap when one secondary node is coming online after reboot in SRX Series high availability environment. [PR1471243](#)
- The kmd process might crash continually after the chassis cluster failover in the IPsec ADVPN scenario. [PR1479738](#)

Resolved Issues: 19.4R1

Application Layer Gateways

- Unexpected forwarding sessions might appear for tenant SIP ALG traffic in the case of cross-tenants. [PR1409748](#)
- The flowd or srpxfe process might stop in SRX Series devices with chassis cluster when SIP ALG is used. [PR1445766](#)
- Packet loss happens during cold synchronization from the secondary node after rebooting. [PR1448252](#)
- After Layer 3 HA is enabled, ALG H.323 group or resource cannot be synchronized to the peer node correctly. [PR1456709](#)

Application Security

- The AAMW diagnostic script generates incorrect error: **Error: Platform does not support SkyATP: srx300.** [PR1423378](#)
- If automatic application-identification download is configured with a start-time specified, the automatic download stops when the time has progressed to the next year and a reboot is done before the start-time is reached that year. [PR1436265](#)
- SSL-based AppID simplification effort (removal of HTTPS, POP3S, IMAPS, SMTPS). [PR1444767](#)
- The flowd process core files might be generated when traffic hits the AppQoS policy. [PR1446080](#)
- The AAMW diagnostic script generates incorrect error when there is Internet latency: **Error: server unreachable is detected, please make sure port 443 is reachable.** [PR1468114](#)

Authentication and Access Control

- Same-source IP sessions are cleared when the IP entry is removed from the UAC table. [PR1457570](#)

Chassis Clustering

- Hardware failure is seen on both nodes in the output of the **show chassis cluster status** command. [PR1452137](#)
- On SRX Series devices with chassis cluster, the control link remains up even though the control link is actually down. [PR1452488](#)

Class of Service

- Frequent issuance of the **show class-of-service spu statistics** command causes the rtlogd process to be busy. [PR1438747](#)

Flow-Based and Packet-Based Processing

- Throughput or latency performance of TCP traffic is dropped when TCP traffic passes from one logical system to another logical system. [PR1403727](#)
- Packet loss is caused by FPGA back pressure on the SPC3 card. [PR1429899](#)
- VPN traffic fails after the primary node is rebooted or powered off. [PR1433336](#)
- Currently, PMI doesn't support the mirror-filter functionality. If mirror filters are configured, PMI flaps all of the traffic to the regular flow path. [PR1434583](#)
- Intermittent packet drop might be observed if IPsec is configured. [PR1434757](#)
- On an SRX4600 device, core file might be generated and SPM might be in present state. [PR1436421](#)
- Security logs cannot be sent to the external syslog server through TCP. [PR1438834](#)
- Decryption traffic doesn't take PMI path after IPsec rekey (initiated by peer) when the loopback interface is configured as an external interface. [PR1438847](#)
- The IKE pass-through packet might be dropped after a NAT operation on the source. [PR1440605](#)

- New CLI option to show only useful group information for an Active Directory user. [PR1442567](#)
- While checking the flow session XML for source NAT under tenant, there is no value identifier for **tenant-name**. [PR1440652](#)
- The flowd or srpxfe process might stop when processing fragmented packets. [PR1443868](#)
- Junos OS: SRX5000 Series: flowd process crash due to receipt of specific TCP packet (CVE-2019-0064). [PR1445480](#)
- J-Flow version 5 stops working after changing the input rate value. [PR1446996](#)
- Packet loss happens during cold synchronization from secondary node after rebooting. [PR1447122](#)
- On the SRX1500 device, automatic installation is removed from CLI. [PR1447796](#)
- SPC3 talus FPGA stuck on 0x3D or 0x69 golden version. [PR1448722](#)
- Host inbound or host outbound traffic on VR does not work when the SRX5000 line of devices works in SPC3 mixed mode. [PR1449059](#)
- SPU priority does not work when PMI is enabled on the SRX5000 line of devices with an SPC3 card. [PR1449587](#)
- All ingress packets are dropped if the traffic transit network is also the same network for LTE mPIM internal management. [PR1450046](#)
- The flowd or srpxfe process might stop when SSL proxy service is used. [PR1450829](#)
- The AAWM policy rules for IMAP traffic sometimes might not get applied when passed through SRX Series devices. [PR1450904](#)
- FTP data cannot pass through SRX320 4G wireless from FTP server to client. [PR1451122](#)
- Traffic forwarding on Q-in-Q port and VLAN tagging are not observed properly on R0. [PR1451474](#)
- The rpd process might stop and restart with the generation of an rpd core file when committing the configuration. [PR1451860](#)
- The SRX Series devices stop and several core files are generated. [PR1455169](#)
- Added some JP APN settings to default list in LTE mPIM. [PR1457838](#)
- Changing the **RESET** configuration button behavior on the SRX1500 does not work. [PR1458323](#)
- The **security flow traceoptions** fills in with RTSP ALG related information. [PR1458578](#)
- Optimizations were made to improve the connections-per-second performance of SPC3. [PR1458727](#)
- The **security-intelligence** CC feed does not block HTTPS traffic based on SNI. [PR1460384](#)
- The AAMWD process exceeds 85 percent RLIMIT_DATA limitation due to memory leak. [PR1460619](#)
- Added command to clear specified associated client. [PR1461577](#)
- The tunnel packets might be dropped because the gr0.0 or st0.0 interface is wrongly calculated after a GRE or VPN route change. [PR1462825](#)

- Fragmented traffic might get looped between the fab interface in a rare case. [PR1465100](#)
- Track Jbuf double free issue. [PR1465286](#)
- HTTP block message stops working after SNI check for HTTPS session. [PR1465626](#)
- The jbuf process usage might increase up to 99 percent after Junos OS upgrade. [PR1467351](#)
- The rpd process might stop after several changes to the **flow-spec** routes [PR1467838](#)
- FTP data connection might be dropped on dl interface. [PR1468570](#)

Interfaces and Chassis

- SCB4 or SCB3 ZF or XF2 fabric plane retraining is needed after switching the fabric redundancy mode. [PR1427119](#)
- MTU change after a CFM session is up can impact Layer 2 Ethernet ping (loopback messages). If the new change is less than the value in the initial incarnation, then Layer 2 Ethernet ping fails. [PR1427589](#)
- The LACP interface might flap while performing a failover. [PR1429712](#)
- LFM remote loopback is not working as expected. [PR1428780](#)
- The number of mgd processes increases as the mgd processes are not closed properly. [PR1439440](#)
- The fxp0 interface might redirect packets not destined to itself. [PR1453154](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srpxfe process crashes and generates a core file. [PR1437569](#)
- CLI helper text was added to the IDP's attack chain expressions. [PR1438620](#)

J-Web

- The default log query time in J-Web monitoring functionality has been reduced. This increases the responsiveness of the landing pages. [PR1423864](#)
- Phone home UI portal to be removed from SRX Series devices. [PR1428717](#)
- Some error messages might be seen when using J-Web. [PR1446081](#)
- The idle-timeout for J-Web access does not work properly. [PR1446990](#)
- J-Web fails to display the traffic log in event mode when stream mode host is configured. [PR1448541](#)
- Editing destination NAT rule in J-Web introduces a non-configured routing-instance field. [PR1461599](#)
- The **Go** button within the J-Web Monitor->Events view now correctly refreshes the logs even when using a blank search query. [PR1464593](#)
- J-Web security resources dashboard widget was not being populated correctly. [PR1464769](#)

Layer 2 Ethernet Services

- DHCP requests might get dropped in a DHCP relay scenario. [PR1435039](#)

Network Address Translation

- The nsd process might stop when SNMP queries deterministic NAT pool information. [PR1436775](#)
- Flowd process core files are generated in the device while testing NAT PBA in AA mode. [PR1443148](#)
- RTSP resource session is not found during NAT64 static mapping. [PR1443222](#)
- A port endian issue in SPU messages between SPC3 and SPC2 results in one redundant NAT binding being created in central point when one binding is allocated in SPC2 SPC. [PR1450929](#)
- Packet loss is observed when multiple source NAT pools and rules are configured. [PR1457904](#)

Network Management and Monitoring

- MIB OID dot3StatsDuplexStatus shows wrong status. [PR1409979](#)
- Snmpd process might generate core files after restarting NSD process by using the **restart network-security gracefully** command. [PR1443675](#)
- Control links are logically down on SRX Series devices with chassis cluster running Junos OS Release 12.3X48. [PR1458314](#)

Platform and Infrastructure

- On SRX4600 platform, when manual RGO failover is performed, sometimes node0 (the original primary node) stays in secondary-hold status for a long time and cannot change back to secondary status. [PR1421242](#)
- Packet drops, replication failure, or ksyncd stops might be seen on the logical system of a Junos OS device after Routing Engine switchover. [PR1427842](#)
- The PICs might go offline and split brain might be seen when interrupt storm happens on internal Ethernet interface em0 or em1. [PR1429181](#)
- REST API does not work properly. [PR1430187](#)
- Unable to launch J-Web when the device is upgraded through USB image. [PR1430941](#)
- Packet Forwarding Engine crashes might be seen on SRX1500 platform. [PR1431380](#)
- The ksyncd process might stop and restart. [PR1440576](#)
- The configured RPM probe server hardware timestamp does not respond with the correct timestamp to the RPM client. [PR1441743](#)
- ARP resolution might fail after **ARP HOLD NHs** are added and deleted continuously [PR1442815](#)
- The SRX300 line of device does not have MIB that can retrieve the fan status. [PR1443649](#)
- IS-IS adjacencies between the GE link are not up. [PR1446533](#)

- The flowd process might stop on SRX Series devices when chassis cluster and IRB interface are configured. [PR1446833](#)
- The **show security flow session** command fails with error messages when SRX4100 or SRX4200 has around 1 million routing entries in FIB. [PR1445791](#)
- LACP cannot work with the encapsulation **flexible-ethernet-services** configuration. [PR1448161](#)
- On certain MPC line cards, cm errors need to be reclassified. [PR1449427](#)
- The REST service might become nonresponsive when the REST API receives several continuous HTTP requests. [PR1449987](#)
- VM core files might be generated if the configured sampling rate is more than 65,535. [PR1461487](#)
- Loading CA certificate causes pkid core file to be generated. [PR1465966](#)

Routing Policy and Firewall Filters

- The NSD process might stop due to a memory corruption issue. [PR1419983](#)
- Two ipfd processes appear in **ps** command and the process pauses. [PR1444472](#)
- During commit, the **nsd_vrf_group_config_ls** log messages are displayed. [PR1446303](#)
- Traffic log shows wrong custom-application name when the **alg ignore** option is used in application configuration. [PR1457029](#)
- The NSD process might get stuck and cause problems. [PR1458639](#)
- The policy detail does not print out policy statistics counter, even when policy count is enabled. [PR1471621](#)

Services Applications

- The flowd process stops when the SRX5000 line of devices works in SPC3 mixed mode with one SPC3 card or seven SPC2 cards. [PR1448395](#)
- The srxpfe lcore-slave core files are generated. [PR1460035](#)

Unified Threat Management

- The **show security utm web-filtering status** command now provides additional context when the status of EWF is **down**. [PR1426748](#)
- Memory issue due to SSL proxy whitelist or whitelist URL category. [PR1430277](#)
- Adjust core allocation ratio for on-box antivirus. [PR1431780](#)

VLAN Infrastructure

- ISSU failed from Junos OS Release 18.4R2.7 to Junos OS Release 19.4, with secondary node PICs in present state after upgrading to Junos OS Release 19.4. [PR1468609](#)

VPNs

- IPsec SA inconsistent on SPCs of node0 and node1 in SRX Series devices with chassis cluster. [PR1351646](#)
- After RG1 failover, IKE phase 1 SA is getting cleared. [PR1352457](#)
- With a large number of IPsec tunnels established, a few tunnels might fail during rekey negotiation if the SRX Series device initiates the rekey. [PR1389607](#)
- Displaying incorrect port number when scale is 1,000 on IKEv1 AutoVPN tunnels. [PR1399147](#)
- The IKE and IPsec configuration under groups is not supported in this release. [PR1405840](#)
- The IKED process stops due to a misconfiguration. [PR1416081](#)
- The VPN tunnel might flap when IKE and IPsec rekey happen simultaneously. [PR1421905](#)
- Old tunnel entries are also seen when new tunnel negotiation happens from peer device after change in IKE gateway configuration at peer side. [PR1423821](#)
- IPsec packet throughput might be impacted if NAT-T is configured and the fragmentation operation of post fragment happens. [PR1424937](#)
- Tunnel does not come up after changing configurations from IPv4 to IPv6 tunnels in the script with **gateway lookup failed** error. [PR1431265](#)
- P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. [PR1432434](#)
- IPsec rekey triggers for when sequence number in AH and ESP packet is about to exhaust. [PR1433343](#)
- P1 or P2 SAs are deleted after RG0 failover. [PR1433355](#)
- IPsec SA in and out key sequence number update missing after cold synchronization. [PR1433424](#)
- Sequence number reset to zero while recovering SA after SPC3 or flowd stops or reboots. [PR1433568](#)
- The kmd log shows resource temporarily unavailable repeatedly and VPNs might be down. [PR1434137](#)
- The IKED process stops on SRX5000 line of devices with SPC3 when IPsec VPN or IKE is configured. [PR1443560](#)
- The IPsec VPN traffic drop might be seen on SRX Series devices with NAT-T scenario. [PR1444730](#)
- Sometimes old SAs are not deleted after rekey and the number of IPsec tunnels shows up more than the configured tunnels. [PR1449296](#)
- Some IPsec tunnels flap after RGs failover on SRX5000 line of devices. [PR1450217](#)
- The VPN flaps on the primary node after a reboot of the secondary node. [PR1455389](#)
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. [PR1455951](#)

- IPsec VPN tunnels are losing routes for traffic selector randomly while tunnel is still up, causing traffic loss of these IPsec VPN tunnels. [PR1456301](#)
- On all SRX Series devices, the **no-anti-replay** option does not take effect immediately. Traffic is not sent out through IPsec VPN after upgrading to Junos OS Release 18.2 or later. [PR1461793](#)
- The IPsec VPN tunnels cannot be established if overlapped subnets are configured in traffic selectors. [PR1463880](#)

SEE ALSO

What's New 265
What's Changed
Known Limitations 279
Open Issues 282
Documentation Updates 295
Migration, Upgrade, and Downgrade Instructions 296

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 296](#)

This section lists the errata and changes in Junos OS Release 19.4R1 for the SRX Series documentation.

Feature Guides Are Renamed As User Guides

- Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the *BGP Feature Guide* is now the *BGP User Guide*. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this [Junos OS Documentation](#).

SEE ALSO

[What's New | 265](#)

[What's Changed](#)

[Known Limitations | 279](#)

[Open Issues | 282](#)

[Resolved Issues | 284](#)

[Migration, Upgrade, and Downgrade Instructions | 296](#)

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.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life 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 15.1X49, 17.3, 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from one Junos OS Release to the next release or one release after the next release. For example you can upgrade from Junos OS Release 15.1X49 to Release 17.3 or 17.4, Junos OS Release 17.4 to Release 18.1 or 18.2, and from Junos OS Release 18.1 to Release 18.2 or 18.3 and so on.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. 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 about software installation and upgrade, see the [Installation and Upgrade Guide for Security Devices](#).

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

SEE ALSO

What's New 265
What's Changed
Known Limitations 279
Open Issues 282
Resolved Issues 284
Documentation Updates 295

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the [High Availability User Guide](#).

For additional information about using ISSU on security devices, see the [Chassis Cluster User Guide for SRX Series Devices](#).

For information about ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\) Web application](#).

Licensing

Starting in 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that have been developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see [Juniper Flex Program](#).

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps us in exploring software feature information to find the right software release and product for your network. <https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home


NOTE: To obtain information about the components that are supported on the devices, and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. apps.juniper.net/compliance/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

21 April 2023—Revision 19, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 July 2022—Revision 18, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 January 2022—Revision 17, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 December 2021—Revision 16, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 October 2021—Revision 15, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 July 2021—Revision 14, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 April 2021—Revision 13, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 March 2021—Revision 12, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 11, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 November 2020—Revision 10, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 October 2020—Revision 9, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 October 2020—Revision 8, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 7, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 July 2020—Revision 6, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 July 2020—Revision 5, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 June 2020—Revision 4, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 May 2020—Revision 3, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 May 2020—Revision 2, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 May 2020—Revision 1, Junos OS Release 19.4R2— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 April 2020—Revision 14, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 April 2020—Revision 13, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 March 2020—Revision 12, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 March 2020—Revision 11, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 March 2020—Revision 10, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

5 March 2020—Revision 9, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 February 2020—Revision 8, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 January 2020—Revision 7, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 January 2020—Revision 6, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 January 2020—Revision 5, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 January 2020—Revision 4, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 January 2020—Revision 3, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 December 2019—Revision 2, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 December 2019—Revision 1, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.