

Release Notes: Junos[®] OS Release 19.1R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion

9 February 2023

Contents	Introduction 11
	Junos OS Release Notes for ACX Series 11
	What's New 12
	What's New in Release 19.1R3 12
	What's New in Release 19.1R2 12
	What's New in Release 19.1R1 12
	What's Changed 14
	What's Changed in 19.1R3 14
	What's Changed in 19.1R2 15
	What's Changed in 19.1R1 15
	Known Limitations 17
	General Routing 17
	Open Issues 18
	General Routing 18
	MPLS 20
	Resolved Issues 20
	Resolved Issues: 19.1R3 20
	Resolved Issues: 19.1R2 22
	Resolved Issues: 19.1R1 24

Documentation Updates | 26

Migration, Upgrade, and Downgrade Instructions | 26

Upgrade and Downgrade Support Policy for Junos OS Releases | 27

Junos OS Release Notes for EX Series Switches | 28

What's New | 28

What's New in Release 19.1R3-S2 | 30

What's New in Release 19.1R3 | 30

What's New in Release 19.1R2 | 30

What's New in Release 19.1R1 | 30

What's Changed | 35

What's Changed in Release 19.1R3-S4 | 36

What's Changed in Release 19.1R3 | 36

What's Changed in Release 19.1R2 | 36

What's Changed in Release 19.1R1 | 38

Known Limitations | 39

EVPN | 40

General Routing | 40

Security | 41

Infrastructure | 41

Virtual Chassis | 41

Open Issues | 41

Authentication and Access Control | 42

General Routing | 42

Infrastructure | 43

Interfaces and Chassis | 44

Junos Fusion Provider Edge | 44

Junos Fusion Satellite Software | 44

Multicast | 44

Network Management and Monitoring | 44

Platform and Infrastructure | 44

Routing Protocols | 45

Subscriber Access Management | 45

Resolved Issues | 45**Resolved Issues: 19.1R3 | 46****Resolved Issues: 19.1R2 | 49****Resolved Issues: 19.1R1 | 55****Documentation Updates | 58****Migration, Upgrade, and Downgrade Instructions | 59****Upgrade and Downgrade Support Policy for Junos OS Releases | 59****Junos OS Release Notes for Junos fusion for enterprise | 60****What's New | 61****Release 19.1R3 New and Changed Features | 61****Release 19.1R2 New and Changed Features | 61****Release 19.1R1 New and Changed Features | 61****What's Changed | 62****Known Limitations | 62****Junos fusion for enterprise | 63****Open Issues | 63****Junos fusion for enterprise | 63****Resolved Issues | 64****Resolved Issues: Release 19.1R3 | 65****Resolved Issues: Release 19.1R2 | 65****Resolved Issues: Release 19.1R1 | 65****Documentation Updates | 66****Migration, Upgrade, and Downgrade Instructions | 66****Basic Procedure for Upgrading Junos OS on an Aggregation Device | 67****Upgrading an Aggregation Device with Redundant Routing Engines | 68****Preparing the Switch for Satellite Device Conversion | 69****Converting a Satellite Device to a Standalone Switch | 70****Upgrade and Downgrade Support Policy for Junos OS Releases | 70****Downgrading from Junos OS | 71****Junos OS Release Notes for Junos fusion for provider edge | 72****What's New | 72****What's New in Release 19.1R3 | 73****What's New in Release 19.1R2 | 73****What's New in Release 19.1R1-S4 | 74**

What's New in Release 19.1R1	74
What's Changed	75
What's Changed in Release 19.1R3	75
What's Changed in Release 19.1R2	75
What's Changed in Release 19.1R1	75
Known Limitations	76
Open Issues	76
Junos fusion for provider edge	77
Resolved Issues	77
Resolved Issues: 19.1R3	78
Resolved Issues: 19.1R2	78
Resolved Issues: 19.1R1	78
Documentation Updates	79
Migration, Upgrade, and Downgrade Instructions	79
Basic Procedure for Upgrading an Aggregation Device	80
Upgrading an Aggregation Device with Redundant Routing Engines	82
Preparing the Switch for Satellite Device Conversion	83
Converting a Satellite Device to a Standalone Device	84
Upgrading an Aggregation Device	86
Upgrade and Downgrade Support Policy for Junos OS Releases	87
Downgrading from Junos OS Release 19.1	87
Junos OS Release Notes for MX Series 5G Universal Routing Platform	88
What's New	88
What's New in Release 19.1R3-S2	89
What's New in Release 19.1R3	89
What's New in Release 19.1R2	90
What's New in Release 19.1R1-S1	90
What's New in Release 19.1R1	90
What's Changed	111
What's Changed in Release 19.1R3	111
What's Changed in Release 19.1R2	113
What's Changed in Release 19.1R1	117

Known Limitations | 121

- Fault Management | 121**
- General Routing | 121**
- Infrastructure | 123**
- Interfaces and Chassis | 123**
- MPLS | 123**
- Platform and Infrastructure | 123**
- Routing Protocols | 124**
- Software-Defined Networking (SDN) | 124**
- Subscriber Management and Services | 124**

Open Issues | 125

- Class of Service (CoS) | 126**
- EVPN | 126**
- Forwarding and Sampling | 127**
- General Routing | 127**
- High Availability (HA) and Resiliency | 134**
- Infrastructure | 134**
- Interfaces and Chassis | 135**
- Junos fusion Enterprise | 136**
- Junos fusion Provider Edge | 136**
- Layer 2 Ethernet Services | 136**
- MPLS | 136**
- Network Management and Monitoring | 137**
- Platform and Infrastructure | 137**
- Routing Protocols | 139**
- Subscriber Access Management | 141**
- User Interface and Configuration | 141**
- VPNs | 141**

Resolved Issues | 142

- Resolved Issues: 19.1R3 | 142**
- Resolved Issues: 19.1R2 | 155**
- Resolved Issues: 19.1R1 | 181**

Documentation Updates | 195

- Spanning Tree Protocol User Guide | 195**

Migration, Upgrade, and Downgrade Instructions | 195

Basic Procedure for Upgrading to Release 19.1 | 196

Procedure to Upgrade to FreeBSD 11.x based Junos OS | 197

Procedure to Upgrade to FreeBSD 6.x based Junos OS | 199

Upgrade and Downgrade Support Policy for Junos OS Releases | 201

Upgrading a Router with Redundant Routing Engines | 201

Downgrading from Release 19.1 | 202

Junos OS Release Notes for NFX Series | 202

What's New | 203

What's New in Release 19.1R3 | 203

What's New in Release 19.1R2 | 203

What's New in Release 19.1R1 | 204

What's Changed | 206

What's Changed in Release 19.1R3 | 207

What's Changed in Release 19.1R2 | 207

What's Changed in Release 19.1R1 | 207

Known Limitations | 207

Security | 208

Platform and Infrastructure | 208

Virtual Network Functions (VNFs) | 208

Open Issues | 209

Interfaces | 209

High Availability | 210

Platform and Infrastructure | 210

Virtualized Network Functions (VNFs) | 211

Resolved Issues | 212

Resolved Issues: 19.1R3 | 212

Resolved Issues: 19.1R2 | 213

Resolved Issues: 19.1R1 | 215

Documentation Updates | 216

Migration, Upgrade, and Downgrade Instructions | 216

Upgrade and Downgrade Support Policy for Junos OS Releases | 216

Basic Procedure for Upgrading to Release 19.1 | 217

Junos OS Release Notes for PTX Series Packet Transport Routers | 218

What's New | 219

What's New in 19.1R3 | 219

What's New in 19.1R2 | 219

What's New in 19.1R1 | 220

What's Changed | 231

What's Changed in Release 19.1R3-S4 | 232

What's Changed in 19.1R3 | 232

What's Changed in 19.1R2 | 232

What's Changed in 19.1R1 | 234

Known Limitations | 236

General Routing | 237

Open Issues | 238

General Routing | 238

Interfaces and Chassis | 240

Routing Protocols | 240

Resolved Issues | 240

Resolved Issues: 19.1R3 | 241

Resolved Issues: 19.1R2 | 242

Resolved Issues: 19.1R1 | 245

Documentation Updates | 247

Migration, Upgrade, and Downgrade Instructions | 248

Basic Procedure for Upgrading to Release 19.1 | 248

Upgrade and Downgrade Support Policy for Junos OS Releases | 251

Upgrading a Router with Redundant Routing Engines | 251

Junos OS Release Notes for the QFX Series | 252

What's New | 252

What's New in Release 19.1R3-S2 | 253

What's New in Release 19.1R3 | 253

What's New in Release 19.1R2 | 253

What's New in Release 19.1R1 | 255

What's Changed | 267

What's Changed in Release 19.1R3 | 268

What's Changed in Release 19.1R2 | 269

What's Changed in Release 19.1R1	270
Known Limitations	272
EVPN	273
General Routing	273
Infrastructure	274
Layer 2 Features	274
MPLS	274
Platform and Infrastructure	274
Routing Protocols	276
Security	276
Virtual Chassis	276
Open Issues	276
EVPN	277
General Routing	277
Infrastructure	280
Interfaces and Chassis	280
Layer 2 Features	280
MPLS	281
Platform and Infrastructure	281
Routing Protocols	281
Resolved Issues	282
Resolved Issues: 19.1R3	282
Resolved Issues: 19.1R2	288
Resolved Issues: 19.1R1	298
Documentation Updates	303
Migration, Upgrade, and Downgrade Instructions	303
Upgrading Software on QFX Series Switches	304
Installing the Software on QFX10002-60C Switches	306
Installing the Software on QFX10002 Switches	306
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	307
Installing the Software on QFX10008 and QFX10016 Switches	309
Performing a Unified ISSU	313
Preparing the Switch for Software Installation	314

Upgrading the Software Using Unified ISSU | 314

Upgrade and Downgrade Support Policy for Junos OS Releases | 316

Junos OS Release Notes for SRX Series | 317

What's New | 318

What's New in Release 19.1R3 | 318

What's New in Release 19.1R2 | 319

What's New in Release 19.1R1 | 319

What's Changed | 329

Changes in Behavior and Syntax: Release 19.1R3 | 330

Changes in Behavior and Syntax: Release 19.1R2 | 333

Changes in Behavior and Syntax: Release 19.1R1 | 335

Known Limitations | 336

Application Security | 338

Installation and Upgrade | 338

J-Web | 338

Platform and Infrastructure | 338

Switching | 338

VPNs | 338

Open Issues | 339

Flow-Based and Packet-Based Processing | 339

Intrusion Detection and Prevention (IDP) | 340

J-Web | 340

Platform and Infrastructure | 340

Routing Policy and Firewall Filters | 340

SSL Proxy | 340

VPNs | 341

Resolved Issues | 341

Resolved Issues: 19.1R3 | 342

Resolved Issues: 19.1R2 | 346

Resolved Issues: 19.1R1 | 355

Documentation Updates | 360

Migration, Upgrade, and Downgrade Instructions | 361

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 361

Upgrading Using ISSU | 363

Licensing | 363

Compliance Advisor | 363

Finding More Information | 364

Documentation Feedback | 364

Requesting Technical Support | 366

Self-Help Online Tools and Resources | 366

Creating a Service Request with JTAC | 367

Revision History | 367

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, and Junos fusion.

These release notes accompany Junos OS Release 19.1R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 12](#)
- [What's Changed | 14](#)
- [Known Limitations | 17](#)
- [Open Issues | 18](#)
- [Resolved Issues | 20](#)
- [Documentation Updates | 26](#)
- [Migration, Upgrade, and Downgrade Instructions | 26](#)

These release notes accompany Junos OS Release 19.1R3 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 19.1R3](#) | 12
- [What's New in Release 19.1R2](#) | 12
- [What's New in Release 19.1R1](#) | 12

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

What's New in Release 19.1R3

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 19.1R3.

What's New in Release 19.1R2

There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 19.1R2.

What's New in Release 19.1R1

Authentication, Authorization, and Accounting (AAA) (RADIUS)

- **Support for SFTP global disablement (ACX Series)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the statement **sftp-server** at the **[edit system services ssh]** hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections were globally enabled by default.

[See [Configuring sftp-server](#).]

Platform and Infrastructure

- **DMA recovery mechanism (ACX Series)**—A recovery mechanism has been introduced that is triggered in case the router enters an Idle state on any DMA channels. The recovery mechanism reboots the Packet Forwarding Engine to recover from Idle state.

The following recovery message is logged in the Routing Engine syslog message:

```
CHASSISD_FPC_ASIC_ERROR: <FPC 0> ASIC Error detected errorno 0x0000ffff FPC
restart initiated
```

The following recovery message is logged in the Packet Forwarding Engine syslog message:

```
BCM DMA channel error detected
Resetting the PFE
```

Routing Protocols

- **Support for BGP graceful shutdown (ACX Series)**— Starting in Junos OS Release 19.1R1, graceful traffic migration from one BGP next hop to another is supported, without traffic interruption. Also, BGP administrative shutdown communication can be sent to the BGP peer.

You can configure both **graceful-shutdown** and **shutdown** statements at the **[edit protocols bgp]**, **[edit protocols bgp group group-name]**, and **[edit protocols bgp group group-name neighbor address]** hierarchy levels.

NOTE: Graceful shutdown is disabled by default.

[See [graceful-shutdown \(Protocols BGP\)](#), [shutdown \(Protocols BGP\)](#).]

SEE ALSO

What's Changed 14
Known Limitations 17
Open Issues 18
Resolved Issues 20
Documentation Updates 26
Migration, Upgrade, and Downgrade Instructions 26

What's Changed

IN THIS SECTION

- [What's Changed in 19.1R3 | 14](#)
- [What's Changed in 19.1R2 | 15](#)
- [What's Changed in 19.1R1 | 15](#)

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

What's Changed in 19.1R3

General Routing

- **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.
- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **`show configuration system arp | grep unicast-mode-on-expire`**.

[See [arp](#).]

What's Changed in 19.1R2

Interfaces and Chassis

- **Support for disabling RS-FEC (ACX6360-OX)**—By default, Junos OS software enables or disables forward error correction based on plugged-in optics. Starting with Junos OS Release 19.1R2, on ACX6360-OX routers used as Transponders, you can now disable Ethernet FEC, also known as RS-FEC or FEC91. Previously, RS-FEC was enabled by default and could not be disabled.

[See [fec](#).]

Network Management and Monitoring

- **The `show system schema` command and `<get-yang-schema>` RPC require specifying an output directory (ACX Series)**—Starting in Junos OS Release 19.1R2, when you issue the `show system schema` operational mode command in the CLI or execute the `<get-yang-schema>` RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the `output-directory` command option in the CLI or the `<output-directory>` element in the RPC. In earlier releases, you can omit the `output-directory` argument when requesting a single module to display the module in standard output.

What's Changed in 19.1R1

Interfaces and Chassis

- **Support for creating Layer 2 logical interface independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, and later releases, ACX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

Network Management and Monitoring

- **NETCONF <kill-session> operation returns different values in <rpc-error> when the session identifier is equal to the current session ID (ACX Series)**—Starting in Junos OS Release 19.1R1, when you execute the <kill-session> NETCONF operation and the session identifier is equal to the current session ID, the values of the <error-type> and <error-tag> elements in the resulting <rpc-error> are **application** and **invalid-value**, respectively. In earlier releases, the <error-type> and <error-tag> values are **protocol** and **operation-failed**.

[See [<kill-session>](#).]

- **sysName.0 MIB object displays the fully qualified domain name (ACX Series)**—Starting in Junos OS Release 19.1R1, the sysName.0 MIB object displays the fully qualified domain name. That is, if the hostname and domain name are configured on the system, both will show up for the sysName.0 MIB object: **host-name.domain-name**. Previously, only the hostname showed up.

[See [show snmp mib](#).]

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when change in number of supported history records is detected (ACX Series and MX Series)**—In Junos OS Release 19.1R1, when Ethernet connectivity fault management starts, it detects the number of history records supported by the existing performance monitoring history database and if there is any change from the number of history records supported (that is, 12) in Release 19.1R1, then the existing performance monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

Routing Policy and Firewall Filters

- **Firewall filters with IPv6 match conditions not supported on ACX6360-OR routers**—Starting in Junos OS Release 19.1R1, firewall filters with Internet Protocol version 6 (IPv6) match conditions are not supported at the **[firewall family inet6 filter name]** hierarchy level on ACX6360-OR routers. Note that different Junos OS releases might have different support limits, for example 19.2R1.

SEE ALSO

[What's New | 12](#)

[Known Limitations | 17](#)

[Open Issues | 18](#)

[Resolved Issues | 20](#)

[Documentation Updates | 26](#)

[Migration, Upgrade, and Downgrade Instructions | 26](#)

Known Limitations

IN THIS SECTION

- [General Routing | 17](#)

Learn about known limitations in this release 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

- For an ACX5448, the theoretical limit of the ARP learning rate is approximately 150 ARP resolutions per second per logical interface. [PR1343221](#)
- All PTP packets go to the best-effort queue instead of the network control queue. [PR1361315](#)
- Junos telemetry interface infrastructure does not support interface filtering capability. After you enable a particular sensor for telemetry, it is turned on for all the interfaces. [PR1371996](#)
- For an et- interface, only the PRE_FEC_SD defect is raised and no OTN alarm is raised. [PR1371997](#)
- On an ACX6360 router, the **static-cak** command encryption does not work between two ACX-OX transponder nodes. [PR1389802](#)
- For ACX6360 TIC, you have only 8 CFP2-DCO ports. The **show chassis beacon** or **request chassis beacon** command for port numbers larger than 7 does not work because the ports do not exist. However, no error is reported. **user@router> request chassis beacon fpc 0 pic-slot 1 port 15 on FPC 0 PIC 1 PORT 15 ON regress@node> show chassis beacon fpc 0 pic-slot 1 port-range lower-limit 0 upper-limit 15 FPC 0 PIC 1 PORT 0 ON FPC 0 PIC 1 PORT 1 ON FPC 0 PIC 1 PORT 2 ON FPC 0 PIC 1 PORT 3 ON FPC 0 PIC 1 PORT 4 ON FPC 0 PIC 1 PORT 5 ON FPC 0 PIC 1 PORT 6 ON FPC 0 PIC 1 PORT 7 ON FPC 0 PIC 1 PORT 8 ON FPC 0 PIC 1 PORT 9 ON FPC 0 PIC 1 PORT 10 OFF FPC 0 PIC 1 PORT 11 OFF FPC 0 PIC 1 PORT 12 OFF FPC 0 PIC 1 PORT 13 OFF FPC 0 PIC 1 PORT 14 OFF FPC 0 PIC 1 PORT 15 ON** [PR1399335](#)
- ACX Series routers support only 900 joins of IGMPv3 users per second. [PR1448146](#)

SEE ALSO

[What's New | 12](#)

[What's Changed | 14](#)
[Open Issues | 18](#)
[Resolved Issues | 20](#)
[Documentation Updates | 26](#)
[Migration, Upgrade, and Downgrade Instructions | 26](#)

Open Issues

IN THIS SECTION

- [General Routing | 18](#)
- [MPLS | 20](#)

Learn about open issues in this release 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

- In the earlier scenario where only the PCS is down and PMD is up, the status is not sent to the MAC and no interruption is caused. However, the status is picked up as part of the 1- second periodic, where the link goes down for a second. In spite of taking the link status from the MAC, you can read the PHY status to avoid the intermittent link down status observed in the PCS as it is a sticky read. [PR1223457](#)
- Upon classifying the Layer 3 packets, DSCP is not preserved and is lost at the egress due to the limitations of a chipset. [PR1322142](#)
- On ACX5448 devices, when a 1-Gigabit SFP transceiver is installed, autonegotiation is enabled by default. There is no functional impact. Only the **show interface extensive** command shows the **Auto-negotiation** field as disabled. [PR1343679](#)
- If the **set interfaces ae<>xaggregated-ether-options link-speed** command is configured in the router, the aggregated Ethernet interface remains down after rebooting. The following error message is observed:

```
===== /kernel: kernel did not add link ae1, link speeds differ 1000000000 10000000000
/kernel: bundle ae1.0: link xe-0/3/0 not added due to speed mismatch.
```

[PR1357012](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)

- The queue statistics are not as expected after configuring the IFD and logical interface shaping with the transmit rate and scheduler-map. [PR1488935](#)
- Loopback status is not shown for OT interfaces on the CLI (available from vty only). [PR1358017](#)
- The SD threshold is normally lower than the SF threshold, so that as errors increase, the SD condition is encountered first. For the ACX6360 optical links, there is no guard code to prevent you from setting the SD threshold above the SF threshold, which would cause increasing errors to trigger the SF alarm before the SD alarm. This does not cause any issues on systems with correctly provisioned SD or SF thresholds. [PR1376869](#)
- On the ACX6360-OR router, enhancement is needed for FRR BER threshold SNMP support. [PR1383303](#)
- On the ACX6360 and PTX10001 routers, Tx power cannot be configured using the + sign. [PR1383980](#)
- The switchover time observed is more than 50 ms under certain soak test conditions with an increased scale and a multiprotocol multirouter topology. [PR1387858](#)
- The circuit cross-connect logs are not compressed after rotation. [PR1398511](#)
- A `jnxIfOtnOperState` trap notification is sent for all ot- interface. [PR1406758](#)
- Packets transmitted in a queue are not as expected when testing the ieee-802.1ad inner classifier at the ingress and ieee-802.1ad rewrite at the egress with various events. [PR1422515](#)
- Protocols get forwarded when you use the nonexisting SSM map source address in IGMPv3 instead of pruning. [PR1435648](#)
- Memory leaks are expected in this release. [PR1438358](#)
- dcpfe generates core files when you enable or disable MACsec through the Toby scripts. [PR1479710](#)
- On the ACX5448 platform, multiple hardware I2C failures might be observed because of the intermittent I2C access failures on the main board internal switches. A system reboot is required to clear this state. [PR1429047](#)
- The FPC might crash with the em2 interface configuration during initialization. [PR1429212](#)
- On the ACX5048 Series router, the egress queue statistics are not working for the aggregated Ethernet interfaces. [PR1472467](#)
- On the ACX5448 Series router, the Layer 2 VPN with the **ethernet-ccc input-vlan-map** or **output-vlan-map** interface might cause traffic to be silently discarded. [PR1485444](#)
- Need to disable the **inet6-precedence** statement under the **[edit class-of-service rewrite-rules]** hierarchy. [PR1344340](#)
- Memory utilization enhancement is needed. [PR1481151](#)
- On the ACX5000 router, the IEEE 802.1p priority and DEI values in the locally generated VLAN-based IP packets might be changed when sourced from the IRB interface. [PR1490966](#)

MPLS

- When two directly connected BGP peers are established over MPLS LSP, if the MTU of the IP layer is smaller than the MTU of the MPLS layer and the BGP packets from the host have the DF bit set, the BGP session might keep flapping because of the use of wrong TCP-MSS. [PR1493431](#)

SEE ALSO

[What's New | 12](#)

[What's Changed | 14](#)

[Known Limitations | 17](#)

[Resolved Issues | 20](#)

[Documentation Updates | 26](#)

[Migration, Upgrade, and Downgrade Instructions | 26](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3 | 20](#)
- [Resolved Issues: 19.1R2 | 22](#)
- [Resolved Issues: 19.1R1 | 24](#)

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: 19.1R3

General Routing

- On the ACX5000 line of routers, the following error message is observed: **MacDrainTimeOut and bcm_port_update failed: Internal error** [PR1284590](#)
- The bcmDPC task is high even though the **START_BY_START** interrupt flag is set to 0. [PR1329656](#)

- On the ACX5448 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet interfaces, link fault signaling does not work. [PR1401718](#)
- High CPU for fxpc processes with CoS, changes on aggregated Ethernet interfaces. [PR1407098](#)
- Optics comes with Tx enabled by default. As the port is administratively disabled, the port is stopped. However, as the port has not been started, it does not disable Tx. [PR1411015](#)
- The l2cpd process might crash and generate a core file when the interfaces flaps. [PR1431355](#)
- On the ACX Series router, packet drop might be seen when chain composite is enabled for Layer 3 VPN. [PR1439317](#)
- Interface on ACX1100 remains down when you use SFP-1FE-FX (740-021487). [PR1439384](#)
- On the ACX5448 router, the **encap flexible VLAN tagging** is not supported with MPLS family. [PR1445046](#)
- On the ACX5448-M router, fans might not run at the correct speed. [PR1448884](#)
- FPC might generate core files after changing the configuration of PTP or Synchronous Ethernet. [PR1451950](#)
- On the ACX5448 device, the Layer 2 VPN with the Ethernet encapsulation type stops passing traffic after a random port is added with VLAN configuration. [PR1456624](#)
- The traffic might be discarded silently during link recovery in an open Ethernet access ring with ERPS configured. [PR1459446](#)
- On the ACX5000 line of routers, the SNMP MIB walk for jnxOperatingTemp does not return anything for FPCs in all the new Junos OS Releases. [PR1460391](#)
- On the ACX Series platforms, the LLDP neighbor does not come up on LAG interfaces after the software upgrade to Junos OS Release 18.2R3-S1. [PR1461831](#)
- Not able to add more than 16 links in a lag. [PR1463253](#)
- RED drops on interface even without any congestion. [PR1470619](#)
- On the ACX5048 routers, the snmp polling gets stalled after the link flaps or the sfp gets replaced and the following error message is observed: **ACX_COS_HALP(acx_cos_gport_sched_set_strict_priority:987): Failed to detach.** [PR1455722](#)
- Memory leakage on the l2cpd process might lead to l2cpd crash. [PR1469635](#)
- The ACX5048 routefails to display the RFC2544 test statistics for the reflected packets. [PR1451759](#)
- FPC might continuously crash after deactivating or activating the loopback filter or rebooting the system after configuring the loopback filter. [PR1477740](#)

Interfaces and Chassis

- The status of the MC-AE interface might appear as unknown when you add the subinterface as a part of the VLAN on the peer MC-AE node. [PR1479012](#)

Routing Protocols

- The rpd might crash continuously due to memory corruption in the IS-IS setup. [PR1455432](#)

Resolved Issues: 19.1R2

General Routing

- The logical interfaces might not come up if scaled logical interfaces exist. [PR1229492](#)
- The 1G copper module interface shows "Link-mode: Half-duplex". [PR1286709](#)
- The fxpc process might use high CPU on ACX5000 after upgrade. [PR1360452](#)
- FPC showing high CPU due to PIC PERIODIC. [PR1360844](#)
- On an ACX Series ring topology, after the link between the ACX Series and MX Series devices flaps, the VPLS RI on the PE device (MX Series) has no MAC address for the CE device over I2circuit. over Layer 2 circuit. [PR1360967](#)
- On an ACX Series devices, the LED on GE interface goes down when speed 10M is added. [PR1385855](#)
- Traffic over the aggregated Ethernet physical interface might get filtered with the filter on one child logical interface on ACX Series. [PR1407855](#)
- The **show services inline stateful-firewall flow** or **show services inline stateful-firewall flow extensive** command might cause the memory leak. [PR1408982](#)
- ACX Series router drops DNS responses that contain an underscore. [PR1410062](#)
- VPLS traffic might stop across ACX5000 with the aggregated Ethernet interface. [PR1412042](#)
- Junos OS PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- The 40-Gigabit Ethernet FEC ACX5448 is ON by default. Need to align with MX Series and QFX Series platforms where FEC is NONE. [PR1414649](#)
- Commit error while configuring firewall with term having log/syslog and accept actions. [PR1417377](#)
- On the ACX448 devices, BFD timer value does not match the configured 900 ms with multiplier 3. It's showing 6000 with multiplier 3 instead for most of the sessions. [PR1418680](#)
- COS table error can sometimes cause traffic outages and SNMP timeouts if the optic is plugged out and inserted back. [PR1418696](#)
- High CPU usage on fxpc process might be seen on ACX5000 platform. [PR1419761](#)
- Slow copy image speed to ACX5448. [PR1422544](#)

- Traffic might forward to wrong bridge-domain if you change interface encapsulation from Ethernet-bridge to VLAN-bridge. [PR1423610](#)
- The JUNIPER_SOURCE LR4T2 optics might not work properly on ACX5448 platforms. [PR1424814](#)
- The FPC/fxpc crash might be observed. [PR1427362](#)
- Chassisd can crash with unsupported hcos configuration when MX104 is used as fusion aggregation device. [PR1430076](#)
- On ACX5448 devices, upon reboot of MC_LAG peer, when peer comes up (but before hardware comes up) there is a 10-20 second traffic hit on node1. [PR1430910](#)
- Auto-RP mapping might periodically time-out on ACX5448 platforms. [PR1432889](#)
- ACX5448 might malfunction in encapsulating small packets if egress link is 40-Gigabit Ethernet or 100-Gigabit Ethernet. [PR1434900](#)
- On ACXp Series latforms, no-vrf-propagate-ttl might not work after you activate or deactivate COS configuration. [PR1435791](#)
- On ACX5448 routers, DHCP packets are not transparent over Layer 2 CIRCUIT. [PR1439518](#)
- ACX5448: Packet buffer error from Packet Forwarding Engine leads to memory leak when IGMP is sent from NNI AC in Layer 2 circuit and VPLS. [PR1442901](#)
- RED drops might be seen after link flaps or CoS configuration changes. [PR1443466](#)
- On ACX Series routers, auto-exported route between VRFs might not reply for ICMP echo requests. [PR1446043](#)
- On the ACX5448 Series router, Layer 2 circuit stops forwarding traffic after LDP flap. [PR1448899](#)
- On the ACX5048 Series router, Layer 2 circuit with backup-neighbor configuration stops passing traffic after link flap and moves from backup neighbor to primary. [PR1449681](#)
- On the ACX5448 Series router, the FPC crashes due to segmentation fault. [PR1453766](#)
- FDB not flushing, silently discards the traffic in the Ethernet ring scenario. [PR1459446](#)

Class of Service

- The dfwd crash can be seen with forwarding-class configuration in policers. [PR1436894](#)

Interfaces and Chassis

- Family inet of the unnumbered interface might be getting deleted when deleting one of the IPs of the binding interface. [PR1412534](#)
- Upgrade from Junos OS Release 17.4R1 and previous releases results in cfmd core file. [PR1425804](#)

Layer 2 Ethernet Services

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

MPLS

- MPLS ingress LSPs for LDP link protection are not coming up after disable/enable of MPLS. [PR1432138](#)

Platform and Infrastructure

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

Routing Protocols

- Loopback address exported into other VRF instance might not work on ACX Series platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even though the LDP neighbor's MAC changes. [PR1451217](#)

Resolved Issues: 19.1R1

General Routing

- SNMP MIB walk/get/set on jnxDomCurrentTable and jnxDomNotifications might fail on ACX Series platforms. [PR1076943](#)
- ACX Series routers support dual-tagged through untagged packets Layer 3 traffic. [PR1307666](#)
- ACX5000: fpc0 **acx_rt_ip_uc_lpm_install:LPM route add failed** error. Reason: Invalid parameter after configuring lpm-profile. [PR1365034](#)
- VPLS with **vlan-id-list** is not working properly in some releases when the link between a PE device and a CE device is an aggregated Ethernet interface with a single member link and child physical interface flap. [PR1365894](#)
- **LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** prints while commit on configuration prompt. [PR1376665](#)
- On ACX5448, channelized ET interface of 25-Gigabit Ethernet interface will not come up after chassis-control restart. [PR1379288](#)
- The L2circuit might stop forwarding traffic when core interface flapping happens. [PR1381487](#)
- Timestamp is incorrect for BER statistics after clearing. [PR1386253](#)
- The **request chassis beacon** CLI command is not working for PIC slot 1 (that is, CFP2 ports). [PR1386711](#)
- ACX 5448:100-Gigabit link FEC is enabled by default on 100-Gigabit LR4. [PR1389518](#)
- On ACX Series platforms, the **forwarding-option dhcp-relay forward-only** command stops working and the DHCP packets are dropped. [PR1392261](#)

- Certain builds of Junos OS do not allow you to upgrade or commit configuration changes when the SI service interface is used. [PR1393729](#)
- MTU is not properly applied, and the output ping mpls l2circuit sweep is giving lower values than expected. [PR1393947](#)
- This model of egress VPLS filter and the output of with **physical-interface-specific** semantic is only to be used to cater to use cases where there is a need to install a "physical-port-based" filter in the egress firewall. [PR1395362](#)
- ACX5048 RPM RFC 2544 benchmarking test is failing to start. [PR1395730](#)
- Error message **ACX_PFE_ERROR: dnx_cfm_bd_endpoint_create: Failed to destroy the remote endpoint, Endpoint id 0x2001001, Entry not found** been logged. [PR1397878](#)
- Error message **ACX_ASIC_PROGRAMMING_ERROR: dnx_cfm_bd_endpoint_create: Failed to create the local endpoint Invalid parameter** been logged on peer node. [PR1397951](#)
- **Output packet error Count** is 40-Gigabit Ethernet and 100-Gigabit Ethernet ports. [PR1398270](#)
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. [PR1398398](#)
- Dynamic tunnels are not supported on ACX Series routers. [PR1398729](#)
- ACX5448: Not able to configure bridge domain more than 1024, using 100-Gigabit and aggregated Ethernet interface in BD. [PR1399214](#)
- FPC might crash after offline/online MIC-3D-16CHE1-T1-CE-H. [PR1402563](#)
- ACX5448 supports TrTCM policer configuration parameters of RFC 4115. [PR1405798](#)
- Aggregated Ethernet interface TWAMP history statistics verification on client is not getting expected "Request Timed Out" error. [PR1411344](#)
- Number of inet-arp policers implemented on ACX5000 has been increased from 16 to 64. [PR1413807](#)
- Swap memory is not initialized on boot on ACX5048. [PR1415898](#)

Infrastructure

- The error of **jlaunchd: disk-monitoring is thrashing, not restarted** might be seen. [PR1380032](#)

Services Applications

- The spd might crash when **any-ip** is configured in the from clause of the NAT rule with the static translation type. [PR1391928](#)

SEE ALSO

[What's New | 12](#)

[What's Changed | 14](#)

Known Limitations 17
Open Issues 18
Documentation Updates 26
Migration, Upgrade, and Downgrade Instructions 26

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for the ACX Series.

SEE ALSO

What's New 12
What's Changed 14
Known Limitations 17
Open Issues 18
Resolved Issues 20
Migration, Upgrade, and Downgrade Instructions 26

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 27](#)

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 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](#).

SEE ALSO

[What's New | 12](#)

[What's Changed | 14](#)

[Known Limitations | 17](#)

[Open Issues | 18](#)

[Resolved Issues | 20](#)

[Documentation Updates | 26](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- [What's New | 28](#)
- [What's Changed | 35](#)
- [Known Limitations | 39](#)
- [Open Issues | 41](#)
- [Resolved Issues | 45](#)
- [Documentation Updates | 58](#)
- [Migration, Upgrade, and Downgrade Instructions | 59](#)

These release notes accompany Junos OS Release 19.1R3 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 Release 19.1R3-S2 | 30](#)
- [What's New in Release 19.1R3 | 30](#)
- [What's New in Release 19.1R2 | 30](#)
- [What's New in Release 19.1R1 | 30](#)

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 19.1R3: EX2300, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in Release 19.1R3-S2

Routing Policy and Firewall Filters

- **Loopback firewall filter scale optimization (QFX5120-48Y and EX4650)**—Starting with Junos OS Release 19.1R3-S2, you can configure up to 768 loopback filter terms for IPv6, and up to 1152 terms for IPv4. To do so, you configure an ingress firewall filter, apply it to the loopback interface, and then enable the **loopback-firewall-optimization** command at the **[edit chassis]** hierarchy level (this triggers the packet forwarding engine (PFE) to restart).

Terms that include a reserved multicast destination (for example 224.0.0.x/24) and terms with a time-to-live (TTL) of 0/1 are not directly supported. Instead, you need to configure a separate filter for these terms. For example, to count OSPF packets on the loopback interface, you would create a separate filter with terms for the protocol (OSPF) to count packets destined to a reserved multicast address (such as 224.0. 0.6).

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

What's New in Release 19.1R3

- There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.1R3.

What's New in Release 19.1R2

- There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.1R2.

What's New in Release 19.1R1

Hardware

- **Support for SFP transceivers on 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ uplink module (model number: EX-UM-4SFPP-MR) (EX4300-48MP and EX4300-48MP-S switches)**—Starting with Junos OS Release 19.1R1, the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet SFP+ uplink module (model number: EX-UM-4SFPP-MR) for EX4300-48MP and EX4300-48MP-S switches support SFP transceivers. You do not need to configure 1-Gigabit Ethernet speed on the uplink module to support SFP transceivers; the uplink module automatically detects the transceiver and creates the interfaces accordingly.

Authentication, Authorization, and Accounting (AAA)

- **RADIUS over TLS (using RadSec) support (EX4300 switches)**—Starting in Junos OS Release 19.1R1, EX4300 switches support RadSec. The RadSec protocol provides secure transport of RADIUS authentication and accounting data across untrusted networks using Transport Layer Security (TLS) over TCP as the transport protocol.

[See [RADIUS over TLS \(RADSEC\)](#).]

- **Support for SFTP global disablement (EX Series)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the **sftp-server** statement at the **[edit system services ssh]** hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections are globally enabled by default.

[See [Configuring sftp-server](#).]

Dynamic Host Configuration Protocol

- **Increased number of DHCP relay servers supported (EX9200 switches)**—Starting in Junos OS Release 19.1R1, server groups can include up to 32 active server IP addresses in a DHCPv4 or DHCPv6 relay configuration.

[See [Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups](#).]

EVPNs

- **Support for proxy MAC addresses in an ARP request (EX9200)**—Starting in Junos OS Release 19.1R1, provider edge (PE) devices in an EVPN network that support ARP proxy can use a proxy MAC address in the ARP replies message to a host. When a PE device receives an ARP request or NDP request, it searches its MAC-IP address binding database for the requested IP address. If the device finds the MAC-IP address entry in its database, it responds to the request with the proxy MAC address. The proxy MAC address is derived from the IRB interface in an EVPN network with edge-routed bridging overlay and from the manually configured MAC address in a centrally routed bridging overlay. If the device does not find an entry, the PE device replaces the MAC and IP address from the customer edge (CE) device in the ARP request with the proxy MAC and IP address of the IRB interface. This allows for enhanced security (that is, Layer 3 filtering) deployments on Layer 3 gateway for both inter-VLAN and intra-VLAN traffic will be routed.

To enable this feature, configure the **proxy-mac [irb | proxy-mac-address]** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy or at the **[edit routing-instances routing-instance-name bridge-domains domain_name]** hierarchy.

[See [ARP and NDP Request with a proxy MAC address](#).]

- **EVPN-VXLAN support (EX4300-48MP switches)**—Starting in Junos OS Release 19.1R1, the EX4300-48MP switch, which functions as a Layer 2 VXLAN gateway in an EVPN-VXLAN network, supports the following features:
 - Active/active multihoming
 - Proxy ARP use and ARP suppression, and NDP use and NDP suppression on non-IRB interfaces
 - Ingress node replication for broadcast, unknown unicast, and multicast (BUM) traffic forwarding

[See [EVPN User Guide](#).]

Interfaces and Chassis

- **Support for 1-Gbps speed on 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module (EX4300-48MP)**—Starting with Junos OS Release 19.1R1, the 4-port 1-Gigabit Ethernet/10-Gigabit Ethernet uplink module (EX-UM-4SFPP-MR) on EX4300-48MP switches supports 1-Gbps speed. You do not need to explicitly configure 1-Gbps speed on the uplink module, because the uplink module automatically identifies the installed 1-Gigabit SFP transceivers and creates the interfaces accordingly.

NOTE: The status LED of the 1-Gigabit Ethernet uplink module port is solid green (instead of blinking green) because of a device limitation. However, there is no impact on device functionality.

[See [speed \(Ethernet\)](#).]

- **Support to channelize 100-Gigabit Ethernet port as four 25-Gigabit Ethernet ports in uplink module (EX4300-48MP)**—Starting with Junos OS Release 19.1R1, in the 2-port QSFP+/1-port QSFP28 uplink module on EX4300-48MP switches, you can channelize the 100-Gigabit Ethernet port to operate as four independent 25-Gigabit Ethernet ports by using breakout cables.

[See [Setting the Mode on 2-port QSFP+/1-port QSFP28 Uplink Module \(CLI Procedure\)](#).]

- **Improved performance of small packets (EX Series)**—Starting in Junos OS Release 19.1R1, the EX9200-40XS and EX9200-12QS line cards provide improved performance of small packets (with a minimum packet size of 64 bytes) in the transmit direction. To enable this feature, reduce the number of active ports (at the PIC level) to the following maximum numbers:
 - Sixteen 10-Gbps ports
 - Four 40-Gbps ports
 - Two 100-Gbps ports (when the line card is in 240-Gbps mode)
 - Three 100-Gbps ports (when the line card is in 400-Gbps mode)

To configure the number of active ports, use the existing command **set chassis fpc slot pic slot number-of-ports *number-of-active-ports***.

NOTE: The command does not change packet performance at the Packet Forwarding Engine level; it improves packet performance in the transmit direction at the port level only.

Junos Telemetry Interface

- **Export of data associated with the Junos kernel through Junos Telemetry Interface (JTI) (EX9200, EX9251, and EX9253)**—Starting in Junos OS Release 19.1R1, you can export data associated with the Junos kernel through remote procedure calls (gRPC) and JTI. Kernel telemetry data includes information on Veriexec state, graceful Routing Engine switchover (GRES), in-service software upgrade (ISSU), and

Routing Engine ifstate. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

Junos kernel sensors introduced in Junos OS Release 19.1R1 support both periodical and ON_CHANGE streaming. The following Junos kernel resource paths support periodical streaming only:

- /junos/kernel-ifstate/dead-ifstates-cnt
- /junos/kernel-ifstate/alive-ifstates-cnt
- /junos/kernel-ifstate/delayed-unrefs-cnt
- /junos/kernel-ifstate/delayed-unrefs-max

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

Operation, Administration, and Maintenance (OAM)

- **LFM support on EX2300 and EX3400 switches** —Starting with Junos OS Release 19.1R1, EX2300 and EX3400 switches support OAM link fault management (LFM). OAM LFM can be configured 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 switches support the following OAM LFM features:
 - Discovery and link monitoring
 - Remote fault detection
 - Remote loopback

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

Routing Policy and Firewall Filters

- **Support for matching IPv6 source addresses from an inet6 egress interface (EX4300)**—Starting in Junos OS Release 19.1R1, you can configure a firewall filter on an IPv6 egress interface to match specified IPv6 source or destination addresses—for example, to protect a third-party device connected to the switch.

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

Routing Protocols

- **Support for BGP graceful shutdown (EX Series)**— Starting in Junos OS Release 19.1R1, graceful traffic migration from one BGP next hop to another is supported, without traffic interruption. Also, BGP administrative shutdown communication can be sent to the BGP peer.

You can configure both **graceful-shutdown** and **shutdown** statements at the **[edit protocols bgp]**, **[edit protocols bgp group group-name]**, and **[edit protocols bgp group group-name neighbor address]** hierarchy levels.

NOTE: Graceful shutdown is disabled by default.

[See [graceful-shutdown \(Protocols BGP\)](#), [shutdown \(Protocols BGP\)](#).]

Security

- **MPLS scaling enhancements (EX4600 switches)**—Starting in Junos OS Release 19.1R1, MPLS scaling is enhanced on the EX4600 switch. For instance, you can increase the scale from its default 1024 to 8192. 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](#).]

Software Installation and Upgrade

- **Phone-home client (EX4300-MP switches)**—Starting with Junos OS Release 19.1R1, you can use either the legacy DHCP-options-based ZTP or the phone-home client (PHC) to provision software for the switch. If the switch boots up and there are DHCP options received from the DHCP server for ZTP, ZTP resumes. If DHCP options are not present, PHC is attempted. PHC enables the switch to securely obtain bootstrapping data, such as a configuration or software image, with no user intervention other than having to physically connect the switch to the network. When the switch first boots, PHC connects to a redirect server, which will redirect to a phone-home server to get the configuration or software image. To initiate either DHCP-options-based ZTP or PCH, either the switch must be in a factory-default state, or you can issue the **request system zeroize** command.

[See [Understanding the Phone-Home Client.](#)]

- **Phone-home client (EX2300-MP switches)**—Starting with Junos OS Release 19.1R1, you can use either the legacy DHCP-options-based ZTP or the phone-home client (PHC) to provision software for the switch. If the switch boots up and there are DHCP options received from the DHCP server for ZTP, ZTP resumes. If DHCP options are not present, PHC is attempted. PHC enables the switch to securely obtain bootstrapping data, such as a configuration or software image, with no user intervention other than having to physically connect the switch to the network. When the switch first boots, PHC connects to a redirect server, which will redirect to a phone-home server to get the configuration or software image. To initiate either DHCP-options-based ZTP or PCH, either the switch must be in a factory-default state, or you can issue the **request system zeroize** command.

[See [Understanding the Phone-Home Client.](#)]

SEE ALSO

[What's Changed | 35](#)

[Known Limitations | 39](#)

[Open Issues | 41](#)

[Resolved Issues | 45](#)

[Documentation Updates | 58](#)

[Migration, Upgrade, and Downgrade Instructions | 59](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.1R3-S4 | 36](#)
- [What's Changed in Release 19.1R3 | 36](#)
- [What's Changed in Release 19.1R2 | 36](#)
- [What's Changed in Release 19.1R1 | 38](#)

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

What's Changed in Release 19.1R3-S4

Infrastructure

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

What's Changed in Release 19.1R3

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 19.1R3, 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.

What's Changed in Release 19.1R2

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, and later, EX Series switches support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces

to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

- **Logical Interface is created along with physical Interface by default (MX Series routers, EX Series switches, and QFX Series switches)**—In Junos OS Release 19.1R2 and later, logical interface is created on **ge**, **et**, **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, earlier when you view the **show interfaces** command, by default, only the physical interface (**ge-0/0/0**), is displayed. Now, the logical interface (**ge-0/0/0.16386**) is also displayed.

- **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.

Layer 2 Features

- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—From Junos OS Release 19.1R2, the configuration statement **input-native-vlan-push** at the **[edit interfaces interface-name]** hierarchy level is introduced. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement **input-vlan-map** with a **push** operation is configured.

[See [input-native-vlan-push](#).]

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (EX Series)**—Starting in Junos OS Release 19.1R2, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.

Routing Protocols

- **Change in the default behavior of advertise-from-main-vpn-tables configuration statement**—BGP now advertises EVPN routes from the main bgp.evpn.0 table. You can no longer configure BGP to advertise the EVPN routes from the routing instance table. In earlier Junos OS Releases, BGP advertised EVPN routes from the routing instance table by default.

[See [advertise-from-main-vpn-tables](#).]

What's Changed in Release 19.1R1

EVPN

- Starting with Junos OS Release 19.1R1, the **no-arp-suppression** configuration statement is no longer supported on EX Series and QFX Series devices.

Interfaces and Chassis

- **No support for performance monitoring on AE Interfaces (EX4300)**—Y.1731 performance monitoring (PM) over aggregated Ethernet interfaces is not supported on EX4300 switches. [See [sla-iterator-profile](#).]

Network Management and Monitoring

- **sysName.0 MIB object displays the fully qualified domain name (EX Series)**—Starting in Junos OS Release 19.1R1, the sysName.0 MIB object displays the fully qualified domain name. That is, if the hostname and domain name are configured on the system, both names are displayed for the sysName.0 MIB object: **host-name.domain-name**. Previously, only the host name was displayed.

[See [show snmp mib](#).]

- **NETCONF <kill-session> operation returns different values in <rpc-error> when the session identifier is equal to the current session ID (EX Series)**—Starting in Junos OS Release 19.1R1, when you execute the <kill-session> NETCONF operation and the session identifier is equal to the current session ID, the values of the <error-type> and <error-tag> elements in the resulting <rpc-error> are **application** and **invalid-value**, respectively. In earlier releases, the <error-type> and <error-tag> values are **protocol** and **operation-failed**.

[See [<kill-session>](#).]

Security

- **Syslog or log action on firewall drops packets (EX4600 switches)**—Starting in Junos OS Release 19.1R1, if you configure a syslog or log action on an ingress firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.

User Interface and Configuration

- **Options for monitor traffic interfaces statement added (EX Series)**—Starting in Junos OS Release 19.1R1, the options **write-file** and **read-file** under the **monitor traffic** command are included in the visible CLI.

[See [monitor traffic](#).]

SEE ALSO

What's New 28
Known Limitations 39
Open Issues 41
Resolved Issues 45
Documentation Updates 58
Migration, Upgrade, and Downgrade Instructions 59

Known Limitations

IN THIS SECTION

- [EVPN | 40](#)
- [General Routing | 40](#)
- [Security | 41](#)

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 19.1R3 for the EX Series.

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

EVPN

- When a VLAN uses an IRB interface as the routing interface, the **vlan-id** parameter must be set to "none" to ensure a proper traffic routing. This issue is platform-independent. [PR1287557](#)

General Routing

- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter will not be installed. [PR1362609](#)
- On EX4650 with 288,000 MAC scale, the Routing Engine **show ethernet-switching table summary** command output shows the learned scale entries after a delay of around 60 seconds. [PR1367538](#)
- Junos OS might hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. As a workaround, you can power cycle the device. [PR1385970](#)
- After you upgrade to Junos OS Release 19.3, the system hangs after executing the **request system software add /var/tmp/<image-gz>** command. As a workaround, power cycle the device. It might resume normal functioning. [PR1405629](#)

Security

- –On EX4600 switches, if a syslog or log action is configured on a firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.

Infrastructure

- If Junos OS panics experiences a file system-related panic, such as 'dup alloc', recovery through the OAM shell might be needed. From the OAM shell, run **fsck** on the root volume until it is marked clean. Only at this point is it safe to reboot to the normal volume. [PR1444941](#)

Virtual Chassis

- A Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on EX4600 and EX4300 Virtual Chassis, a minimal traffic disruption or traffic loop (greater than 2 seconds) might be seen. [PR1347902](#)

SEE ALSO

What's New	 28
What's Changed	 35
Open Issues	 41
Resolved Issues	 45
Documentation Updates	 58
Migration, Upgrade, and Downgrade Instructions	 59

Open Issues

IN THIS SECTION

- [Authentication and Access Control](#) | 42
- [General Routing](#) | 42
- [Infrastructure](#) | 43
- [Interfaces and Chassis](#) | 44
- [Junos Fusion Provider Edge](#) | 44

- Junos Fusion Satellite Software | 44
- Multicast | 44
- Network Management and Monitoring | 44
- Platform and Infrastructure | 44
- Routing Protocols | 45
- Subscriber Access Management | 45

Learn about open issues in Junos OS Release 19.1R3 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

- Before running the **load ssl-certificate path PATHNAME** command, configure the path using the CLI **set protocols dot1x ssl-certificate-path PATHNAME** command if the default pathname is not **/var/tmp/**. [PR1431086](#)

General Routing

- The ARP queue limit has been changed from 100 pps to 3000 pps. [PR1165757](#)
- On an EX2300 switch, the output of the **show chassis routing-engine** command might display an incorrect value of **mac reset** in the **last reboot reason** field. [PR1331264](#)
- Traffic flooding happens instead of routing when VRRP scaled more than 150. [PR1371520](#)
- An EX4300 switch configured with a firewall filter on lo0 and dhcp-security on VLAN simultaneously might drop legitimate DHCP renew requests from clients on the corresponding VLANs. [PR1376454](#)
- On an EX9208 switch, a few 10-Gigabit Ethernet 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](#)
- After the Media Access Control security (MACsec) session is deleted, the corresponding interfaces might lose their MACsec function if LACP is enabled on them and the **exclude lacp** statement is configured under the **[security macsec]** hierarchy. [PR1378710](#)
- On an EX4650 switch, DCPFE does not come up in some instances when the switch is abruptly powered off and powered on. Power cycle the device or reboot the host to recover the device. [PR1393554](#)
- There is a possibility of seeing multiple reconnect logs, JTASK_IO_CONNECT_FAILED, during the device initialization. There is no functionality impact because of these messages. These messages can be ignored. [PR1408995](#)

- On an EX9200 switch with MC-LAG configuration and other features enabled, there is a loss of approximately 20 seconds during restart of the routing daemon. This traffic loss varies with the configuration that is done. [PR1409773](#)
- On EX2300 and EX4650 platforms, uRPF check in strict mode might not work properly. [PR1417546](#)
- For an EVPN instance with integrated routing and bridging (IRB) interface, the proxy ARP and ARP suppression is enabled by default. With EVPN proxy ARP, ARP suppression enabled, and Junos OS Release 17.2R3 and later or 17.3R3 and later (correlated service release is also affected), the kernel process on master Routing Engine might crash because of a software defect on packet handling. This is a rare issue. [PR1443903](#)
- On an EX9214 switch, if the MACsec-enabled link flaps after reboot, the error `errorlib_set_error_log(): err_id(-1718026239)` is observed. [PR1448368](#)
- On EX3400 Virtual Chassis, during reboot or upgrade, because of a high CPU load in the slow path of FXPC, the TCP keepalive message is not sent. Hence, it is observed that some Virtual Chassis members might take longer to join the Virtual Chassis. [PR1467707](#)
- On EX2300 switches with **mac-limit** and **drop-and-log** action configured, when the limit threshold is reached, a syslog message is triggered, but no SNMP trap is generated. [PR1482709](#)

Infrastructure

- On EX3400 and EX2300 line of switches, during zero-touch provisioning (ZTP) with configuration and image upgrade with FTP, image upgrade is successful, but sometimes VM core files might be generated. [PR1377721](#)
- On EX Series switches, if you configure a large-scale number of firewall filters on some interfaces, FPC crash and generation of core files might be seen. [PR1434927](#)
- On an EX4300 switch, the CLI configuration **on-disk-failure** is not supported. [PR1450093](#)

Interfaces and Chassis

- After GRES, the VSTP port cost on aggregated Ethernet interfaces might get changed, leading to a topology change. [PR1174213](#)

Junos Fusion Provider Edge

- On a Junos fusion for enterprise system, intermediate traffic drop is seen between AD and SD when sFlow is enabled on an ingress interface. This is not seen always. When sFlow is enabled, the original packet is getting corrupted for those packets that hit the sFlow filter. This is because a few packets transmitted from the egress of AD1 are short of FCS (4 bytes) + 2 bytes of data, due to which the drop occurs. It is seen that the normal data packets are of size 128 bytes (4 bytes FCS + 14 bytes Ethernet header + 20 bytes IP header + 90 bytes data), while the corrupted packet is 122 bytes (14 bytes Ethernet header + 20 bytes IP header + 88 bytes data). [PR1450373](#)

Junos Fusion Satellite Software

- In Junos fusion SP setup, EX4300 acting as satellite devices is generating temperature sensor alarm on multiple satellite devices modules connected to same aggregation device. [PR1466324](#)

Multicast

- IGMP query packets might be duplicated between Layer 2 interfaces when IGMP snooping is enabled. [PR1391753](#)

Network Management and Monitoring

- Trace files are not closed properly; as a result, writing of traceoptions to the log file suddenly stops. [PR1380764](#)

Platform and Infrastructure

- There are multiple failures when events such as node reboots and ICL or ICCP flaps; and even with enhanced convergence configured there is no guarantee that subsecond convergence will be achieved. [PR1371493](#)
- The image size grows over a period of time and subsequently storage might be insufficient to install images when upgrading EX2300 and EX3400 platforms. The upgrade might fail with the error message **not enough space to unpack**. [PR1464808](#)

- In a Virtual Chassis with the VRRP master initially and if the VRRP IRB is disabled, then traffic destined to the VRRP VIP on the new VRRP master might be dropped if this is in the traffic path. [PR1491348](#)

Routing Protocols

- Error messages `pimd_rtrequest_v4(1133), IS_MASTER_RE: 1, Process: rpd, RTM_ID: 5, error: 17, errmsg: rt exists; ifindex = 340` are cosmetic and expected logs. These logs are not harmful and have no functional impact, but they just show the state of PIM register messages. These logs are already LOG_DEBUG for external builds, you do not need to make any change in any of the components. [PR1371431](#)
- In 19.1R3 release, MUX state of LACP interface might not change when **force-up** configured. [PR1484523](#)

Subscriber Access Management

- The authd reuses address quickly before jdhcpd has completely cleaned up the old subscriber, which results in syslog errors: `jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815`. [PR1402653](#)

SEE ALSO

What's New	 28
What's Changed	 35
Known Limitations	 39
Resolved Issues	 45
Documentation Updates	 58
Migration, Upgrade, and Downgrade Instructions	 59

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3](#) | [46](#)
- [Resolved Issues: 19.1R2](#) | [49](#)
- [Resolved Issues: 19.1R1](#) | [55](#)

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

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

Resolved Issues: 19.1R3

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 extended port even if it comes in with different DSCP marking. [PR1433252](#)

EVPN

- The ARP request or neighbor solicitation message might be sent back to the local segment by the DF router. [PR1459830](#)
- The rpd might crash after the EVPN-related configuration is changed. [PR1467309](#)

Forwarding and Sampling

- Memory leak might be observed in the l2ald process. [PR1455034](#)
- Type 1 ESI/AD route might not be generated locally on EVPN PE in the **all-active** mode. [PR1464778](#)

General Routing

- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- Micro-BFD 3x100 ms flap is seen upon inserting a QSFP transceiver into other port. [PR1435221](#)
- IRB over VTEP unicast traffic might get dropped on EX9200 platforms. [PR1436924](#)
- The rpd might generate a core file during router bootup due to a file pointer issue because there are two code paths that can close the file. [PR1438597](#)
- On EX4600, the Virtual Chassis does not come up after you replace the Virtual Chassis port fiber connection with a DAC cable. [PR1440062](#)
- MAC addresses learned on RTG might not be aged out after a Virtual Chassis member is rebooted. [PR1440574](#)
- On EX3400 with half-duplex mode on 10M or 100M speed at medium traffic rates due to PHY side MAC buffer inconsistent state, MAC pause frames are seen on the port and egress traffic on the port stops to flow. [PR1452209](#)
- The l2ald and eventd processes are hogging 100 percent after **clear ethernet-switching table** command is issued. [PR1452738](#)
- A firewall filter might not be able to be applied in a particular Virtual Chassis or Virtual Chassis Fabric member as TCAM space running out of space. [PR1455177](#)

- Packet drop might be seen after removing and reinserting the SFP transceiver on the 40-Gigabit Ethernet uplink module ports. [PR1456039](#)
- Link-up delay and traffic drop might be seen on mixed service provider Layer 2/Layer 3 and enterprise style Layer 2 type configurations. [PR1456336](#)
- The syslog **timeout connecting to peer database-replication** is generated when the **show version detail** command is issued. [PR1457284](#)
- Overtemperature SNMP trap messages are displayed after update even though the temperatures are within the system thresholds. [PR1457456](#)
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used. [PR1458559](#)
- The fxpc process might crash due to BGP IPV6 session flaps. [PR1459759](#)
- On EX2300 and EX3400 platforms, storage space limitation leads to image installation failure during phone home. [PR1460087](#)
- MAC addresses learned on RTG might not be aged out after aging time. [PR1461293](#)
- RTG link faces nearly 20 seconds down during backup node rebooting. [PR1461554](#)
- Configuring any combination of VLANs and interfaces under VSTP/MSTP might prevent any VSTP/MSTP-related configuration from being committed. [PR1463251](#)
- There are some command lines to disable MAC learning and some of them 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](#)
- On EX2300 switches, an FXPC core file is generated after mastership election based on user priority. [PR1465526](#)
- The broadcast and multicast traffic might be dropped over an IRB or a LAG interface in a EX Virtual Chassis scenario. [PR1466423](#)
- The MAC move message might have an incorrect "from" interface when MAC moves rapidly. [PR1467459](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- FPC might be down when configuring **vxlan-routing**. [PR1468736](#)
- With autonegotiation enabled, EX3400 advertises only 100 Mbps whenever we configure the 100-Mbps speed. [PR1471931](#)
- The shaping of CoS does not work after a reboot. [PR1472223](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)
- The dhcpd process might crash in a Junos fusion environment. [PR1478375](#)

- TFTP installation from loader prompt might not succeed on the EX Series devices. [PR1480348](#)
- In EVPN-VXLAN scenario, RP request packets for unknown host might get dropped in remote PE. [PR1480776](#)

Infrastructure

- On EX4600 and EX4300 Virtual Chassis, the queue buffers might get stuck. As a result, error messages related to soft reset of port is seen. [PR1462106](#)
- The traffic dropped on EX4300-48MP device acting as a leaf in Layer 2 IP fabric EVPN VXLAN environment. [PR1463318](#)
- In an EX2300 Virtual Chassis scenario, continuous **dcpfe error** messages and **eventd** process hog might be seen. [PR1474808](#)

Interfaces and Chassis

- The traffic might be forwarded to incorrect interfaces in an MC-LAG scenario. [PR1465077](#)
- Executing commit might become unresponsive due to a stuck device control process. [PR1470622](#)

Junos fusion for enterprise

- Inconsistency in the loop detect feature on newly provisioned SDs. [PR1460209](#)

Junos fusion Satellite Software

- The dpd crash might be observed on satellite devices in Junos fusion for enterprise. [PR1460607](#)

Layer 2 Features

- Physical layer and MAC/ARP learning might not work for copper base SFP-T on EX4600 switches. [PR1437577](#)
- The LLDP function might fail when a Juniper device connects to a non-Juniper device. [PR1462171](#)
- After rebooting, the FXPC core file might be seen when committing the configuration. [PR1467763](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)

Platform and Infrastructure

- The laser Tx might be enabled while the interface is disabled. [PR1445626](#)
- NSSU causes traffic loss again after the backup to master transitions. [PR1448607](#)
- On certain MPCs, cm errors need to be reclassified. [PR1449427](#)
- The IRB traffic might get dropped after mastership switchover. [PR1453025](#)
- The OSPF neighbor might go down when mDNS/PTP traffic is received at a rate higher than 1400 pps. [PR1459210](#)
- ERP might not revert to IDLE state after reload or reboot of multiple switches. [PR1461434](#)

- On EX4300 Virtual Chassis, traffic loss may be observed longer than 20 seconds when performing NSSU. [PR1461983](#)
- IGMP reports are dropped with mixed enterprise or 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 EX4300 switches, input firewall filter attached to isolated or community VLANs not matching dot1p bits on VLAN header. [PR1478240](#)

Routing Protocols

- Host destined packets with filter log action might not reach to the Routing Engine if log/syslog is enabled. [PR1379718](#)
- EX9208: BGP v4/v6 convergence and RIB install/delete time degraded in Junos OS Release 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The traffic with destination UDP port 520 (RIP) or 521 (RIPng) drops on EX4600 switches. [PR1429543](#)

User Interface and Configuration

- Problem with access to J-Web after update from Junos OS Release 18.2R2 to 18.2R3. [PR1454150](#)

Resolved Issues: 19.1R2

EVPN

- The device might proxy the ARP probe packets in an EVPN environment. [PR1427109](#)
- Configuring ESI on a single-homed 25-Gigabit Ethernet port might not work. [PR1438227](#)
- ARP request/NS might be sent back to the local segment by DF router. [PR1459830](#)

General Routing

- Certain EX Series devices are vulnerable to 'Etherleak' memory disclosure in Ethernet padding data. [PR1063645](#)
- Transit OSPF traffic over Q-in-Q tunneling might be dropped if a firewall filter is applied to Lo0. interface [PR1355111](#)
- l2ald process might crash and generate a core file on EX2300 Virtual Chassis when a trunk port is converted to 802.1x access port with tagged traffic flowing. [PR1362587](#)
- IPv6 router advertisement (RA) messages might increase internal kernel memory usage. [PR1369638](#)
- RIPv2 update packets might not send when IGMP-snooping is enabled. [PR1375332](#)
- Interface flapping on an EX Series Virtual Chassis might cause high CPU utilization and multicast traffic delay. [PR1393405](#)

- The fxpc core file might be generated if a scaled number of filter-based forwarding (FBF) filters are configured. [PR1398256](#)
- EX3400 might not learn 30,000 MAC addresses when it sends MAC learning traffic. [PR1399575](#)
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface. [PR1403528](#)
- MAC address movement might not happen in Flexible Ethernet Services mode when the families **inet** or **inet6** and **vlan-bridge** are configured on the same physical interface. [PR1408230](#)
- On EX2300-24P, the error message **dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family** is seen. [PR1410717](#)
- EX Series PEM alarm for backup FPC will remain on master FPC although the backup FPC is detached from the Virtual Chassis. [PR1412429](#)
- On EX4300-48MP switches, the chassis status LED is lit yellow instead of amber. [PR1413194](#)
- The upgrade of the PoE firmware might fail on EX3400. [PR1413802](#)
- VXLAN encapsulation next hop (VENH) does not get installed during BGP flapping or when routing is restarted. [PR1415450](#)
- On EX3400, the **show chassis environment** command repeats **OK** and **Failed** at short intervals. [PR1417839](#)
- The EX3400 Virtual Chassis status might be unstable during the bootup of Virtual Chassis or after the Virtual Chassis port flapping. [PR1418490](#)
- Traffic drop might be observed when a transit static LSP is configured on EX4650 and QFX5120 platforms. [PR1420370](#)
- Virtual Chassis might become unstable and FXPC might generate a core file when there are a lot of configured filter entries. [PR1422132](#)
- The interface on failed member FPC of EX2300 or EX3400 Virtual Chassis might stay up for 120 seconds. [PR1422507](#)
- Ensure Phone-home works in factory default configuration. [PR1423015](#)
- IPv6 multicast traffic received on one Virtual Chassis member might be dropped when exiting another Virtual Chassis member if MLD snooping is enabled. [PR1423310](#)
- On EX3400 auto-negotiation status shows incomplete on ge-0/2/0 using SFP-SX. [PR1423469](#)
- Multicast traffic might be silently dropped on an ingress port with **igmp-snooping** enabled. [PR1423556](#)
- MACsec connection on EX4600 platforms might not come back up after interface is disconnected and then reconnected. [PR1423597](#)
- On MX204 optics **SFP-1GE-FE-E-T**, I2C read errors are seen when an SFP-T transceiver is inserted into a disabled state port. [PR1423858](#)

- EX2300 or EX4300 platforms might fail to get an image/configuration from a phone-home server. [PR1424321](#)
- MAC overlaps between different switches. [PR1425123](#)
- The jdhcpd process might consume 100 percent CPU and crash if **dhcp-security** is configured. [PR1425206](#)
- VC split after network topology is changed. [PR1427075](#)
- The fxpc or Packet Forwarding Engine might crash on EX2300 and EX3400 platforms. [PR1427391](#)
- Rebooting or halting a Virtual Chassis member might cause 30 seconds down on the RTG link. [PR1427500](#)
- On EX2300-24P, the l2ald core files are generated after the removal and re-addition of multiple supplicant mode with PVLAN on interface. [PR1428469](#)
- Verification of ND inspection with a dynamically bound client, moved to a different VLAN on the same port, is failing. [PR1428769](#)
- EX4300-48MP switch cannot learn MAC addresses through some access ports that are directly connected to a host when auto-negotiation is used. [PR1430109](#)
- Incorrect model information while polling through SNMP from Virtual Chassis. [PR1431135](#)
- Packet drop might be seen if native VLAN is configured along with flexible VLAN tagging. [PR1434646](#)
- The mc-ae interface may get stuck in waiting state in a dual mc-ae scenario. [PR1435874](#)
- i40e NVM upgrade support for EX9200 platform. [PR1436223](#)
- LED turns on even after the VC members are powered off. [PR1438252](#)
- The DHCP snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it. [PR1438351](#)
- The dot1x might not work when captive-port is also configured on the interface on a backup or non-master FPC. [PR1439200](#)
- DHCPv6 relay binding is not up while verifying the DHCP snooping along with DHCPv6 relay. [PR1439844](#)
- EX4600 Virtual Chassis does not come up after replacing Virtual Chassis port from fiber connection to DAC cable. [PR1440062](#)
- CPU might hang or interface might get stuck on a particular 100G port on EX Series switches. [PR1440526](#)
- MAC addresses learned on RTG might not be aged out after a Virtual Chassis member is rebooted. [PR1440574](#)
- Clients in an isolated VLAN might not get IP addresses after completing authentication when both **dhcp-security** and **dot1x** are configured. [PR1442078](#)
- EX3400 FAN alarm (Fan X not spinning) appears and disappears repeatedly after removal of the fantray (absent). [PR1442134](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- Nondesignated port is not moving to backup port role. [PR1443489](#)

- `/var/host/motd` does not exist message is flooded every 5 seconds in chassisd logs. [PR1444903](#)
- Packets drop might be seen after removing and reinserting the SFP of the 40G uplink module ports. [PR1456039](#)
- On EX4300-MP, the following log is generated continuously: `rpd [6550]: task_connect: task AGENTD I/O.128.0.0.1+9500 addr 128.0.0.1+9500: Connection refused.` [PR1445618](#)
- Major alarm log messages for temperature conditions for EX4600 at 56 degrees Celsius. [PR1446363](#)
- Traffic might be dropped when a firewall filter rule uses 'then vlan' as the action in a VC scenario. [PR1446844](#)
- The phone-home feature might fail on EX4300 switches because sysctl cannot read the device serial number. [PR1447291](#)
- Added the **on-disk-failure** CLI configuration on the EX3400 switches. [PR1447853](#)
- Unicast ARP requests are not replied to with the **no-arp-trap** option. [PR1448071](#)
- On EX3400, IPv6 routes received through BGP do not show the correct age time. [PR1449305](#)
- DHCP snooping static binding is not effective after the configuration is deleted and added back. [PR1451688](#)
- Configuration change in **VLAN all** option might affect the per-VLAN configuration. [PR1453505](#)
- Version compare in PHC might fail, making PHC to download the same image. [PR1453535](#)
- Timeout connecting to peer 'database-replication'. [PR1457284](#)
- SNMP trap messages are shown up after upgrade even when the temperature are within the system thresholds. [PR1457456](#)

Infrastructure

- The Packet Forwarding Engine is flooded with messages `// pkt rx on ifd NULL unit 0.` [PR1381151](#)
- The traffic to the NLB server may not be forwarded if the NLB cluster works on multicast mode. [PR1411549](#)
- Some Junos OS releases might not be installed successfully on EX2300-C platform. [PR1414688](#)
- The operations on console might not work if the **system ports console log-out-on-disconnect** statement is configured. [PR1433224](#)
- EX3400 might reboot suddenly generating VM core files. [PR1456668](#)
- The traffic dropped on EX4300-48MP device acting as a leaf in Layer 2 IP fabric EVPN VXLAN environment. [PR1463318](#)

Interfaces and Chassis

- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces a misleading error message. [PR1402606](#)
- The logical interfaces in EVPN routing instances might flap after committing configurations. [PR1425339](#)
- On EX9200 line of Ethernet switches, the unexpected commit error **duplicate VLAN-ID** occurs. [PR1430966](#)
- VRRP-v6 state is flapping with init and idle states after configuring **vlan-tagging**. [PR1445370](#)
- The traffic might be forwarded to wrong interfaces in MC-LAG scenario. [PR1465077](#)

J-Web

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

Junos fusion for enterprise

- PoE over LLDP negotiation is not supported in Junos fusion for enterprise setup. [PR1366106](#)
- Traffic might be silently dropped and discarded in a Junos fusion for enterprise with dual-AD. [PR1417139](#)
- 1-Gigabit SFP transceiver in a 10-Gigabit Ethernet upstream port on EX3400/EX4300 satellite device may not come up. [PR1420343](#)
- Loop-detect feature not working in Junos fusion for enterprise. [PR1426757](#)

Layer 2 Ethernet Services

- BFD might flap when some of underlay ECMP interfaces are disabled in the leaf nodes. [PR1416941](#)
- The malfunction of the core isolation feature in EVPN-VXLAN scenarios causes traffic to be silently dropped and discarded. [PR1417729](#)
- The DHCP DECLINE packets are not forwarded to the DHCP server when forward-only is set within dhcp-reply. [PR1429456](#)
- DHCP request might get dropped in a DHCP relay scenario. [PR1435039](#)
- On EX9200, DHCP-Relay is stripping the **GIADDR** field in messages toward the DHCP clients. [PR1443516](#)

Layer 2 Features

- ERPS nodes do not converge to idle state after failure recovery or reboot. [PR1431262](#)
- The MAC ARP learning might not work for copper base SFP-T on QFX5100, QFX5110, and EX4600 line of switches. [PR1437577](#)

Network Management and Monitoring

- Overtemperature trap is not sent out even though there is a Temperature Hot alarm. [PR1412161](#)

Platform and Infrastructure

- On EX4300, OAM LFM might not work on **extended-vlan-bridge** interface with **native vlan** configured. [PR1399864](#)
- On EX9251, EX9253, and EX9208, DDoS violation for LLDP, MVRP, provider MVRP, and 802.1x is incorrectly reported as LACP DDoS violation. [PR1409626](#)
- Untagged traffic is single-tagged in Q-in-Q scenario on EX4300 platforms. [PR1413700](#)
- In EX4300 few ports might remain in dot1x connecting state and fail to transition to authenticated state. [PR1417270](#)
- Overtemperature SNMP trap is generated incorrectly for LC (EX4300-48P) based on master Routing Engine (EX4300-48MP) temperature threshold value. [PR1419300](#)
- EX4300: Runt counter never incremented. [PR1419724](#)
- EX4300 does not send Fragmentation needed message when MTU is exceeded with DF bit set. [PR1419893](#)
- The pfex process might crash and generate core files when the SFP transceiver is reinserted. [PR1421257](#)
- Traffic loss when one of the logical interfaces on a LAG is deactivated or deleted. [PR1422920](#)
- The auditd crashes when accounting RADIUS server is not reachable. [PR1424030](#)
- SNMP (ifHighSpeed) value is not displayed properly only for VCP interfaces, and it appears as zero. [PR1425167](#)
- Interface flapping scenario might lead to ECMP next-hop install failure on EX4300 switches. [PR1426760](#)
- IPv6 traffic might be dropped when static /64 IPv6 routes are configured. [PR1427866](#)
- VIP might not forward the traffic if VRRP is configured on an aggregated Ethernet interface. [PR1428124](#)
- EX4300 does not drop FCS frames with CRC error on 10-Gigabit Ethernet interfaces. [PR1429865](#)
- Unicast ARP requests are not replied to with the "no-arp-trap" option. [PR1429964](#)
- EX4300 enables the soft error recovery feature on the Packet Forwarding Engine, which can automatically detect the Packet Forwarding Engine parity error and recover by itself. [PR1430079](#)
- The ERPS failover does not work as expected on an EX4300 device. [PR1432397](#)
- The **/var/db/scripts** directory might be deleted after executing **request system zeroize**. [PR1436773](#)
- PoE might not work after upgrading the PoE firmware on EX4300 platforms. [PR1446915](#)
- REST API process is non-responsive when a number of requests arrive with a high rate. [PR1449987](#)
- ERP might not revert back to IDLE state after reload/reboot of multiple switches [PR1461434](#)

Routing Protocols

- Host-directed packets with filter log action might not reach the Routing Engine if log/syslog is enabled. [PR1379718](#)
- ICMPv6 RA packets generated by the Routing Engine might be dropped on the backup member of the Virtual Chassis if **igmp-snooping** is configured. [PR1413543](#)
- The EX Series switch may not install all IRB MAC addresses during initialization. [PR1416025](#)
- Sometimes, IGMP snooping may not work. Workaround is to restart the multicast-snooping process. [PR1420921](#)
- Error message **RPD_DYN_CFG_GET_PROF_NAME_FAILED: Get profile name for session XXX failed: -7**, might be seen in syslog after restarting routing daemon. [PR1439514](#)
- The bandwidth value of the DDoS-protection might cause packet loss after the device is rebooted. [PR1440847](#)
- Traffic might be dropped after the QinQ enabled interface is flapped or a change is made to the vlan-id-list. [PR1441402](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- Junos OS BFD sessions with authentication flap after a certain time. [PR1448649](#)
- Loopback address exported into other VRF instance might not work on ACX Series, EX Series, and QFX Series platforms. [PR1449410](#)

Subscriber Access Management

- On EX4300, /var showing full; the /var/log/dfcd_enc file grows in size. [PR1425000](#)

Virtual Chassis

- Current MAC address might change when you delete one of the multiple Layer 3 interfaces. [PR1449206](#)

VPNs

- MVPN using PIM dense mode does not prune the OIF when PIM prune is received. [PR1425876](#)

Resolved Issues: 19.1R1

EVPN

- A few minutes traffic loss might be observed during recovery from link failure. [PR1396597](#)

General Routing

- FPM board status is missing in SNMP MIB walk result. [PR1364246](#)
- The **OAM Ethernet connectivity-fault-management** configured on aggregated Ethernet interfaces is not supported; and no commit error is seen. [PR1367588](#)

- ARP request packets might be sent out with 802.1Q VLAN tag. [PR1379138](#)
- The 40-Gigabit Ethernet and 100- Gigabit Ethernet uplink port number options show incorrect number ranges. [PR1382578](#)
- Commit error is observed for the first time while loading the **mini-PDT base** configurations. [PR1383469](#)
- On the EX4650 switch, occasionally two of the channelized 25-Gigabit Ethernet ports that are using 4x25G breakout cable will not come up after Junos OS reboots. [PR1384898](#)
- ARP and ethernet-table entry in pointing to an aggregated Ethernet interface whose state is down. [PR1385199](#)
- On EX4300-48MP, the **session-option** stanza under the **[access profile]** hierarchy for EX Series and QFX Series platforms is not applicable. [PR1385229](#)
- On EX9200 platforms, the warning message **prefer-status-control-active** with **status-control standby** might be seen whenever you commit an operation. [PR1386479](#)
- On EX2300 with stacked VLAN, **flexible-vlan-tagging** is unable to obtain DHCP IP for IRB after a reboot/power-cycle. [PR1387039](#)
- On EX3400 Virtual Chassis, **Error tvp_status_led_set** and **Error:tvp_optics_diag_eeprom_read** syslog errors are seen. [PR1389407](#)
- MAC learning might stop working on some LAG interfaces. [PR1389411](#)
- On EX4300-48MP, need to remove messages **Recommend power cycle the device to complete the upgrade** and **Please power cycle the device to complete the upgrade** after ssd firmware upgrade. [PR1389543](#)
- "Input rate pps" is not increased on EX2300-MP uplink ports if the packet is a pure Layer 2 packet like non-etherII or non-EtherSnap. [PR1389908](#)
- The smid core file is generated during sanity script execution on QFX5100 and EX4300 switches. [PR1391909](#)
- PTP over Ethernet traffic might be dropped when IGMP and PTP TC are configured together. [PR1395186](#)
- DOT1XD core file is generated at **pnac_bd_create pnac_bdm_handler knl_async_receive_and_process**. [PR1395384](#)
- On EX2300, MAC table is not populated after interface-mode change. [PR1396422](#)
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. [PR1398398](#)
- After upgrading Junos OS Release 15.1X53 to Junos OS Release 18.2R1.9, the EX3400 cannot learn 30,000 MAC addresses. [PR1399575](#)
- The FBF routing-instance instance-type "forwarding" is missed for EX Series (EX3400). [PR1400163](#)
- MAC-limit with persistent MAC is not working after reboot. [PR1400507](#)
- The authd might crash when you issue the **show network-access requests pending** command during authd restart. [PR1401249](#)

- On EX4300-48mp, packets are dropped after configuring traffic filter and routing instance. [PR1407424](#)
- The l2cpd might crash if the VSTP **traceoptions** and VSTP VLAN **all** commands are configured. [PR1407469](#)
- The chassisd output power budget is received continuously for 5 seconds without any alarm after upgrading to Junos OS Release 18.1R3. [PR1414267](#)
- VXLAN Encapsulation nexthop (VENH) does not get installed during BGP flap or restart routing. [PR1415450](#)

Infrastructure

- IfSpeed and IfHighSpeed are erroneously reported as zero on EX2300. [PR1326902](#)

Junos fusion for enterprise

- An error **peer_daemon: bad daemon: scpd** is seen on EX9251 switch running Junos OS Releases 18.1R1 and 18.1R2. [PR1369646](#)
- Cannot login to SD cluster though it is recognized by AD properly. [PR1395570](#)
- The l2ald process might generate a core file when persistent MAC addresses are cleared from the switching table. [PR1409403](#)
- Extended ports do not adjust MTU in Junos fusion for enterprise on VOIP-enabled ports. [PR1411179](#)

Layer 2 Features

- RTG MAC refresh packets are sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. [PR1389695](#)

Layer 3 Features

- The l2ald might crash when the **clear ethernet-switching table persistent-learning** command is issued. [PR1381739](#)

Platform and Infrastructure

- Ping does not go through device after WTR timer expires in ERPS scenario. [PR1132770](#)
- EX4300 upgrade fails during validation of slax script. [PR1376750](#)
- ECMP route installation failure with log messages such as unilist install failure might be observed on EX4300 device. [PR1376804](#)
- Packet drops on interface if the statement **gigether-options loopback** is configured. [PR1380746](#)
- Traffic loss is seen in Layer 2 VPN with GRE tunnel. [PR1381740](#)
- EX4300 device chooses incorrect bridge-id as RSTP bridge-id. [PR1383356](#)
- On EX4300-48MP switch mixed Virtual Chassis, PoE interface maximum power configuration on member EX4300 gives an error if configured more than 30. [PR1383717](#)
- Unicast DHCP request get misforwarded to backup RTG link on EX4300-VC. [PR1388211](#)

- ICMPV6 packets are not classified with static or multifield forwarding-class mapping. [PR1388324](#)
- Layer 3 IP route might be deleted after a Layer 2 next-hop change is seen. [PR1389688](#)
- Continuous log messages get printed in EX4300: **17.4 / MCSNOOPD ICCP Context./var/run/iccpd_control addr /var/run/iccpd_control: Connection refused.** [PR1391942](#)
- On EX4300 switches, tcpdump shows that the kernel is sending out the ARP response on receiving the ARP request, but that the response does not get on the wire. [PR1405168](#)
- The policer might not work when it is applied through the dynamic filter. [PR1410973](#)

Routing Protocols

- The PPM mode for BFD session in EX4300 is centralized and not distributed by default. [PR1361800](#)
- On EX4300-48MP, stale VLAN entries are seen after continuous script is run involving split, merge, and reboot. [PR1363739](#)
- On EX4650 switches, the command output for the **show pfe route summary hw** statement shows different scale values for the IPv4 and IPv6 LPM routes rather than the supported scale. [PR1366579](#)
- Host-destined packets with filter log action might reach the Routing Engine. [PR1379718](#)
- EX4300 might drop the incoming IS-IS hello packets when IGMP or MLD snooping is configured. [PR1400838](#)

SEE ALSO

What's New 28
What's Changed 35
Known Limitations 39
Open Issues 41
Documentation Updates 58
Migration, Upgrade, and Downgrade Instructions 59

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for the EX Series switches.

SEE ALSO

What's New	28
What's Changed	35
Known Limitations	39
Open Issues	41
Resolved Issues	45
Migration, Upgrade, and Downgrade Instructions	59

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 59

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 | 28](#)
- [What's Changed | 35](#)
- [Known Limitations | 39](#)
- [Open Issues | 41](#)
- [Resolved Issues | 45](#)
- [Documentation Updates | 58](#)

Junos OS Release Notes for Junos fusion for enterprise

IN THIS SECTION

- [What's New | 61](#)
- [What's Changed | 62](#)
- [Known Limitations | 62](#)
- [Open Issues | 63](#)
- [Resolved Issues | 64](#)
- [Documentation Updates | 66](#)
- [Migration, Upgrade, and Downgrade Instructions | 66](#)

These release notes accompany Junos OS Release 19.1R3 for Junos fusion for enterprise. Junos fusion for 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 for enterprise, including which Juniper Networks devices can function as satellite devices, see [Understanding Junos fusion for 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

IN THIS SECTION

- [Release 19.1R3 New and Changed Features](#) | 61
- [Release 19.1R2 New and Changed Features](#) | 61
- [Release 19.1R1 New and Changed Features](#) | 61

Learn about new features introduced in the Junos OS main and maintenance releases for Junos fusion for enterprise.

NOTE: For more information about the Junos fusion for enterprise features, see the [Junos fusion for enterprise User Guide](#).

Release 19.1R3 New and Changed Features

There are no new or changed features in Junos OS Release 19.1R3 for Junos fusion for enterprise.

Release 19.1R2 New and Changed Features

There are no new or changed features in Junos OS Release 19.1R2 for Junos fusion for enterprise.

Release 19.1R1 New and Changed Features

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for SFTP global disablement (Junos fusion for enterprise)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the **sftp-server** statement at the **[edit system services ssh]** hierarchy level. In releases before Junos OS Release 19.1R1, the incoming SFTP connections are globally enabled by default.

[See [Configuring sftp-server](#).]

SEE ALSO

What's Changed	 62
Known Limitations	 62
Open Issues	 63
Resolved Issues	 64
Documentation Updates	 66
Migration, Upgrade, and Downgrade Instructions	 66

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 19.1R3 for Junos fusion for enterprise.

SEE ALSO

What's New	 61
Known Limitations	 62
Open Issues	 63
Resolved Issues	 64
Documentation Updates	 66
Migration, Upgrade, and Downgrade Instructions	 66

Known Limitations

IN THIS SECTION

- [Junos fusion for enterprise](#) | 63

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R1 for Junos fusion Enterprise.

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

Junos fusion for enterprise

- When using an EX9200 switch as the aggregation device (AD), you might see the error "jnh_dot1br_ktree_entry_create(1098): Entry Already Exists" continuously. This is caused by EX2300 satellite devices (SDs) having duplicate ECID values for the cluster/extended ports across members of the same cluster devices. [PR1408947](#)

SEE ALSO

What's New	 61
What's Changed	 62
Open Issues	 62
Resolved Issues	 64
Documentation Updates	 66
Migration, Upgrade, and Downgrade Instructions	 66

Open Issues

IN THIS SECTION

- [Junos fusion for enterprise](#) | [63](#)

This section lists the known issues in hardware and software in Junos OS Release 19.1R3 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

- The sdpd process generates a core file at **vfpc_all_eports_deletion_complete** **vfpc_dampen_fpc_timer_expiry** on the following triggers:
 - Deleting and adding the same satellite device with some attribute change in a single commit.
 - Interchange of slot ID between two satellite devices in a single commit.

- As a workaround, instead of making these satellite device configuration changes in a single commit, break it into multiple commits. [PR1454335](#)
- On a Junos fusion system, intermediate traffic drop might be seen between the aggregation and satellite device when SFlow is enabled on the ingress interface. When SFlow is enabled, the original packet is corrupted for those packets which hit the SFlow filter. This is because the packets egressing the aggregation device are short of 4 bytes of FCS 2 bytes of data. Normal data packets are 128 bytes (4 bytes for FCS, 14 bytes for Ethernet header, 20 bytes for IP header and 90 bytes for data). The corrupted packets are 122 bytes (14 bytes for Ethernet header, 20 bytes for IP header, and 88 bytes for data). [PR1450373](#)
 - In a Junos Fusion scenario with the EX4300 switch acting as the satellite device, the temperature sensor alarms and logs might be seen due to an incorrect reading about 2-3 degrees lower than expected. There is no CLI command to clear the alarm and logs. [PR1466324](#)

SEE ALSO

What's New 61
What's Changed 62
Known Limitations 62
Resolved Issues 64
Documentation Updates 66
Migration, Upgrade, and Downgrade Instructions 66

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 19.1R3 | 65](#)
- [Resolved Issues: Release 19.1R2 | 65](#)
- [Resolved Issues: Release 19.1R1 | 65](#)

Learn which issues were resolved in the Junos OS main and maintenance releases 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.

Resolved Issues: Release 19.1R3

- There is an inconsistency in the loop detect feature on newly provisioned satellite devices. [PR1460209](#)
- The dpd process might generate a core file on satellite devices in Junos fusion for enterprise. [PR1460607](#)

Resolved Issues: Release 19.1R2

- Traffic might be dropped in Junos fusion for enterprise scenario with dual aggregation devices. [PR1417139](#)
- The 1-gigabit SFP transceiver in the 10-Gigabit Ethernet upstream port on EX3400 and EX4300 satellite devices might not come up. [PR1420343](#)
- The loop-detect feature does not work in Junos fusion for enterprise. [PR1426757](#)
- Reachability issue of the host connected to the satellite device might be affected in Junos fusion for enterprise environment with EX9200 series devices as aggregation devices. [PR1447873](#)

Resolved Issues: Release 19.1R1

- PoE-over-LLDP negotiation is not supported on a Junos fusion for enterprise setup. [PR1366106](#)
- An error `peer_daemon: bad daemon: scpd` is seen on an EX9251 switch running Junos OS Releases 18.1R1 and 18.1R2. [PR1369646](#)
- Cannot log in to satellite device cluster although it is recognized by the aggregation device properly. [PR1395570](#)
- The l2ald process might generate a core file when persistent MAC addresses are cleared from the switching table. [PR1409403](#)
- VoIP-enabled extended ports (on satellite devices) do not adjust MTU in Junos fusion for enterprise. [PR1411179](#)

SEE ALSO

[What's New | 61](#)

[What's Changed | 62](#)

[Known Limitations | 62](#)

[Open Issues | 63](#)

[Documentation Updates | 66](#)

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 for documentation for Junos fusion for enterprise.

SEE ALSO

What's New	 61
What's Changed	 62
Known Limitations	 62
Open Issues	 63
Resolved Issues	 64
Migration, Upgrade, and Downgrade Instructions	 66

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | 67
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | 68
- [Preparing the Switch for Satellite Device Conversion](#) | 69
- [Converting a Satellite Device to a Standalone Switch](#) | 70
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 70
- [Downgrading from Junos OS](#) | 71

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 for 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, where *n* is the spin number.

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-18.3B1.n.tgz
```

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

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-18.3B1.n-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 for enterprise. See [Configuring or Expanding a Junos fusion for enterprise](#).

For satellite device hardware and software requirements, see [Understanding Junos fusion for 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 for 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 for 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 for enterprise from Junos OS Release 18.3R1, follow the procedure for upgrading, but replace the 18.3 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

[What's New | 61](#)

[What's Changed | 62](#)

[Known Limitations | 62](#)

[Open Issues | 63](#)

[Resolved Issues | 64](#)

[Documentation Updates | 66](#)

Junos OS Release Notes for Junos fusion for provider edge

IN THIS SECTION

- What's New | 72
- What's Changed | 75
- Known Limitations | 76
- Open Issues | 76
- Resolved Issues | 77
- Documentation Updates | 79
- Migration, Upgrade, and Downgrade Instructions | 79

These release notes accompany Junos OS Release 19.1R3 for the Junos fusion for 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

IN THIS SECTION

- What's New in Release 19.1R3 | 73
- What's New in Release 19.1R2 | 73
- What's New in Release 19.1R1-S4 | 74
- What's New in Release 19.1R1 | 74

Learn about new features introduced in the Junos OS main and maintenance releases for Junos fusion for provider edge.

What's New in Release 19.1R3

Hardware

- **Support for QFX5110 as a satellite device in a Junos fusion for provider edge on a GNF (MX480, MX960, MX2010, and MX2020)**—With Junos Node Slicing, you can create guest network functions (GNFs), partitions where an aggregation device can be configured. The aggregation device on a GNF supports a maximum of 10 satellite devices. In Junos OS Release 19.1R3, Junos OS supports QFX5110 switches as satellite devices in Junos fusion for provider edge on a GNF.

[See [Understanding Junos Fusion Provider Edge Software and Hardware Requirements](#) and [Junos Node Slicing Overview](#).]

What's New in Release 19.1R2

There are no new features in Junos OS Release 19.1R2 for Junos fusion for provider edge.

What's New in Release 19.1R1-S4

Hardware

- **Support for QFX5110 as a satellite device in a Junos fusion for provider edge on a GNF(MX480 and MX960)**—With Junos Node Slicing, you can create guest network functions (GNFs), partitions where an aggregation device can be configured. The aggregation device on a GNF supports a maximum of 10 satellite devices. Starting in Junos OS Release 19.1R1-S4, Junos OS supports QFX5110 switches as satellite devices in Junos fusion for provider edge on a GNF.

[See [Understanding Junos Fusion Provider Edge Software and Hardware Requirements](#) and [Junos Node Slicing Overview](#).]

What's New in Release 19.1R1

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for SFTP global disablement (Junos fusion for provider edge)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the statement **sftp-server** at the **[edit system services ssh]** hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections were globally enabled by default.

[See [Configuring sftp-server](#)]

SEE ALSO

What's Changed 75
Known Limitations 76
Open Issues 76
Resolved Issues 77
Documentation Updates 79
Migration, Upgrade, and Downgrade Instructions 79

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.1R3 | 75](#)
- [What's Changed in Release 19.1R2 | 75](#)
- [What's Changed in Release 19.1R1 | 75](#)

Learn about what changed in the Junos OS main and maintenance releases for Junos fusion for provider edge.

What's Changed in Release 19.1R3

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 19.1R3 for Junos fusion for provider edge.

What's Changed in Release 19.1R2

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 19.1R2 for Junos fusion for provider edge.

What's Changed in Release 19.1R1

Junos fusion

- Starting with Junos OS Release 19.1R1, configuration changes that lead to catastrophic operations for subscribers on a static VLAN demux interface will now fail and an error message displayed.

NOTE: Junos currently checks configuration changes for subscribers on static and dynamic VLAN interfaces and displays an error message for catastrophic operation changes.

SEE ALSO

[What's New | 72](#)

[Known Limitations | 76](#)

[Open Issues | 76](#)[Resolved Issues | 77](#)[Documentation Updates | 79](#)[Migration, Upgrade, and Downgrade Instructions | 79](#)

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 19.1R3 for Junos fusion for 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 | 72](#)[What's Changed | 75](#)[Open Issues | 76](#)[Resolved Issues | 77](#)[Documentation Updates | 79](#)[Migration, Upgrade, and Downgrade Instructions | 79](#)

Open Issues

IN THIS SECTION

- [Junos fusion for provider edge | 77](#)

Learn about open issues in this release for Junos fusion for provider edge.

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 provider edge

- In a Junos fusion scenario, if there are more than 12 cascade ports configured on a satellite device (SD), the satellite discovery and provisioning process (sdpd) might continuously crash after the configuration is committed. As a result, the SD cannot be managed from the aggregation device (AD). Traffic loss might not be observed right after the sdpd crash, and because the crash is continuous without interruption, the related FPCs on the AD device might reach 100 percent CPU utilization. [PR1437387](#)
- On an MX Series Junos fusion setup with an aggregated Ethernet interface configured with **link-speed**, the aggregated Ethernet interface might flap whenever a new logical interface is added to it and a commit is performed. The issue results in service flapping on the aggregated Ethernet interface. [PR1441869](#)

SEE ALSO

[What's New | 72](#)

[What's Changed | 75](#)

[Known Limitations | 76](#)

[Resolved Issues | 77](#)

[Documentation Updates | 79](#)

[Migration, Upgrade, and Downgrade Instructions | 79](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3 | 78](#)
- [Resolved Issues: 19.1R2 | 78](#)
- [Resolved Issues: 19.1R1 | 78](#)

Learn which issues were resolved in the Junos OS main and maintenance releases for Junos fusion for 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: 19.1R3

Junos fusion for provider edge

- There are no fixed issues in the Junos OS Release 19.1R3 for Junos fusion for provider edge.

Resolved Issues: 19.1R2

Junos fusion for provider edge

- auto-negotiation is not disabled in hardware after setting the **no-auto-negotiation** option in CLI. [PR1411852](#)
- Junos fusion: Incorrect power values for extended optical ports. [PR1412781](#)
- The sdpd process might continuously crash if there are more than 12 **cascade-ports** configured to a satellite device. [PR1437387](#)
- The ae interface might flap whenever a new IFL is added to it. [PR1441869](#)
- Deprecate Junos fusion Support. [PR1448245](#)

Resolved Issues: 19.1R1

Junos fusion for provider edge

- Broadcast, Unknown Unicast and Multicast(BUM) traffic might get dropped on peer fusion Aggregation Device when link between Satellite Device and local Aggregate Device goes down. [PR1384440](#)

Junos fusion Satellite Software

- Extended Port (EP) LAG might go down on the Satellite Devices (SDs) if the related Cascade Port (CP) links to an Aggregation Device (AD) goes down. [PR1397992](#)

SEE ALSO

[What's New | 72](#)

[What's Changed | 75](#)

[Known Limitations | 76](#)

[Open Issues | 76](#)

[Documentation Updates | 79](#)

[Migration, Upgrade, and Downgrade Instructions | 79](#)

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for Junos fusion for provider edge.

SEE ALSO

[What's New | 72](#)[What's Changed | 75](#)[Known Limitations | 76](#)[Open Issues | 76](#)[Resolved Issues | 77](#)[Migration, Upgrade, and Downgrade Instructions | 79](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 80](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 82](#)
- [Preparing the Switch for Satellite Device Conversion | 83](#)
- [Converting a Satellite Device to a Standalone Device | 84](#)
- [Upgrading an Aggregation Device | 86](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 87](#)
- [Downgrading from Junos OS Release 19.1 | 87](#)

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 19.1R3 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 for provider edge.

```
user@host> request system software add validate reboot  
source/jinstall64-19.1R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot  
source/jinstall-19.1R3.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 for provider edge.

```
user@host> request system software add validate reboot  
source/jinstall64-19.1R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot  
source/jinstall-19.1R3.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 19.1R3 **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 and 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 for 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 19.1R3, 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 19.1

To downgrade from Release 19.1 to another supported release, follow the procedure for upgrading, but replace the 19.1 **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 | 72](#)

[What's Changed | 75](#)

[Known Limitations | 76](#)

[Open Issues | 76](#)

[Resolved Issues | 77](#)

[Documentation Updates | 79](#)

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

IN THIS SECTION

- What's New | **88**
- What's Changed | **111**
- Known Limitations | **121**
- Open Issues | **125**
- Resolved Issues | **142**
- Documentation Updates | **195**
- Migration, Upgrade, and Downgrade Instructions | **195**

These release notes accompany Junos OS Release 19.1R3 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 Release 19.1R3-S2 | **89**
- What's New in Release 19.1R3 | **89**
- What's New in Release 19.1R2 | **90**
- What's New in Release 19.1R1-S1 | **90**
- What's New in Release 19.1R1 | **90**

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

What's New in Release 19.1R3-S2

High Availability (HA) and Resiliency

- **LACP inline support during ISSU for multivendor networks (MX104, MX240, MX480, MX960, MX10003)**—Starting with Junos OS Release 19.1R3-S2, ISSU supports LACP interoperability with other vendor devices for fast periodic interval sessions. LACP sessions in full scale scenarios with interoperability will no longer experience timeouts during ISSU.

Use the **set protocols lacp ppm inline** command to enable LACP inline support.

[See [Getting Started with Unified In-Service Software Upgrade](#)]

What's New in Release 19.1R3

There are no new features or enhancements to existing features for MX Series in Junos OS Release 19.1R3.

What's New in Release 19.1R2

Routing Protocols

- **MPC10 Inline BFD support (MX Series)**—Starting in Junos OS Release 19.1R2, MPC10 MPCs support inline BFD features, excluding micro BFD and BFD sessions with authentication.

[See [Understanding BFD for Static Routes](#).]

Subscriber Management and Services

- **CoA messages support Session-Timeout attribute (MX Series)**—Starting in Junos OS Release 19.1R2, 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](#).]

What's New in Release 19.1R1-S1

Interfaces and Chassis

- **MPC10 Distributed LACP Support in PPM AFT (MX Series)**—Starting in Junos OS Release 19.1R1-S1 and 19.1R2, MPC10E-15C-MRATE and MPC10E-10C-MRATE MPCs support distributed LACP in Periodic Packet Manager (ppman) Advanced Forwarding Toolkit (AFT).

What's New in Release 19.1R1

Authentication, Authorization, and Accounting (AAA) (RADIUS)

- **Support for SFTP global disablement (MX Series)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the statement `sftp-server` at the `[edit system services ssh]` hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections were globally enabled by default.

[See [Configuring sftp-server](#).]

Class of Service (CoS)

- **Support for CoS features (classifiers, rewrites, port queuing, L3 interfaces only) (MX Series)**—Starting with Junos OS Release 19.1R1, you can configure the standard CoS forwarding (classifiers, rewrites, port queuing, L3 interfaces only) on MPC10E-15C-MRATE line cards.

[See [Understanding Class of Service](#).]

- **Support for Real-time Transport Protocol (RTP) payload types 96 through 127 on inline video monitoring (MX Series)**—Starting with Junos OS 19.1R1, you can configure MX Series Routers for inline video

monitoring of uncompressed HD or 4K stream video (Payload Type 96 through 127). MDI functionality has been extended to video flows such as ST 2000-5 (RTP PT 98) and ST 2000-6 (RTP PT 99). These are non-MPEG video flows over IP/UDP/RTP and are constant bit rate flows. The operator would specify proper IP addresses and UDP ports so that nonvideo flows over RTP will not go through MDI processing.

[See [Understanding Inline Video Monitoring on MX Series Routers.](#)]

EVPN

- **Support for auto-derived route target on EVPN-MPLS (MX Series)**—Starting in Junos OS Release 19.1R1, Junos OS supports the automatic derivation of a route target on EVPN-MPLS. With this feature, the route target is automatically derived from the VLAN ID for EVPN type 2 and EVPN type 3 routes. The auto-derived route targets have higher precedence over manually configured RT in vrf-target, vrf-export policies, and vrf-import policies.

To enable auto-derived route target, include the **auto** statement at the **[edit routing-instances routing-instance-name protocols evpn vrf-target]** hierarchy level.

[See [Auto-derived Route Targets.](#)]

- **Support for proxy MAC addresses in an ARP request (MX Series)**—Starting in Junos OS Release 19.1R1, provider edge (PE) devices in an EVPN network that support ARP proxy can use a proxy MAC address in the ARP replies message to a host. When a PE device receives an ARP request or NDP request, it searches its MAC-IP address binding database for the requested IP address. If the device finds the MAC-IP address entry in its database, it responds to the request with the proxy MAC address. The proxy MAC address is derived from the IRB interface in an EVPN network with edge-routed bridging overlay and from the manually configured MAC address in a centrally routed bridging overlay. If the device does not find an entry, the PE device replaces the MAC and IP address from the CE device in the ARP request with the proxy MAC and IP address of the IRB interface. This allows for enhanced security (for example, L3 filtering) deployments on L3 gateway for both inter-VLAN and intra-VLAN traffic will be routed.

To enable this feature, configure the **proxy-mac [irb | proxy-mac-address]** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy or at the **[edit routing-instances routing-instance-name bridge-domains domain_name]** hierarchy.

[See [ARP and NDP Request with a proxy MAC address.](#)]

- **Support for asynchronous notification on EVPN-VPWS (MX Series)**—Starting in Junos OS Release 19.1R1, asynchronous-notification is supported on interfaces on EVPN-VPWS. You can enable the asynchronous notification command to send a loss of signal (LOS) alarm to the CE device when the circuit cross-connect link between a customer edge and provider edge device goes down. Asynchronous notification supports ethernet-ccc, ethernet-vpls, or vlan-ccc encapsulation.

To enable this feature, include the **asynchronous-notification** statement at the **[edit interfaces interface-name]** hierarchy level.

[See [Configuring Gigabit Ethernet Notification of Link Down Alarm.](#)]

Forwarding and Sampling

- **Support for tracking static RPM routes across multiple next hops (MX Series)**—Starting in Junos OS Release 19.1R1, you can use **rpm-tracking** to track up to 16 next hops for RPM-controlled static routes. This feature supports both IPv4 and IPv6 static rpm-tracked routes, and extends the single hop [rpm-tracking](#) introduced in Junos OS Release 18.4.

[See [show route rpm-tracking](#).]

- **Support for using IP addresses in an SR-TE LSP segment list (MX Series)**—Starting in Junos OS Release 19.1R1, you can use IP addresses (IPv4 or IPv6) for next hops in a segment routing traffic engineering (SR-TE) list of label-switched paths (LSPs). This work extends the support for traffic steering based on a segment routing policy that was introduced in Junos OS Release 17.4R1, wherein the controller can specify a segment routing policy consisting of multiple paths to steer labeled or IP traffic.

[See **auto-translate** in [segment-list](#) and **retry-timer** in [source-packet-routing](#).]

Hardware

- **New fixed-configuration Modular Port Concentrator (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.1R1, the MPC10E-15C-MRATE is a new Modular Port Concentrator (MPC) that is supported on MX240, MX480, and MX960 routers.

The MPC10E-15C-MRATE features the following:

- Line-rate throughput of up to 1.5 Tbps.
- Twelve QSFP28 ports—Port numbers 0/0 through 0/3, 1/0 through 1/3, and 2/0 through 2/3. The ports can be configured as 10-Gbps, 40-Gbps, or 100-Gbps Ethernet ports.
- Three QSFP56-DD ports—Port numbers 0/4, 1/4, and 2/4. The ports can be configured as 10-Gbps, 40-Gbps, 100-Gbps Ethernet ports.

[See [MX Series 5G Universal Routing Platform Interface Module Reference](#).]

Interfaces and Chassis

- **Support for MPC10E-15C-MRATE line card (MX240, MX480, and MX960)**—Starting with Junos OS Release 19.1R1, the MX240, MX480, and MX960 routers support the MPC10E-15C-MRATE line card. This fixed-port line card is capable of delivering a bandwidth of up to 1.5 Tbps per MPC slot. It supports three MICs (one per Packet Forwarding Engine), each of which can deliver a throughput of up to 500 Gbps. Each MIC comprises five ports that support 100 Gbps (the default), 40 Gbps, and 10 Gbps speeds through the use of QSFP28+ and QSFP+ optics. You enable 10 Gbps speed (four 10 Gbps channels) by using breakout cables.

NOTE:

- The MPC10E-15C-MRATE is powered on only if the MX Series router has an enhanced Switch Control Board (MX-SCBE3) installed.
- The MPC10E-15C-MRATE is supported only with the high-capacity AC and DC power entry modules (PEMs) and the high-capacity fan trays used in MX Series routers.
- The MPC10E-15C-MRATE is powered on only if the router operates in **enhanced-ip** or **enhanced-ethernet** mode.
- The MPC10E-15C-MRATE is not supported on the MX2000 and MX10000 lines of routers.

[See [MPC10E-15C-MRATE](#), [Understanding Interface Naming Conventions for MPC10E-15C-MRATE MPC](#), [MPC10E-15C-MRATE Rate-Selectability Overview](#), [Supported Active Physical Ports for Rate Selectability to Prevent Oversubscription on MPC10E-15C-MRATE](#), and [Configuring Rate Selectability on MPC10E-15C-MRATE to Enable Different Port Speeds](#).]

- **Chassis and power management for MPC10E-15C-MRATE (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.1R1, the MX240, MX480, and MX960 routers with the MPC10E-15C-MRATE line card support chassis management features, including field-replaceable unit (FRU) management, power budgeting and management, and environmental monitoring. The MPC10E-15C-MRATE line card supports configuration of ambient temperature (25°C, 40°C, and 55°C) and dynamic power management. The default ambient temperature value supported is 40°C. The MPC10E-15C-MRATE line card supports both hyper mode (the default mode) and normal mode.

NOTE:

- The MPC10E-15C-MRATE is powered on only if the MX Series router has an enhanced Switch Control Board (MX-SCBE3) installed.
- The MPC10E-15C-MRATE is powered on only if the router operates in **enhanced-ip** or **enhanced-ethernet** mode.
- The MPC10E-15C-MRATE will be powered on only when the MX Series router is installed with enhanced Fan Trays.
- The MPC10E-15C-MRATE will be supported only when the MX Series router is installed with Enhanced PEMs.
- On MX960 routers with enhanced Midplane on the slot 1, the MPC10E-15C-MRATE will not be powered on.

[See [Understanding How Configuring Ambient Temperature Helps Optimize Power Utilization](#) and [Understanding How Dynamic Power Management Enables Better Utilization of Power.](#)]

- **Packet Forwarding Engine power on and power off support for MPC10E-15C-MRATE (MX240, MX480, and MX960)**—Starting Junos OS Release 19.1R1, on MX240, MX480, and MX960 devices with MPC10E-15C support, you can power on or power off a Packet Forwarding Engine using the command **set chassis fpc slot-number pfe slot-number power (on | off)**.

The **show chassis fpc FPC Slot detail** displays the Packet Forwarding Engine power ON/OFF status and bandwidth for the individual PFEs in an MPC10E-15C-MRATE.

[See [show chassis fpc](#).]

- **Support for ETH-ED (MX Series)**—Starting with Junos OS Release 19.1R1, when a unified in-service software upgrade (ISSU) is about to start, the peer maintenance association end point (MEP) is notified to suppress the remote defect indication (RDI) and loss of adjacency alarms for a specified duration. To ensure that the notification is sent before the upgrade starts, you must configure the Ethernet expected defect (ETH-ED) function by including the **expected-defect** statement at the **[edit protocols oam ethernet connectivity-fault-management expected-defect]** hierarchy level.

[See [connectivity-fault-management](#).]

- **Improved performance of small packets (MX Series)** —Starting in Junos OS Release 19.1R1, the MPC7E-MRATE, MPC7E-10G, MPC8E, MPC9E, MX10003 MPC, MX204, and JNP10K-LC2101 line cards provide improved performance of small packets (with a minimum packet size of 64 bytes) in transmit direction. To enable this feature, reduce the number of active ports (at the PIC level) to the following maximum numbers:
 - Sixteen 10-Gbps ports
 - Four 40-Gbps ports

- Two 100-Gbps ports (when the line card is in 240-Gbps mode)
- Three 100-Gbps ports (when the line card is in 400-Gbps mode)

To configure the number of active ports, use the existing command **set chassis fpc slot pic slot number-of-ports number-of-active-ports**.

NOTE:

- The command does not change packet performance at the Packet Forwarding Engine level; it improves packet performance in transmit direction at the port level only.
- For an MX10003 MPC, in 40-Gbps and 10-Gbps PIC modes, if both the PICs are used, the number of ports cannot exceed six on either PIC. If only PIC 1 is used, you can set the number of ports to 12. For an MX204 MPC, in 10-Gbps PIC mode, if both the PICs are used, the sum of the interfaces created on the PICs cannot exceed 16. If only PIC 0 is used, you can set the number of ports to 4 (4 interfaces per port). If only PIC 1 is used, you can set the number of ports to 8 (1 interface per port).

[See [Understanding Rate Selectability](#).]

- **Support for inline LACP PDU transmission processing (MX Series routers with MPCs)**—Starting in Junos OS Release 19.1R1, MX Series routers with MPCs support inline LACP PDU transmission processing for periodic packet management (on the Packet Forwarding Engine). To enable the inline processing method instead of using the default LACP PDU transmission processing, issue the **set protocols lacp ppm inline** command.

[See [inline](#).]

IPsec

- **Distinguished name support in IPsec (MX Series)**—Starting with Junos OS Release 19.1R1, distinguished name (DN) support is added to the IKE identification (IKE ID) that is used for validation of VPN peer devices during IKE negotiation. The IKE ID received by an MX Series router from a remote peer can be an IPv4 or an IPv6 address, a hostname, a fully qualified domain name (FQDN), or a DN. The IKE ID sent by the remote peer needs to match what is expected by the MX Series router. Otherwise, IKE ID validation fails and the VPN is not established.

A distinguished name (DN) is a name used with digital certificates to uniquely identify a user. You can use a container keyword to specify the order of the fields in a distinguished name and their values must exactly match the configured distinguished name, or use a wildcard keyword to specify that the values of fields must match but the order of the fields does not matter.

[See [Understanding Junos VPN Site Secure](#).]

- **Support for IPsec and Group VPN services (MX2010 and MX2020)**—Starting in Junos OS Release 19.1R1, Junos OS supports IPsec and Group VPN services on MX2010 and MX2020 routers. Group VPNs eliminate the need for point-to-point VPN tunnels in a mesh architecture. They provide a set of

features that are necessary to secure unicast traffic over a private WAN that originates on or flows through a router.

[See [Group VPNv2 Overview](#).]

Junos Telemetry Interface

- **RSVP interface OpenConfig model support and self-ping logs on Junos telemetry interface (JTI) (MX960 and PTX10003)**—Starting in Junos OS Release 19.1R1, JTI sensor support is enhanced for RSVP interfaces to include delivery of more statistics. The level of support is equivalent to the output delivered when using the **show rsvp interface detail** operational mode command.

To configure the sensor for statistics to be issued to an outside collector, include the following path for gRPC streaming:

- **/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/interface-attributes/interfaces/interface/***

To provision the 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 JTI.

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

- **Delegated RPM statistics sensor support for Junos telemetry interface (JTI) (MX Series)**—Starting with Junos OS Release 19.1R1, for MX Series routers operating with MS-MIC and MS-MPC, a new sensor allows customers to monitor delegated RPM service statistics on the router and export them to outside collectors at configurable intervals encoded in Google Protocol Buffer (GPB) format.

Delegated RPM is a mode where RPM probe generation and measurement calculation are done by MS-MIC and MS-MPC cards. This hardware assistance allows a very high scale of concurrent RPM probes. JTI sensor support for other RPM modes was added in Junos OS Release 18.3R1.

You can use the resulting data from this sensor to improve network design and optimize traffic engineering. Data can also be used to detect problems in individual devices as well as in the overall network and the traffic carried by it.

Monitor delegated RPM service statistics by configuring the **/junos/services/spu/delegated-rpm/** sensor for the **sensor** configuration statement.

For exporting statistics, configure parameters at the **[edit services analytics]** hierarchy level.

[See [sensor \(Junos Telemetry Interface\)](#) and [Configuring a Junos Telemetry Interface Sensor \(CLI Procedure\)](#) and [Junos OS Release Notes for MX Series 5G Universal Routing Platform, 18.3R1](#).]

- **Export of data associated with the Junos kernel through Junos telemetry interface (JTI) (EX9200, EX9251, EX9253, MX Series, and PTX Series)**—Starting in Junos OS Release 19.1R1, you can export data associated with the Junos kernel through remote procedure calls (gRPC) and JTI. Kernel telemetry data includes information on Veriexec state, graceful Routing Engine switchover (GRES), in-service

software upgrade (ISSU), and Routing Engine ifstate. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

Junos kernel sensors introduced in Junos OS Release 19.1R1 support both periodical and ON_CHANGE streaming. The following Junos kernel resource paths support periodical streaming only:

- /junos/kernel-ifstate/dead-ifstates-cnt
- /junos/kernel-ifstate/alive-ifstates-cnt
- /junos/kernel-ifstate/delayed-unrefs-cnt
- /junos/kernel-ifstate/delayed-unrefs-max

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

- **SR-TE telemetry statistics and BINDING-SID routes support for uncolored SR-TE policies (MX Series)**—Starting in Release 19.1R1, Junos OS supports SR-TE telemetry statistics and BINDING-SID routes for uncolored SR-TE policies. Uncolored SR-TE LSP is characterized by the absence of **color** statement in its configuration.

Junos OS now allows collection of traffic statistics for both ingress IP traffic and transit MPLS traffic that take uncolored SR-TE paths. Also, you can install BINDING-SID labels even if the first hop of the segment list is a label. Prior to Junos OS 19.1R1 Release, the installation of BSID routes was not supported if the first hop of the segment list was a label, and a commit check was done.

The **show spring-traffic-engineering lsp** command is enhanced to provide the source by which the SRTE policy was provisioned. For example, Static, Path Computation Element Protocol. Also, the **show spring-traffic-engineering lsp detail** command is enhanced to provide information on the source of the tunnel configuration and statistics.

By default, traffic sensors and statistic collection are disabled for static SR-TE routes. To enable provisioning of Junos telemetry interface traffic sensors in Junos OS data plane to stream out traffic statistics on segment routing policies and their Binding-SID routes, use the existing **statistics** statement at the **[edit source-packet-routing telemetry]** hierarchy level.

Layer 3

- **Support for Layer 3 features on the MPC10E-15C-MRATE (MX Series)**—Starting in Junos OS Release 19.1R1, Junos OS supports the following Layer 3 features on the MPC10E-15C-MRATE.
 - BGP (multipath/IPv4 and IPv6 labelled unicast)
 - IPv4 (forwarding and options)
 - IPv6 (forwarding and route accounting)
 - Load balancing (ECMP and FRR)

Options supported: enhanced -hash-key family inet/inet6/mpls

- L2VPN, CCC, and Layer 2 circuit
- MPLS (push/pop/swap, LDP, RSVP-Aggregate, RSVP-TE Admin Groups, RSVP-TE, OAM - LSP/VPN ping, trace route, auto bandwidth, and MPLS-FRR link node protection.
Options supported: No Decrement Ttl, No Propagate Ttl, MTU-signaling Splitting-merging, Primary/Secondary, ICMP Tunneling, IPv6 Tunneling, LDP Tunneling, Optimize Timer, Explicit-Null, UHP and PHP support.
- OSPF (node-link-protection and node-link-degradation)
- Protocols (IS-IS, OSPF, OSPF V3 for V6, BGP + BGP-v6, BGP LU, BGP-LS, BGP optimal-route-reflection (ORR), BFD (Centralized), Micro BFD(centralized), ICMP and ICMPv6 error handling, and LLDP).
- Routing instance logical system VRF
- Tunnel (generic routing encapsulation (GRE), logical tunnel (LT), and virtual tunnel (VT))
[See [Tunnel Services Overview](#).]

Management

- **Tracing support for individual JET application files (MX Series)**—Previously you could configure trace options for all applications. Starting in Junos OS Release 19.1R1, you can also configure trace options for an individual application file. If you configure trace options both globally (all applications) and locally (by application file), the local configuration has the higher priority. You must commit global trace options and the daemonized application configurations at the same time for the global trace options for the daemonized application to take effect.
[See [application](#).]

MPLS

- **Flexible MPLS label stack depth (MX Series with MPC and MIC)**—Currently, Junos OS supports push of up to a maximum of five labels per component of the next-hop chain, even though the underlying device capability can be higher. Starting in Junos OS Release 19.1R1, the device capability of pushing more than 5 labels can be leveraged for features, such as, segment routing traffic-engineering (TE) LSPs and RSVP-TE pop-and-forward LSPs.

The number of labels that can be pushed for an MPLS next hop is the number of labels the device is capable of pushing, or the maximum-labels configured under **family mpls** of the outgoing interface, whichever is smaller.

[See [Configuring the Maximum Number of MPLS Labels, maximum-labels](#).]

- **Support for MPLS ping and traceroute for segment routing (MX Series)**—Starting in Junos OS Release 19.1R1, MPLS ping and traceroute are supported for segment routing (SR) for protocols IS-IS and OSPF over IPv4. This feature also supports ECMP traceroute for protocols IS-IS and OSPF.

On MX Series routers, MPLS ping and traceroute for segment routing is supported with enhanced-ip mode only. Segment routing with IS-IS tunnels are stitched to LDP tunnels. Ping and traceroute for segment routing over RSVP is supported.

In Junos OS Release 19.1R1, MPLS ping and traceroute for segment routing supports IPv4 IGP-Prefix segment FEC validation. FEC validation for IGP-Adjacency Segment ID is not supported.

[See [ping mpls segment routing isis](#), [ping mpls segment routing ospf](#), [traceroute mpls segment-routing ospf](#), [traceroute mpls segment-routing isis](#).]

- **Support for MPC10E-15C-MRATE (MX Series)**—Starting in Junos OS Release 19.1R1, a new MPC10E-15C-MRATE is introduced.

The following MPLS features are supported on MPC10E-15C-MRATE in 19.1R1:

- Static, RSVP and LDP LSPs
- LSP statistics
- LSP ping and traceroute
- LSP TTL commands: no-propagate-ttl, no-decrement-ttl
- L2Circuit and L2VPN with or without control word
- L3VPN with chain-composite-nexthop
- L3VPN with vrf-table-label
- MPLS link protection, node protection and FRR
- 6VPE

The following MPLS features are not supported on the MPC10E-15C-MRATE: in 19.1R1:

- VCCV BFD
- L2CKT/L2VPN interworking (iw interface)
- Translational cross-connect (TCC)
- Flow-aware transport label
- Entropy label

- **Enhancements to MPLS for LSP path selection (MX Series)**—Starting in Junos OS Release 19.1R1, the following enhancements to MPLS have been added for LSP path selection and optimization:

- Earlier, when LSP active paths were modified, the LSP path got cleared and gets resigaled immediately. From Junos OS Release 19.1R1 onward, if a secondary path is available, and then Junos OS selects the secondary path as active, clears and resignals the primary path after the expiry of the **optimize-hold-dead-delay** timer. When the primary LSP path is established, the **revert-timer** gets started. After the **revert-timer** expires, the primary LSP path becomes active.

If the primary LSP path is not active with **revert-timer** on and when there is a change to the primary LSP path, then the LSP path gets cleared and resigaled immediately. When the primary LSP path is established, the revert-timer gets restarted.

- Earlier if there was any Constrained Shortest Path First (CSPF) failure then the current LSP path becomes invalid because it did not match with the configured constraints. In this case, the current LSP

path gets cleared immediately. From Junos OS Release 19.1R1 onwards, if a secondary LSP path is available, then Junos OS selects the secondary LSP path as active and clears the primary path after the expiry of the **optimize-hold-dead-delay** timer.

- The CLI command **no-bypass-statistics-polling** added under the **[edit protocols mpls statistics]** hierarchy now provides information on bypass LSP statistics.
- A new CLI command **delay** has been introduced under the **[edit protocols mpls optimize-adaptive-teardown]** hierarchy and the value for delay is in the range of (3..65535 seconds). When the **adaptive-teardown** configuration is triggered, the **delay** CLI command further delays the tearing down of old optimized LSP paths based on the configured value.

[See [statistics \(Protocols MPLS\)](#), [optimize-adaptive-teardown](#).]

- **Control transport address used for targeted-LDP session (MX Series)**—Currently, only the router-ID or interface address is used as the LDP transport address. Starting in Junos OS Release 19.1R1, you can configure any other IP address as the transport address of targeted LDP sessions, session-groups, and interfaces. This new configuration is applicable only for configured LDP neighbors that have Layer 2 circuit, MPLS, and VPLS adjacencies.

This feature is beneficial when you have multiple loopback interface addresses, and different IGPs associated with LDP interfaces, and you can control the session established between targeted LDP neighbors with the configured transport address.

[See [Control Transport Address Used for Targeted-LDP Session](#).]

- **MPLS egress traffic statistics for label IS-IS routes at ingress device (MX Series with MPC and MIC)**—Currently, sensors are available for collecting segment routing statistics for MPLS transit traffic, which is MPLS-to-MPLS in nature. Starting in Junos OS Release 19.1R1, additional sensors are introduced to collect segment routing statistics for MPLS egress traffic at the ingress provider edge (PE) device, which is IP-to-MPLS in nature.

With this feature, you can enable sensors for label IS-IS segment routing egress traffic only, and stream the statistics to a gRPC client.

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

- **Policy-based multipath routes (MX Series)**—In segment routing networks with multiple protocols in the core, you can combine segment routing traffic-engineered (SR-TE) LDP routes and SR-TE IP routes to create a multipath route that is installed in the routing information base (also known as routing table). You can resolve BGP service routes over the multipath route through policy configuration and steer traffic differently for different prefixes.

[See [Policy-Based Multipath Routes Overview](#).]

- **Support of Layer2 and Layer3 VPN services over non-colored Segment Routing for Traffic Engineering (SR TE) (MX Series)**— Starting from Junos OS Release 19.1R1, you can use BGP-based Layer2 and Layer3 VPN services over non-colored Segment Routing for Traffic Engineering (SR TE). You can also use other features such as un-balanced ECMP (wecmp), and multi-level weighted ECMP (h-wecmp).

To use hierarchical multi-level weighted ECMP (h-wecmp), configure the following route resolution import-policy:

set policy-options policy-statement mpath then multipath-resolve

set routing-options resolution rib bgp.l3vpn.0 inet-import mpath

set routing-options resolution rib bgp.l2vpn.0 inet-import mpath

set routing-options resolution rib mpls.0 inet-import mpath

[See [Static Segment Routing Label Switched Path](#)]

- **Routing Engine-based S-BFD for segment-routing traffic engineering (MX Series)**—Starting in Junos OS Release 19.1R1, you can run Routing Engine-based seamless BFD (S-BFD) over non-colored and colored label-switched paths (LSPs) with first-hop label resolution and use S-BFD as a fast mechanism to detect path failures.

[See [Routing Engine-based S-BFD for Segment-Routing Traffic Engineering with First-Hop Label Resolution](#).]

- **Use of SID labels as first hop for resolving non-colored static segment routing LSPs (MX Series)**—Currently, for a static non-colored segment routing traffic-engineered LSP to be usable, the first hop of the segment list must be an IP address. Only the second to *n*th hop could be segment identifier (SID) labels. Starting in Junos OS Release 19.1R1, this requirement does not apply. You can now configure SID labels as the first hop in the segment list.

With this configuration, static non-colored segment routing LSPs are resolved using MPLS fast reroute (FRR) and weighted equal-cost multipath. Without this configuration, by default, the LSPs are resolved using IP address.

[See [Static Segment Routing Label Switched Path](#).]

- **Support of install statement for segment routing LSPs (MX Series)**—The **install destination-prefix** statement which is currently supported at the **[edit protocols mpls label-switched-path *lsp-name*]** and **[edit protocols mpls static-label-switched-path *lsp-name* ingress]** hierarchy levels is now also supported at the **[edit protocols source-packet-routing source-routing-path *lsp-name*]** hierarchy level for both colored and non-colored static segment routing label-switched paths (LSPs).

You can associate one or more prefixes with a segment routing LSP using the **install** statement. When the LSP is up, all the prefixes are installed as entries into the **inet.3** or **inet6.3** routing table.

[See [install \(Protocols MPLS\)](#).]

Multicast

- **Support for BGP-MVPN Inter-AS option B (MX Series and T Series)**—Starting in Junos OS Release 19.1R1, for improved security and scalability, Juniper supports Inter-AS option B for BGP multicast virtual private networks (MVPNs) and segmented provider tunnels. Only specific configurations are supported, so for example, static tunnels are not supported, nor are PIM any-source multicast (ASM) and PIM source-specific multicast (SSM) tunnels.

In the supported configuration, BGP-MVPN sites can span multiple autonomous system (AS) boundaries (that is, domains). Each AS can implement its own p-tunnel (they don't have to be the same). Per-VPN subinterfaces are not shared between ASBRs. Likewise, provider edge (PE) routers from one AS cannot be reached from another AS, and the AS topology of one site is not exposed to any others.

[See [inter-as \(Routing Instances\)](#) and [BGP-MVPN Inter-AS Option B Overview](#).]

- **Support for multicast forwarding on MPC10E-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.1R1, multicast forwarding is fully supported on MPC10E-MRATE line cards for MX Series routers.

Network Management and Monitoring

- **Error handling and resiliency support for MPC10E (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.1R1, the MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE) line cards support error handling and software resiliency. The MPC10E supports detecting errors, reporting them through alarms, and triggering resultant actions. Use the existing commands **show chassis errors active**, **show chassis errors active details**, and **show chassis fpc errors** to view more details of the errors. MPC10E-15C-MRATE also supports powering on or off Packet Forwarding Engine (pfe2), by using the command **set chassis fpc slot pfe slot power (on|off)**, in case of errors such as hardware components issues in Packet Forwarding Engine (pfe2).

[See [show chassis fpc errors](#) and [clear chassis fpc errors](#).]

- **sFlow performance improvements (MX Series)**—Starting in Junos OS Release 19.1R1, the following improvements have been added to the sFlow technology feature:
 - For MX Series, the maximum number of samples per second per line card is raised from 950 pps to 9500 pps. Junos OS also introduces an adaptive sampling fallback feature, which decreases the sampling load when the traffic load decreases after adaptive sampling has taken place.
 - For MX Series, PTX Series, and QFX Series, you can configure forwarding class and DSCP values per collector.
 - For MX Series, dual vlans are supported.
 - For MX Series, true output interface (OIF) is supported.
 - Enhancements are made to the following CLI commands: **show sflow collector**, **show sflow collector address ip-address**, and **show sflow interface**.

[See [Understanding How to Use sFlow Technology for Network Monitoring](#), [collector](#), [agent-id](#), [source-ip](#), [show flow collector](#), and [show flow interface](#).]

Port Security

- **Media Access Control Security (MACsec) support (MX Series)**—Starting with Junos OS Release 19.1R1, MACsec is supported on all QSFP interfaces on the MPC10E-15C and MPC10E-10C line cards when installed in an MX Series router. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is capable of identifying and preventing most security threats, and can be used in combination with other security protocols to provide end-to-end network security.

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

Routing Policy and Firewall Filters

- **Support for firewall forwarding on MPC10E-MRATE line cards (MX Series)**—Starting in Junos OS Release 19.1R1, firewall forwarding is fully supported on MPC10E-MRATE line cards for MX Series routers.

Routing Protocols

- **Support for BGP graceful shutdown (MX Series)**— Starting in Junos OS Release 19.1R1, graceful traffic migration from one BGP next hop to another is supported, without traffic interruption. Also, BGP administrative shutdown communication can be sent to the BGP peer.

You can configure both **graceful-shutdown** and **shutdown** statements at the `[edit protocols bgp]`, `[edit protocols bgp group group-name]`, and `[edit protocols bgp group group-name neighbor address]` hierarchy levels.

NOTE: Graceful shutdown is disabled by default.

[See: [graceful-shutdown \(Protocols BGP\)](#), [shutdown \(Protocols BGP\)](#).]

- **Support for anycast and prefix segments in SPRING for OSPF protocols (MX Series)**—Starting in Junos OS Release 19.1R1, anycast and prefix segments are supported in SPRING. An anycast segment enforces forwarding based on the equal-cost multipath-aware shortest-path toward the closest node of the anycast set. Within an anycast group, all the routers advertise the same prefix with the same SID value, which facilitates load balancing. You can designate prefix segment indexes to prefix SIDs, both anycast and node SIDs, that are advertised in OSPF through policy configuration. Remote routers use this index to consolidate prefixes into respective SRGBs and to derive the segment identifier and forward the traffic destined for a specific prefix.

You can also configure **explicit-NULL** flag on all prefix SID advertisements and configure **shortcut** statement for SPRING routes using family inet (for IPv4 OSPF routes) or family inet-mpls (for IPv4 L-OSPF routes).

[See: [Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING](#).]

- **Support for configurable SRGB used by SPRING in OSPF protocols (MX Series)**— Starting in Junos OS Release 19.1R1, 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 19.1R1, 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:

- SRGB under OSPF
- SRGB under MPLS
- Node-segment implementation of 256 label block

[See: [source-packet-routing \(Protocols IS-IS and OSPF\)](#).]

- **Support for static adjacency segment identifier in OSPF protocols (MX Series)**—Starting in Junos OS Release 19.1R1, static adjacency segment identifiers (SIDs) are supported for OSPFv2 protocols.

For static adjacency SIDs, the labels are picked from either a static reserved label pool or from an OSPF segment routing global block (SRGB). You can reserve a label range to be used for static allocation of labels using the following configuration: **set protocols mpls label-range static-label-range start-value end-value**

The static pool can be used by any protocol to allocate a label in this range. You need to ensure that no two protocols use the same static label. OSPF adjacency SIDs can be allocated from this label block through the configuration using the keyword **label**.

[See: [Static Adjacency Segment Identifier for OSPF](#).]

- **Support for export of BGP Adjacency-RIB-Out through BGP Monitoring Protocol (BMP) (MX Series and vMX)**— Starting in Junos OS Release 19.1R1, BMP is enhanced to support route monitoring of pre and post **rib-out** policy.

You can configure **post-policy** and **pre-policy** under **rib-out** statement at **[edit protocols bgp bmp]**, **[edit protocols bgp group group-name bmp]**, and **[edit protocols bgp group group-name neighbor address bmp]** hierarchies.

NOTE: The default monitoring mode of rib-out is **pre-policy**.

[See: [Understanding the BGP Monitoring Protocol](#).]

- **Support for TCP authentication to BGP peers (MX Series)**— Starting in Release 19.1R1, Junos OS extends support for TCP authentication to BGP peers that are discovered through allowed prefix subnets configured in a BGP group.

In releases before Junos OS Release 19.1, BGP supports TCP authentication at the **[edit protocols bgp group group-name neighbor address]** and **[edit protocols bgp group group-name]** hierarchy levels. Starting

in Junos OS Release 19.1, you can configure TCP authentication under allow statements at the **[edit protocols bgp group group-name dynamic-neighbor dyn-name]** hierarchy level.

[See: [Understanding Router Authentication for BGP.](#)]

- **Support for stitching of OSPF LDP and segment routing (MX Series)**—Starting in Junos OS Release 19.1R1, segment routing-LDP border router can stitch segment routing traffic to LDP next hop and vice versa.

In an LDP network with deployment of segment routing, there can be islands of devices that support either only LDP, or only segment routing. For the devices to interwork, the LDP mapping server feature is required to be configured on any device in the segment routing network.

[See: [LDP Mapping Server for Interoperability of Segment Routing with LDP Overview.](#)]

- **Support for BGP link-state distribution with SPRING extensions (MX Series)**—Starting in Junos OS Release 19.1R1, BGP link-state extensions export segment routing topology information to software-defined networking controllers. Controllers can get the topology information by either being a part of an interior gateway protocol (IGP) domain or through BGP link-state distribution.

BGP link-state distribution is supported on inter-domain networks and provides a scalable mechanism to export the topology information. This feature benefits networks that are moving to source packet routing in networking (SPRING) but also have RSVP deployed, and continue to use both SPRING and RSVP in their networks.

In this release, OSPF link-state protocol is supported which pushes SPRING information to the BGP link-state address family.

[See [Link-State Distribution Using BGP Overview.](#)]

- **MPLS transit route installation as primary MPLS fast reroute (FRR) for BGP labeled unicast prefixes (MX Series)**—Starting in Junos OS Release 19.1R1, when a peer autonomous system (AS) boundary router or a link fails, traffic traversing through an inter-AS link can be rerouted provided a loop-free path is available. In networks with node protection enabled, MPLS transit routes are installed as primary backup path for BGP labeled unicast prefixes learned from external BGP multi-hop sessions. This feature facilitates quicker route resolution and BGP convergence for BGP labeled unicast prefixes.

To enable node protection in an inter-AS environment for BGP labeled unicast prefixes, include the existing configuration statement **protection** at the **[edit protocols bgp group family inet labeled-unicast]** hierarchy level in **enhanced-ip network-services** mode.

- **Support for creating IS-IS topology-independent LFA for prefix-SIDs learned from LDP mapping server (MX Series)**—Starting in Junos OS Release 19.1R1, you can configure a point of local repair to create a topology-independent loop-free alternate backup path for prefix-SIDs derived from LDP mapping server advertisements in an IS-IS network. In a network configured with segment routing, IS-IS uses the LDP mapping server advertisements to derive prefix-SIDs. LDP Mapping server advertisements for IPv6 are currently not supported.

To attach flags to LDP mapping server advertisements, include the **attached** statement at the **[edit routing-options source-packet-routing mapping-server-entry mapping-server-name]** hierarchy level.

[See [prefix-segment-range](#).]

Services Applications

- **Support for tunnel interfaces on the MPC10E line card (MX Series)**—Starting in Junos OS Release 19.1R1, Junos OS supports three tunnel interfaces: generic routing encapsulation (GRE) tunnel, logical tunnel (LT), and virtual tunnel (VT) on the MPC10E line card.
 - The GRE tunnel interface supports the **tunnel** statement with these options: **destination**, **key**, **source**, **traffic-class** and **ttl**. The **copy-tos-to-outer-ip-header** statement is also supported.
 - The LT interface supports **family inet**, **family inet6**, and **family iso** options. The **encapsulation** statement supports the Ethernet and VLAN physical interface options only.
 - The VT interface supports the **family inet** option only.

[See [Tunnel Services Overview](#)]

- **Support for Port Mirroring on the MPC10E line card (MX Series)**—Starting in Junos OS Release 19.1R1, Junos OS supports port mirroring on the MPC10E line card. The MPC10E supports IPv4 (inet) and IPv6 (inet6) address families only.

[See [Configuring Port Mirroring](#)]

- **Support for inline flow monitoring on the MPE10E line card (MX Series)**—Starting in Junos OS Release 19.1R1, Junos OS supports inline active flow monitoring. Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Version 9 template is supported for IPv4, IPv6, MPLS, and MPLS-IPv4. IPFIX template is supported for IPv4, IPv6, MPLS, MPLS-IPv4, and VPLS flows. Both IPFIX and version 9 templates use UDP as the transport protocol.

[See [Understanding Inline Active Flow Monitoring](#)]

- **Support for automatic restart of Two-Way Active Measurement Protocol (TWAMP) Client**—Starting in Junos OS Release 19.1R1, the TWAMP client restarts automatically after a network failure, a configuration change, or an IP connectivity issue. However, for the client to reconnect to the TWAMP server automatically, you must use 0 as the **test-count** value in the **set rpm twamp client control-connection test-count** command. Also, at the TWAMP server side, the default value of **max-connection-duration** in the **set rpm twamp server max-connection-duration** must also be 0. You can display the test results after the network recovers, or after the server is reachable, by using the **set services rpm twamp client control-connection c1 persistent-results** command.

[See [Understanding TWAMP Auto-Restart](#)].

- **Support for Layer 2 services over GRE tunnel interfaces with IPv6 transport (MX Series routers with MPCs)**—Starting in Release 19.1R1, Junos OS supports Layer 2 Ethernet services over GRE interfaces with IPv6 traffic. After GRE encapsulates the packets, it redirects them to an intermediate host, where they are de-encapsulated and routed to their final destination. Support for bridging over GRE enables you to configure bridge domain families on gr- interfaces and also enable integrated routing and bridging (IRB) on gr- interfaces. To configure the bridge domain family on gr- interfaces, include the **family bridge** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level.

[See [Layer 2 Services over GRE Tunnel Interfaces on MX Series with MPCs](#).]

Software Defined Networking (SDN)

- **Support for in-chassis Junos Node Slicing (MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 19.1R1, Junos Node Slicing supports an in-chassis model, which allows all Junos Node Slicing components, such as Juniper Device Manager (JDM), base system (BSYS), as well as guest network functions (GNFs), to run within the Routing Engine of the MX Series router. To configure in-chassis Junos Node Slicing, ensure that the MX Series router has one of the following Routing Engines installed:

- RE-S-2X00x6-128 (used in MX480 and MX960 routers)
- RE-MX200X8-128G (used in MX2010 and MX2020 routers)

[See [Junos Node Slicing Overview](#) and [Configuring MX Series Router to Operate in In-Chassis Mode](#).]

- **Support for VXLAN on GNFs (MX480, MX960, MX2010, MX2020, and MX2008)**—Starting in Junos OS Release 19.1R1, guest network functions (GNFs) support EVPN with VXLAN encapsulation. This support enables you to configure GNFs to function as VXLAN Layer 2 or Layer 3 gateways. This support is also available on MX Series routers in LAN mode.

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#) and [Components of Junos Node Slicing](#).]

- **Abstracted fabric interface support for MS-MPC, 16X10GE MPC, MPC2E, MPC3E, MPC4E (MX480, MX960, MX2010, MX2020, and MX2008)**—Starting in Junos OS Release 19.1R1, Abstracted fabric (af) interfaces interoperate with the following line cards:

- Multiservices MPC (MS-MPC)
- 16x10GE MPC
- MPC2E
- MPC3E
- 32x10GE MPC4E
- 2x100GE + 8x10GE MPC4E

An abstracted fabric interface is a pseudointerface that facilitates routing control and management traffic between guest network functions (GNFs) through the switch fabric.

[See [Abstracted Fabric \(AF\) Interface](#).]

- **MS-MIC and MS-MPC support for in-chassis Junos Node Slicing (MX480, MX960, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 19.1R1, in-chassis Junos Node Slicing supports assignment of MS-MICs and MS-MPCs to guest network functions (GNFs). MS-MICs and MS-MPCs provide improved scaling and high performance, and possess enhanced memory and processing capabilities. The MS-MIC supports the Layer 3 services such as stateful firewall, NAT, IPsec, active flow monitoring, RPM, and graceful Routing Engine switchover (GRES). In-chassis Junos Node Slicing also support inline Layer 2 and Layer 3 services.

[See [Multiservices MIC and Multiservices MPC \(MS-MIC and MS-MPC\) Overview](#).]

- **Software resiliency support for in-chassis Junos Node Slicing (MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 19.1R1, software resiliency is enabled for guest network functions (GNFs) in the in-chassis Junos Node Slicing model. Resiliency enables the software to recover from certain types of failures. The in-chassis model allows all Junos Node Slicing components, such as Juniper Device Manager (JDM), base system (BSYS), as well as guest network functions (GNFs), to run within the Routing Engine of the MX Series router.

[See [Junos Node Slicing Overview](#).]

- **Multiversion software support for in-chassis Junos Node Slicing (MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 19.1R1, in-chassis Junos Node Slicing supports multi-version software interoperability, enabling the BSYS to interoperate with a guest network function (GNF), which runs a Junos OS version that is later than the software version on the base system (BSYS). This feature supports a difference of up to two versions between the GNF and the BSYS. That is, the GNF software can be up to two versions later than the BSYS software.

NOTE: The multiversion software compatibility support is limited to major releases only.

[See [Understanding Multi-Version Software Compatibility](#).]

- **Programmable flexible VXLAN tunnels (MX80, MX104, MX204, MX10003, and vMX)**—Starting in Junos OS Release 19.1R1, we support flexible VXLAN tunnels in a data center environment that includes one or more controllers. In this environment, one or more of the supported MX Series routers can function as data center edge gateways that exchange Layer 2 traffic with hosts in a data center. Through the use of static routes and tunnel encapsulation and de-encapsulation profiles, the Layer 2 traffic is dynamically tunneled over an intervening IPv4 or IPv6 network.

The controllers in the data center environment enable you to program a large volume of static routes and tunnel profiles on the gateway devices through the Juniper Extension Toolkit (JET) APIs.

[See [Understanding Programmable Flexible VXLAN Tunnels](#) and the [JET API Guide](#).]

Subscriber Management and Services

- **Control plane resiliency enhancements (MX Series)**—Starting in Junos OS Release 19.1R1, the following enhancements are available:
 - The master and standby Routing Engines exchange detailed information about session database replication. This exchange enables the Routing Engines to better determine whether the replication is correct.
 - You can configure the router to detect shared memory corruption and to automatically recover by rebooting the master or standby Routing Engine, or both. In earlier releases, a manual reboot is required to clear the corrupted shared memory; otherwise, it remains corrupted, causing processes that share the memory to generate core errors.

- You can monitor Routing Engine resiliency with the new **show system subscriber-management resiliency** command. The **summary** version indicates whether the system is functioning normally or an unexpected condition exists. The **detail** and **extensive** versions provide detailed information about the shared memory per Routing Engine.

[See [Junos OS Enhanced Subscriber Management](#) and [show system subscriber-management resiliency](#).]

- **Subscriber management support for in-chassis Junos Node Slicing (MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 19.1R1, in-chassis Junos Node Slicing supports all subscriber management features and services. Subscriber management provides capabilities such as subscriber access, authentication, and service creation, activation, and deactivation. The subscriber management services include DHCP, PPP, L2TP, VLAN, and pseudowire.

[See [Subscriber Management Overview](#).]

- **DHCP active leasequery for live updates of binding information (MX Series)**—Starting in Junos OS Release 19.1R1, you can configure active leasequery so that DHCP servers can provide an update to DHCP relay agents whenever the DHCP binding information changes. Individual and bulk leasequery provide updates only in a response to a query; subsequent changes are not reported to the relay agent until another query is made. Active leasequery also enables redundancy between relay agents to restore subscriber information if one of the peer relay agents reboots.

[See [DHCP Leasequery Methods](#).]

- **Display RPF check statistics for dynamic logical interfaces (MX Series)**—Starting in Junos OS Release 19.1R1, the **show interfaces statistics logical-interface-name detail** command can display byte and packet statistics for unicast RPF failures. These statistics are only displayed for dynamic IPv4 or IPv6 logical interfaces where RPF check is configured with the **rpf-check** or **rpf-check mode loose** statement. The **clear interfaces statistics logical-interface-name** command clears RPF statistics.

[See [Unicast RPF in Dynamic Profiles for Subscriber Interfaces](#).]

- **Additional encapsulations added to pseudowire subscriber logical interfaces (MX Series with MPC and MIC)**—Currently, the only supported encapsulation type on the pseudowire subscriber interfaces include:

- **Transport logical interfaces**—Circuit cross-connect (CCC) encapsulation.
- **Service logical interfaces:**
 - Ethernet VPLS encapsulation
 - VLAN bridge encapsulation
 - VLAN VPLS encapsulation

Starting in Junos OS Release 19.1R1, in addition to the existing encapsulation types, the following support is provided:

- **Transport logical interfaces**—Ethernet VPLS encapsulation, and provision for terminating the interface on the l2backhaul-vpn routing-instance.

- **Service logical interfaces**—Circuit cross-connect (CCC) encapsulation, and provision for terminating the interface on locally switched Layer 2 circuits.

[See [Pseudowire Subscriber Logical Interfaces Overview](#).]

- **Insert identifier tags in HTTP GET headers (MX Series)**—Starting in Junos OS Release 19.1R1, you can configure HTTP redirect service filters to insert tags into the headers of HTTP GET messages. You can specify one or more destination addresses in the service rule to identify traffic for tagging. The tagged message is forwarded to the HTTP server where the server can accept or reject access based on the tag values.

[See [Inserting GET Header Tags That the HTTP Server Can Use to Control Content Access](#).]

System Logging

- **Support for TCP/TLS transport for syslog (MX240, MX480, and MX960)**—Starting with Junos OS Release 19.1R1, you can configure multiple TLS syslog servers for a service on the MS-MPC or MS-MIC services cards. You can configure a maximum of four syslog servers for each set of services, and send encrypted data to the servers. The source address for the logs sent to remote hosts uses the configured source address of TCP/TLS host. See [TCP/TLS Transport Protocol for Syslog Messages Configuration Overview](#).

System Management

- **Support for SFTP global disablement (MX Series)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections from the CLI by default. You can enable incoming SFTP connections globally by configuring the statement **sftp-server** at the **[edit system services ssh]** hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections were globally enabled by default.

See [Configuring sftp-server](#)

SEE ALSO

[What's Changed | 111](#)

[Known Limitations | 121](#)

[Open Issues | 125](#)

[Resolved Issues | 142](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 195](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.1R3 | 111](#)
- [What's Changed in Release 19.1R2 | 113](#)
- [What's Changed in Release 19.1R1 | 117](#)

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

What's Changed in Release 19.1R3

General Routing

- **Change in startup notification after GRES (MX Series routers)**—The master Routing Engine sends a **coldStart** notification when a device comes up. The master Routing Engine also sends a **warmStart** notifications for subsequent restarts of the SNMP daemon. After GRES, the new master Routing Engine sends a single **warmStart** notification and the backup Routing Engine does not send any notification. In earlier releases, after GRES, the new master Routing Engine would sometimes send two notifications or a single notification. Of these, the first notification was always a **coldStart** notification and the second was either a **coldStart** notification or a **warmStart** notification.

[See [Standard SNMP Traps Supported by Junos OS](#).]

- **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.
- **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 lsdist.0 and lsdist.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 lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router ID.
- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view 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.

- **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.

Interfaces and Chassis

- **Change in error severity (MX960, MX240, MX2020, MX480, MX2008, and MX2010)**—Starting in Junos OS Release 19.1R3, we have reduced the severity of the CRC error 'XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR' from Fatal to Major. Earlier, these errors caused the line card to be reset, if the **interasic-linkerror-recovery-enable** statement was configured. Now, these errors only disable the Packet Forwarding Engines that are affected. With this change, the **interasic-linkerror-recovery-enable** statement has no effect when this error occurs because we have reduced its severity to Major.

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

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 19.1R3, 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](#).]

What's Changed in Release 19.1R2

EVPN

- **Support for disabling automatic ESI generation (MX Series and QFX Series)**—Starting with Junos OS Release 19.1R2, Junos OS supports disabling the automatic ESI generation for virtual gateway addresses. We recommend that you disable the automatic ESI generation for EVPN networks with edge-routed bridging to improve performance. To disable automatic ESI generation, include the **no-auto-virtual-gateway-esi** statement at the **[edit interfaces name irb unit logical-unit-number]** hierarchy level.

General Routing

- **User confirmation prompt for configuring the sub-options of request vmhost commands (MX Series and PTX series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm a **[yes,no]** for the sub-options also.
 - **request vmhost reboot**
 - **request vmhost poweroff**
 - **request vmhost halt**

In previous releases, the confirmation prompt was available for only the main options.

Interfaces and Chassis

- **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.

- **Logical Interface is created along with physical Interface by default (MX Series routers)**—In Junos OS Release 19.1R2 and later, by default logical interfaces are created on ge-, et-, and xe- interfaces along with the physical interface. In earlier Junos OS Releases, by default, only physical interfaces are created.

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

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (MX Series)**—Starting in Junos OS Release 19.1R2, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.

Routing Protocols

- **Change in the default behavior of advertise-from-main-vpn-tables configuration statement**—BGP now advertises EVPN routes from the main bgp.evpn .0 table. You can no longer configure BGP to advertise the EVPN routes from the routing instance table. In earlier Junos OS Releases, BGP advertised EVPN routes from the routing instance table by default.

[See [advertise-from-main-vpn-tables](#).]

Security

- On MX960 routers, the decapsulate GRE action now de-encapsulates GRE, IP-in-IP and IPv6-in-IP tunneling packets. You configure this action at the **[edit firewall family inet filter filter-name term term-name]** hierarchy level.

Services Applications

- **Change in NAT port block syslog message display (MX Series Routers)**—When you configure a software prefix other than 128, all the JSERVICES_NAT_PORT_BLOCK logs now displays the prefixed B4 address. The following JSERVICES_NAT_PORT_BLOCK are modified:

- JSERVICES_NAT_PORT_BLOCK_ALLOC
- JSERVICES_NAT_PORT_BLOCK_RELEASE
- JSERVICES_NAT_PORT_BLOCK_ACTIVE

In earlier releases of Junos OS, when a software prefix was configured, some of the B4 addresses displayed in the JSERVICES_NAT_PORT_BLOCK log were /128 addresses(irrespective of the configured prefix). This change is not observed when the software prefix is not configured.

- **New syslog message displayed during NAT port allocation error (MX Series Routers with MS MPC)**—With address pooling paired (APP) enabled, an internal host is mapped to a particular NAT pool address. In case, all the ports under a NAT pool address are exhausted, further port allocation requests from the internal host results in a port allocation failure. The following new syslog message is displayed during such conditions:

JSERVICES_NAT_OUTOF_PORTS_APP

This syslog message is generated only once per NAT pool address.

Software Defined Networking (SDN)

- **Increase in the maximum value of delegation-cleanup-timeout (MX 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](#).]

Subscriber Management and Services

- **Enhancement to commands to display reason for Routing Engine disconnect (MX Series)**—Starting in Junos OS Release 19.1R2, several commands display the reason when the master and standby Routing Engines disconnect because of a memory mismatch error. On a chassis with two Routing Engines, a DRAM size mismatch 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.

The **show database-replication summary** and **show system subscriber-management summary** commands display the DRAM mismatch as the reason in the Disconnection field. The **request chassis routing-engine master switch check** command displays an error message if the DRAM size is different for the two Routing Engines.

- **XML output format change for test aaa type user commands (MX Series)**—Starting in Junos OS Release 19.1R2, the XML output format changes for the **test aaa authd-lite user**, **test aaa dhcp user**, and **test aaa ppp user** commands. Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the <radius-server-data> tag. The new tag makes it easier to recognize the name/value pairs, both for operators and API clients. You may have to change any scripts that use the XML output to work properly with the new format.

[See [AAA Testing and Troubleshooting](#).]

- **Support for Pseudowire Physical Interface for ANCP Autoconfiguration (MX Series)**—Starting in Junos OS Release 19.1R2, you can associate an ANCP neighbor with a subscriber-facing interface pseudowire physical interface for ANCP autoconfiguration of VLANs. When configured, ANCP Port Up and Port

Down messages received on the interface trigger notifications to the auto-configuration daemon (autoconfd) to initiate VLAN creation (Port Up) or removal (Port Down). In earlier releases, ANCP supports only the following physical interface types for this feature: aggregated Ethernet (ae), Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), 100-Gigabit Ethernet (et), and demux.

- **Out-of-address SNMP trap requires thresholds to be configured (MX Series)**—Starting in Junos OS Release 19.1R2, the behavior has changed for generating an out-of-address SNMP trap for an address pool configured at the **[edit access address-assignment]** or **[edit routing-instance name address-assignment]** hierarchy levels. You must now configure both the high-utilization and abated-utilization thresholds. When the number of assigned addresses surpasses the high-utilization threshold, a high-utilization trap is generated. If all the addresses are assigned from the pool, an out-of-address trap is generated and an out-of-address syslog message is sent.

In earlier releases, an out-of-address trap is generated when the address pool is exhausted, regardless of whether the thresholds are configured.

If the number of assigned addresses subsequently drops below the abated-utilization threshold, an abate-high-utilization trap is generated; this behavior is unchanged.

- **Prevent queue-based throttling from stopping subscriber login (MX Series)**—Starting in Junos OS Release 19.1R2, 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.
- **Changing attributes of physical interface with active subscribers (MX Series)**—Starting in Junos OS Release 19.1R2, the commit check fails when you change any attribute of the physical interface, such as the MTU, when subscribers are active. This affects only aggregated Ethernet physical interfaces with targeted distribution configured. In earlier releases, the commit check does not fail and the attribute change brings down the physical interface and all subscribers using that interface.

What's Changed in Release 19.1R1

EVPN

- **Changes in encoding the ESI label field (MX Series)**—Starting in 19.1R1, Junos OS switched from using lower-order bits to higher-order bits in encoding the ESI label field. This results in BUM traffic loss and duplication in traffic. If you encounter this, and you wish to use a mix of Junos OS releases, you must include the `es-label-oldstyle` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy on the device that is running the Junos OS release that supports higher-order bit encoding of the ESI label.

General Routing

- **Root XML tag change for show rsvp pop-and-forward | display xml command (MX480)**—We've 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.](#)]

Interfaces and Chassis

- In MX204 routers, the error messages are logged when **vlan-tagging** for a trunk interface that is not configured. These error messages were previously logged with severity level "critical" even though they were not critical enough to require immediate action. The maximum transmission unit (MTU) of interface with or without VLAN-tagging is now logged in as the informational error message (instead of critical error message).
- **IRB not supported on pseudowire subscriber (PS) logical Interface in bridge-domain (MX Series)**—In Junos OS Release 19.1R1, integrated routing and bridging (IRB) is not supported on pseudowire subscriber (PS) logical Interface. Hence, you cannot add IRB to a bridge domain with PS interface, that is, you cannot configure IRB and PS interface in the same bridge domain.

Note that adding IRB to a bridge domain having pseudowire subscriber (PS) logical interface causes kernel crash and continuous reboot of the router until the configuration is rolled back.

NOTE: IRB is not supported on PS only in bridge-domain.

[See [bridge-domain.](#)]

- **Support for MAP-E encapsulation and de-encapsulation on inline service interfaces (MX2010)**—Starting in Junos OS Release 19.1R1, the MX2010 routers support encapsulation and de-encapsulation of the following ICMP message types for inline service (si) interfaces:
 - Time exceeded (type 11)
 - Destination unreachable (type 3)

- Source quench (type 4)
- Parameter problem (type 12)
- Address mask request and Address mask reply (type 17 and type 18)
- Redirect (type 5)
- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (MX Series)**—Starting in Junos OS Release 19.1R1, the **show lacp interfaces | display xml** command displays a new XML tag element <lacp-hold-up-state>. The <lacp-hold-up-state> displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces was in a single <lacp-hold-up-information> XML tag. Now, for each interface it is displayed in a separate <lacp-hold-up-information> XML tag.
- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, and later, MX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.
- **Support to get Optics loopback status for QSFP-100GE-DWDM2 transceivers (MX Series)**—Starting in Junos OS Releases 19.1R1, 19.1R2, and later on MX Series routers, you can get the optics loopback status of QSFP-100GE-DWDM2 transceivers along with the regular Ethernet loopback status by issuing the **show interfaces interface-name** or **show interfaces interface-name brief** command. The new output field **Optics Loopback** is added under **Link-level type** when the **show interfaces interface-name** CLI command is executed.

MPLS

- Starting in Junos OS Release 18.4R1 and 19.1R1, the remote procedure call (RPC) protocol XML tag for **mpls-label-value** is renamed as **mpls-history-label-value**, **mpls-usage-label-value**, and **mpls-label-id-value** depending on the context of command usage.
- **New debug statistics counter (MX Series)**—The **show system statistics mpls** command has a new output field, called **Packets dropped, over p2mp composite nexthop**, to record the packet drops over composite point-to-multipoint next hops.
- Starting in Junos OS Release 19.1R1, the **bfd-liveness-detection** statement is not supported at the **[edit protocols source-packet-routing segment-list]** hierarchy level.

Network Management and Monitoring

- **NETCONF <kill-session> operation returns different values in <rpc-error> when the session identifier is equal to the current session ID (MX Series)**—Starting in Junos OS Release 19.1R1, when you execute

the **<kill-session>** NETCONF operation and the session identifier is equal to the current session ID, the values of the **<error-type>** and **<error-tag>** elements in the resulting **<rpc-error>** are **application** and **invalid-value**, respectively. In earlier releases, the **<error-type>** and **<error-tag>** values are **protocol** and **operation-failed**.

[See [<kill-session>](#).]

- **sysName.0 MIB object displays the fully qualified domain name (MX Series)**—Starting in Junos OS Release 19.1R1, the sysName.0 MIB object displays the fully qualified domain name. That is, if the hostname and domain name are configured on the system, both will show up for the sysName.0 MIB object: **host-name.domain-name**. Previously, only the host name showed up.

[see [show snmp mib](#).]

Network Operations and Troubleshooting Automation

- Starting in Junos OS Release 19.1, the RPC XML tag for **mpls-label-value** is renamed as **mpls-history-label-value**, **mpls-usage-label-value**, and **mpls-label-id-value** depending on the context of command usage.

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when a change in number of supported history records is detected (ACX Series and MX Series)**—In Junos OS Release 19.1R1, when Ethernet connectivity fault management starts, it detects the number of history records supported by the existing performance monitoring history database and if there is any change from the number of history records supported (that is, 12) in Release 19.1R1, then the existing performance monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

Routing Protocols

- **Support for BGP LU link protection for a multihop EBGp peer (MX-Series)**—Starting in Junos OS Release 19.1R1, you can enable BGP Labeled unicast protection for an indirect next hop for logical-interface-based FRR. In earlier Junos Releases, Junos OS did not compute a backup path for the active indirect next hop failure and caused link failure for EBGp multihop cases where EBGp is chosen as a primary route for BGP LU protection on affected routes.

To configure BGP link protection for a multihop EBGp peer, enable **protection** at the **[edit protocols bgp group group-name family inet labeled-unicast]** hierarchy level.

Services Applications

- **Change in error message displayed while fragmenting or defragmenting IPv6 GRE tunnel interface (MX Series routers)**—In Junos OS Release 19.1R1, on an IPv6 GRE tunnel interface, when you enable fragmentation using the **allow-fragmentation** command or disable fragmentation using the **do-not-fragment** command, the following error message is displayed:

Fragmentation for V6 tunnels is not supported

In earlier Junos OS releases, the following message was displayed:

dcd_config_ifl_tunnel:Fragmentation for V6 tunnels is not supported

- **Support for host generated traffic on a GRE over GRE tunnel (MX Series)**—In Junos OS Release 19.1R1, you can send host generated traffic on a GRE over GRE tunnel. However, when path maximum transmission unit (PMTU) is updated for the outer GRE tunnel, MTU for inner GRE tunnel is not corrected.
- **Deprecated IPsec manual security association option (MX Series)**—In Junos Release 19.1R1 and later releases, the option **hmac-sha2-256** under the **services ipsec-vpn rule *rule-name* term *term-name* then manual direction (bidirectional | inbound | outbound) authentication algorithm** statement is deprecated. Use the **hmac-sha-256-128** option instead.

Software-Defined Networking (SDN)

- Starting in Junos OS Release 18.2X75-D30 and 19.1R1, the maximum value for service identifier (SID) depth for PCEP segment routing (SR) LSP has been increased to more than 5 labels. The supported range of **max-sid-depth** is 1 through 16 with a default value of 5 labels.

[See [pce](#).]

Subscriber Management and Services

- **ICMP error message rate limit increased (MX Series)**—Starting in Junos OS Release 19.1R1, the maximum rate limit for generating ICMP messages for IPv4 and IPv6 packet errors is increased from 50 pps to 1000 pps. The rate limit applies only to non-ttl-expired packets.

[See [Configuring the Rate Limit for ICMPv4 Error Messages](#) and [Configuring the Rate Limit for ICMPv6 Error Messages](#),]

- **Subscribers allowed to log in with bad framed route (MX Series)**—Starting in Junos OS Release 19.1R1, users are allowed to log in if the framed route received from RADIUS is bad; for example, if the format is incorrect. In earlier releases, the subscriber is not allowed to log in. For customers that use multiple framed routes, the new behavior enables the subscriber to have partial access to the network using the routes that are accepted instead of not being allowed any access.

User Interface and Configuration

- **Options for monitor traffic interfaces statement added (MX Series)**—Starting in Junos OS Release 19.1R1, the options **write-file** and **read-file** under the **monitor traffic** command are included in the visible CLI.

[See [monitor traffic](#).]

SEE ALSO

[What's New](#) | 88

[Known Limitations](#) | 121

[Open Issues](#) | 125

[Resolved Issues | 142](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 195](#)

Known Limitations

IN THIS SECTION

- [Fault Management | 121](#)
- [General Routing | 121](#)
- [Infrastructure | 123](#)
- [Interfaces and Chassis | 123](#)
- [MPLS | 123](#)
- [Platform and Infrastructure | 123](#)
- [Routing Protocols | 124](#)
- [Software-Defined Networking \(SDN\) | 124](#)
- [Subscriber Management and Services | 124](#)

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.

Fault Management

- The **Cmerror Op Set** log message is missed for the **bringup jspec** command-based error simulation. [PR1430300](#)

General Routing

- CFM is not supported for a Layer2-over-GRE tunnel. CCM can pass through as transit traffic through GRE interfaces transparently using the data path. Link trace functionality uses MAC learning and re-injecting LTM on the GRE interface in case the bridge is configured with CFM. This is not a supported feature. [PR1275833](#)
- The system watchdog might not work with secured BIOS. [PR1343131](#)

- The Routing Engine boots from the secondary disk when you:
 - Press the reset button on the RCB front panel, while the Routing Engine boots up before Junos OS reboots.
 - Upgrade the software by booting from the network using the **request vmhost reboot network** command and the system fails to boot from the network.
 - Upgrade the BIOS and it fails.
 - Reboot the system and it hangs before Junos OS reboots. [PR1344342](#)
- If MTU is configured to a value higher than 9500, which is the maximum permissible value, configuration succeeds. However, the actual value sets back to 1518B without any error. The DCD log can be checked to verify the occurrence. [PR1372690](#)
- The MIC-MACSEC-20G supports 10 Gbps speed through the **set chassis fpc x pic y pic-mode 10G** configuration applied to both the PICs in that MIC. You must remove any other PIC mode configuration and then apply the 10 Gbps PIC mode configuration. [PR1374680](#)
- The memory usage of processes increases significantly between Junos OS Release 14.1X53-Dxx to Junos OS Release 17.x. [PR1390226](#)
- The OSPF adjacencies flap after GRES, with LACP fast timer configuration. [PR1392316](#)
- On an MX2008 platform with MPC9E, in a line-rate traffic with a redundant SFB2 scenario, if there is one offline redundant SFB2, there might be tail drops or sometimes WRED drops in MPC9E. This results in partial traffic loss. Under normal circumstances, the SFBs should automatically fail over if one of them fails and there should be only a few packets dropped momentarily. [PR1395591](#)
- Configuration database can remain locked after the SSH session is halted. [PR1410322](#)
- When SRTE policies have segment lists that have a single label or three or more labels, the IS-IS interface statistics are not incremented even when SRTE routes take these IS-IS next hops. The kstat-based states are enabled in IS-IS by the set **protocols isis source-packet-routing traffic-statistics** command. [PR1410682](#)
- Line cards get rebooted when BSYS performs Junos OS-only reboot.
 - Ensure that before performing BSYS Junos OS-only reboot in **in-chassis node-slicing** mode, switch the mastership to the other Routing Engine.
 - Ensure that before performing BSYS Junos OS image upgrade and reboot in **in-chassis node-slicing** mode, switch the mastership to the other Routing Engine.
 - Note that BSYS Junos OS-only unified ISSU is not supported in **in-chassis node-slicing** mode; you need to use the **request vmhost software in-service-upgrade** option for this. [PR1413810](#)
- The MX Series Packet Forwarding Engine does not account for the labels pushed onto the packet on the egress Packet Forwarding Engine, while PTX Series Packet Forwarding Engine does. This results in slight difference in the byte count for the same traffic stream across these two platforms. The packet-count is still the same across the platforms. Currently, this issue is noticed for uncolored SRTE policies. [PR1416738](#)

- The repd generates core files during the third GRES on RE1 with reference to **BbeStatsSessionDb::repdConvert (this=0x886b018, rep_msg=0x89a3814)** at `../../../../../../src/junos/lib/libbbe-stats/bbeStatsSession.cc:395`. [PR1417732](#)
- In a large-scale setup (such as, large number of routing-instances or interfaces), if there are frequent changes in configuration and interface flapping when the rpd is restarted (through deactivate/activate logical system or restart routing), the rpd might crash. [PR1438049](#)
- The **FPC x Voltage Tolerance Exceeded** alarm is raised and cleared upon bootup of JNP10K-LC2101. [PR1415671](#)

Infrastructure

- On Juniper Networks Routing Engines with Hagiwara CompactFlash card installed, after the upgrade to Junos OS Release 15.1 and later, the following error message is observed: **smartd[xxxx]: Device: /dev/ada1, failed to read SMART Attribute Data might appear on message logs**. [PR1333855](#)

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](#)

MPLS

- With NSR enabled, when the master rpd is restarted, occasionally, **out-of-order** add and delete messages can arrive on the backup Routing Engine. As a result, label assignment collisions occur, which cause the backup rpd to crash. [PR1401813](#)
- The rpd process might crash. [PR1461468](#)

Platform and Infrastructure

- On all Junos OS platforms, execution of Python scripts through enhanced automation does not work on veriexec images. [PR1334425](#)
- When only peer 1 advertises routes, that peer might not install the de-encapsulated next-hop (NH) route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any type-5 route. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. [PR1386423](#)

Routing Protocols

- When 32,000 SRTE policies are configured at once, during configuration time there might be scheduler slips. [PR1339829](#)
- When you replace the simple VLAN interfaces with PS over RLT, you might see an increase in FPC CPU usage. This is in keeping with the increased processing and resources needed to support these types of interfaces, which are similar in this regard to the requirements for an aggregated Ethernet interface. [PR1396925](#)
- The rpd process might crash during convergence or route update due to incorrect path selection algorithm. [PR1352697](#)

Software-Defined Networking (SDN)

- The MX Series platform type of the guest network function (GNF) configured on an MX Series chassis does not automatically change if the Routing Engine is installed on a different MX Series chassis type. To fix this issue, you need to delete the GNF and configure it from the start on the new chassis in which Routing Engine is installed.
- When guest network functions (GNFs) are rebooted for different reasons, the **show chassis routing-engine** command may incorrectly display the reboot reason as **Router rebooted after a normal shutdown**. To find the actual reboot reasons, see the log messages of GNFs.
- External Ethernet port LEDs on Control Board of MX2020 and MX2010 routers do not turn off when the network-slices configuration is deleted or deactivated.
- If you try to install Juniper Device Manager (JDM) after performing **request vmhost zeroize**, the installation is unsuccessful. As a workaround, you can delete JDM and install it again.
- The PS interface maximum transmission unit (MTU) size at times has an incorrect default value. As a workaround, you can delete the PS interface and configure it again.
- Junos OS Release 19.1R1 does not interoperate with earlier releases of Junos OS that support Junos node slicing. To run the 19.1R1 version of Junos node slicing on any GNF, the BSYS and all other GNFs must also run Junos OS Release 19.1R1.

Subscriber Management and Services

- For dual-stacked clients over the same PPP-over-L2TP LNS session, enhanced subscriber management does not support configurations where both of the following are true:
 - The CPE sends separate DHCPv6 solicit messages for the IA_NA and the IA_PD.
 - The solicit messages specify a type 2 or type 3 DUID (link-layer address).

As a workaround, you must configure the CPE to send a single solicit message for both IA_NA and IA_PD when the other configuration elements are present. [PR1441801](#)

SEE ALSO

[What's New | 88](#)

[What's Changed | 111](#)

[Open Issues | 125](#)

[Resolved Issues | 142](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 195](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 126](#)
- [EVPN | 126](#)
- [Forwarding and Sampling | 127](#)
- [General Routing | 127](#)
- [High Availability \(HA\) and Resiliency | 134](#)
- [Infrastructure | 134](#)
- [Interfaces and Chassis | 135](#)
- [Junos fusion Enterprise | 136](#)
- [Junos fusion Provider Edge | 136](#)
- [Layer 2 Ethernet Services | 136](#)
- [MPLS | 136](#)
- [Network Management and Monitoring | 137](#)
- [Platform and Infrastructure | 137](#)
- [Routing Protocols | 139](#)
- [Subscriber Access Management | 141](#)
- [User Interface and Configuration | 141](#)
- [VPNs | 141](#)

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)

- This issue is related to the configuration of the hidden statement **rate-limit-burst** in the **class-of-service** hierarchy. The commit needs to push an update for CoS code handling on all the Packet Forwarding Engines and during this time, an interface setting (internal attributes for an interface) was found to be NULL. Interface settings are usually stored in a memory location and the pointer to it became NULL because cosd did not check for the NULL values and resulted in segmentation fault. The channelized interface setting was found to be NULL for channelized interfaces, but the CoS code handling the configuration **rate-limit-burst** in Packet Forwarding Engine de-referenced the setting without doing a NULL check, resulting in core files. [PR1425667](#)

EVPN

- There might be a few duplicate packets seen in an A/A EVPN scenario when the remote PE device sends packets with IM label due to MAC not being learned on remote PE device but being learned on the A/A local PE device. The non-DF sends the IM-labeled encapsulated packet to the PE-CE interface after MAC lookups instead of dropping the packet, which causes duplicate packets on the CE side. [PR1245316](#)
- Due to timing condition, the dead next hops in the flood group of EVPN-MPLS are seen after remote PE devices bounce. [PR1484296](#)
- On the MX, QFX10002, QFX10008, and QFX10016 devices within an EVPN-MPLS active-active multi-homing mode, when adding and deleting the virtual-gateway-address, the ARP entry to CE might disappear from the kernel. [PR1485377](#)
- When you try to execute VXLAN ping overlay through the RPC command, an RPC error is observed. [PR1373025](#)
- On the MX104 platform, the **chained-composite-next-hop ingress evpn** statement is missing in the **junos-defaults** group. This statement must be configured to make EVPN work prior to Junos OS Release 18.3R1-S1/18.3R2, otherwise EVPN does not work as expected. [PR1415466](#)
- The rpd process might generate a core file when the Routing Engine switches over after disabling the BGP protocol globally. [PR1490953](#)
- The VXLAN bridge domain might lose VTEP logical interface after restarting chassisd. [PR1495098](#)
- EVPN core isolation might not work if the BGP GR is (Graceful Restart) enabled. [PR1496229](#)

Forwarding and Sampling

- The skip-service configuration does not work with IPv6 NDP negotiation or ping. [PR1074853](#)
- In a BGP scenario with sampling enabled, incorrect ASN might be returned for the traffic originated from an internal prefix. This is because some AS paths and routes do not hold the latest information in the message buffers that the sampling route-record daemon uses to send to the clients. [PR1439630](#)
- Additional information for the **set firewall flexible-match source-ipv6-match ...** command must be provided to avoid confusion. [PR1389103](#)
- On a Junos fusion, ingress policing on the SD is broken (MX Series and QFX Series; ingress on AD and SD) and the **set interfaces layer2-policer input-policer** command is not supported in this release. [PR1395217](#)
- For Junos OS Releases 18.4R1 and 18.3R2, if an IPv4 prefix is added on a prefix list referred by an IPv6 firewall filter, the following log message is not seen in this particular release: **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized.** [PR1395923](#)
- For Junos OS Release 19.1R1, for physical interface policer for **ip-option traffic**, the traffic rate is more than 10 percent. [PR1398728](#)
- The **show firewall filter global-edge-filter** command does not return any values. [PR1422407](#)
- When you verify the selective local-switching functionality with 4000 VLANs, error messages related to traffic not being policed as expected after locally switching it to VLAN 100 and 101 are observed. [PR1436343](#)
- After routing is restarted, the remote mask (indicating from which remote PE devices MAC-IP entries are learned) that the routing daemon sends can be different from the existing remote mask that the Layer 2 learning daemon had prior to restart. This causes a mismatch between the Layer 2 learning and routing daemons' interpretation as to where the MAC-IP entries are learned (local or remote). This leads to the mac-ip table being out of synchronization. [PR1452990](#)
- Packet length for ICMPv6 is displayed as 0 in the output of the **show firewall log detail** command. [PR1184624](#)
- The following syslog error message might be observed due to SSD hardware failure: **Failed connecting to DFWD, error checking reply - Operation timed out.** [PR1397171](#)

General Routing

- If a Layer 3 interface is receiving a GRE-encapsulated packet and the interface has two filters attached at ingress as follows:
 - (a) family **any** with action as **mirror**.

- (b) family **inet** with action as **decapsulate gre**, then the expected behavior is that the mirrored copy must have the GRE headers as well. However, that is not working as expected because of the presence of filter (b). If you are interested in mirroring the entire packet that came on the interface (that includes GRE header as well), then the workaround is to deactivate or disable the **decapsulate gre** action of filter (b). [PR1090854](#)
- In a BGP or MPLS scenario, if the next hop type of label route is indirect, then the following changing events about the next hop interface MPLS family might cause the route to be in the Dead state, and the route remains in that state even when the family MPLS is again activated:
 - Deactivating and activating the interface family mpls.
 - Deleting and adding back the interface family mpls.
 - Changing maximum labels for the interface.

When a labelled route is resolved over an interface, that interface must have family mpls configured for the route to be successfully resolved. Otherwise, the route does not get resolved. [PR1242589](#)

- In an EVPN-MPLS scenario, if the core-facing interface (mpls interface) and the CE-facing interface are on different PEs, and the traffic from core is not continuous and DMAC (Dynamic MAC) ages out, due to an incorrect flood next-hop programming across different PFEs, packet loss might be observed after device rebooted or I2ald restarted. [PR1484468](#)
- The **ptp-clock-global-freq-tracable** leaf value becomes false and does not changes to true) when the internal lock is in the **Acquiring** state. [PR1493743](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- The input errors on the MX150 router devices might be zero under the **show interfaces extensive** output when there are CRC or Align errors on the interface. [PR1485706](#)
- In some scenarios with PTP hybrid mode and MPC5E, continuous resetting of the Playback Engine log message occurs. Playback engine resides inside the MPC5E FPGA and it is responsible for maintaining the PTP states corresponding. [PR1420335](#)
- On the MX Series platforms, if the clock frequency slowly changes on CB0 (slow drift), the clock source for MPC-3D-16XGE-SFPP might not be changed to CB1, which cause interfaces on it to go down and remain in the **Down** state. [PR1433948](#)
- On the MPC7E, MPC8E, and MPC9E line cards, hardware sensor information is logged on the syslog and **/var/log/messages** every 30 minutes. [PR1478816](#)
- The following syslog message appears: **fpcX user.notice logrotate: ALERT exited abnormally with [1]**. [PR1471006](#)
- Changing the framing modes on a CHE1T1 MIC between E1 and T1 on a MPC3E NG HQoS line card causes the PIC to go offline. [PR1474449](#)
- MPC10E line card might crash due to inconsistencies during firewall filter add or delete operations that results in traffic being silently discarded. [PR1465153](#)

- IPv6 accounting stop attributes are not correct for MLPPP subscribers. [PR1455175](#)
- The following error message is observed that has no functional impact: **MIB2D_RTSLIB_READ_FAILURE: check_rtsock_rc: failed in reading mac_db(xe-5/0/1:0): 0 (Invalid argument)** [PR1461289](#)
- SIP session fails when the IPv4 SIP client in a public network initiates a SIP call with the IPv6 SIP client in the private network. [PR1139008](#)
- When performing a Routing Engine switchover without the support of nonstop routing (NSR), occasionally the Layer 2 Control Protocol daemon reports a slips of a few seconds (1 to 10) in its scheduled run and a similar log message as the following is displayed: **l2cpd[32770]: JTASK_SCHED_SLIP: 8 sec scheduler slip, user: 0 sec 2180 usec, system: 0 sec, 2188 usec**. This delayed run has no functionality or operational effect on any of the Layer 2 protocols controlled by Layer 2 Control Protocol daemon. This is because the STP task delegates transmit or receive BPDUs to a separate dedicated PPMD daemon and the LLDP task transmit or receive PDUs are dealt from the daemon itself. However, the advertisement interval is 30 seconds, with the hold timer value for neighbors' LLDP PDUs being 120 seconds, so the time to recover the few seconds of slips is plenty and enough to absorb it. [PR1203977](#)
- The following cosmetic error is observed as the output: **mshpmand[190]: msvcs_session_send: Plugin id 3 not present in the svc chain for session. Please open a JTAC case to confirm.** [PR1258970](#)
- On the vMX platform, performance of the Intel X710 NIC is lower compared to the performance of Intel 82599 NIC. This issue occurs because 10-Gbps rate can be achieved at 512 byte packet size for X710 NICs, whereas the same can be achieved at 256 bytes for 82599 NICs. [PR1281366](#)
- If a vmhost snapshot is taken on the alternate disk and there is no further vmhost software image upgrade, the expectation is that on the current vmhost image that gets corrupted, the system boots with the alternate disk so you can recover the primary disk to restore the state. However, under the condition where corruption is with host root file system, the node boots with the previous vmhost software instead of the alternate disk. [PR1281554](#)
- On MX204 and MX10003, the Routing Engine might get stuck and boot from the other SSD after vmhost reboot. This is a race condition during BIOS handoff to Junos OS. You must boot the Routing Engine from the primary SSD. [PR1295219](#)
- The user-configured packet hashing options for **inet** family under **enhanced-hash-key** might not take effect for the TRIO based FPCs. FPC keeps using the default behavior for hash calculation for IPv4 packets. [PR1302637](#)
- The **show dynamic-tunnels database summary** command does not show an accurate tunnels summary when the anchor Packet Forwarding Engine line card is not in the up state. As a workaround, use the **show dynamic-tunnels database** and **show dynamic-tunnels database terse** commands. [PR1314763](#)
- The customer does not use the **chain-composite** statement. The **chain-composite** statement does not bring in a lot of gain, because the TCNH is based on the ingress rewrite premise. Without this statement, things work just fine. [PR1318984](#)
- In JDM (running on secondary server), the jdmd daemon might generate core files if GNF add-image is aborted by pressing Ctrl+C. [PR1321803](#)

- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infra relies on the integrity of the TCP connections. The reactions to failed situations might not be handled in a graceful way---for example, the TCP connection might time out because of jlock hog crossing the boundary value (5 seconds), which might cause bad consequences for the MX Series Virtual Chassis. Currently, there aren't any easy solutions to reduce this jlock hog besides enabling marker infra in the MX Series Virtual Chassis setup. [PR1332765](#)
- The first packet pertaining to J-Flow Packet Forwarding Engine sensor in UDP mode is missing after line-card reboot. [PR1344755](#)
- With Graceful Routing Engine switchover (GRES) enabled in a subscriber environment, if subscribers are logging in or out very quickly, the service sessions in the session database of the backup Routing Engine might be leaked. If the problem is not detected for a long time, the backup Routing Engine might not be able to come back into synchronization with the master Routing Engine and might not be ready for GRES. [PR1346300](#)
- The backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- During unified ISSU that warrants host upgrade, if the router is configured with 8 million IPv4 or IPv6 routes or more, the unified ISSU might fail, resulting in FPC restart. [PR1348825](#)
- In some cases, online insertion and removal (OIR) of a MIC installed in an MPC might lead to the silent dropping of traffic destined to the MPC. The only way to recover from this is to restart the MPC. The issue is not be seen if you use the corresponding CLI commands to take the MIC offline and then bring it back online. [PR1350103](#)
- The following error messages are observed with Junos OS Release 17.3 throttle image:
localtcp_offload_tx_errcheck: failed to send packet 4 times in last one second. [PR1359149](#)
- 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, this scaled configuration on MPC7E takes settling time of more than 4 minutes. [PR1359286](#)
- For the INH -> FWD NH{List NH} -> {Chain NH} scenario, when rpd reads next hops from the kernel on restart, rpd must not create an old-style list next hop for the forwarding next hop. [PR1360354](#)
- 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 also shows alarm. [PR1360426](#)
- When an FPC is booting up (during unified ISSU, router reboot, or FPC restart), the Layer 2 circuit timeout error messages for the SFP transceiver are observed. These errors are seen when the Layer 2 circuit action is not completed because the device was busy. When the FPC is up, all the Layer 2 circuit transactions to the device were all right, so no periodic failure is observed. There is no functional impact and these errors can be ignored. [PR1369382](#)

- If any log messages continue to appear in the MPC console, they indicate the presence of a faulty SFP/SFP+ transceiver. 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 2x10GE SFPP/20x10GE SFP MACSEC in PIC 0 that results in unexpected behavior:

I2C Failed device: group 0xa0 address 0x70Failed 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: 0mic_sfp_id_read: Failed to select link 0

For example, the link does not come up or the MIC itself is not booting up on restart.

The only way to recover from these failures is to detect and replace the faulty SFP/SFP+ transceiver plugged into the 2x10GE SFPP/20x10GE SFP MACSEC ports. [PR1375674](#)

- A few xe- interfaces go down with the following error message: **error if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error.** [PR1377840](#)
- In low-end 32-bit systems, rpd has a lower level of available memory. We need to display a log message to alert customers when the average memory usage or transient memory usage exceeds thresholds. [PR1387465](#)
- Control Plane Switch Management (CPSM) daemon memory leak occurs in vmhost. It might also cause the log rotate to stop working and cause large CPSM log size. [PR1387903](#)
- On an MX Series platform enabled with enhanced subscriber management, if the filter service is enabled for each subscriber and if there is a large scale of broadband edge (bbe) subscribers (for example, 10,000) logging in and out repeatedly, the FPC might crash due to this rare issue. [PR1388120](#)
- The virtio throughput remains same for the multiqueue and single-queue deployments. [PR1389338](#)
- If the **persist-groups-inheritance** statement is configured when you try to add additional sites to an existing group and routing instance configuration, you might observe an error that can cause the commit to fail after issuing commit check. [PR1391668](#)
- On MX Series platforms, if the channelized OC MIC (such as, 1xCOC12/4xCOC3 CH-CE) is used, the MPC/AFEB/TFEB (Forwarding Engine Board) might crash, generating core files. This is not easily reproducible. The traffic through the MIC might be impacted. [PR1396538](#)
- The Junos OS rpd daemon has facilities to attempt to trap certain classes of nonfatal bugs by continuing to run, but leaves a soft generated core file. Leaving a soft generated core file is intended to be nondisruptive to routing and forwarding. You must have a mechanism by which you can disable the soft generated core files. [PR1396935](#)
- Router advertises the ESMC QL of PRC even though the current clock status is holdover. [PR1398129](#)
- In a BGP-PIC case, if route R1 resolves on top of multipath-route R2, where R2 has primary and backup indirect next hops, it is better if the backup leg is not used for the resolution of R1. There is no impact on any existing CLI commands. The backup path must not be used when the primary path is available. [PR1401322](#)

- Generated core file and rpd reboot is observed when the **condition-manager** policy is configured for the routing table and the same table is repeatedly deleted and read. [PR1401396](#)
- For ON_CHANGE subscriptions of /interfaces/ sensors, the sync_response is sent to the collector before the complete data is sent to the collector. However, there is no behavioral impact because the collector still receives complete data and the synchronization response received earlier impacts the ONCE mode of subscriptions. [PR1403672](#)
- On vMX-based platforms including MX150, when you run the **clear pim join instance instance-name all** command, it might result in stopping of the riot process on the system. [PR1409527](#)
- For MPC10 and MPC11 line cards, firmware upgrade for INphi modules should use updated version of scripts that support these line cards. The old script fails to perform the firmware upgrade. [PR1410133](#)
- The configuration database can remain locked after the SSH session is halted. [PR1410322](#)
- After ISSU, some of L2TP and PPPoE subscribers do not come up. [PR1412818](#)
- A small number of tunneled subscribers might be terminated during unified ISSU to Junos OS Release 19.1R1 software due to momentary loss of IP connectivity between the LAC and LNS devices. [PR1414928](#)
- When the Routing Engine software is not able to access the fabric chip, the corresponding plane goes into fault state. The following log message is observed: **CHASSISD_FASIC_PIO_READ_ERROR: Fchip (CB 1, ID 1): read error in sfchip_init() for link#100 at address 0 in register FCHIP_FTOP_CONFIG.** [PR1416814](#)
- On the NG-MPC's KATS, tests does not run post system reboot and the chassisd restarts with 2x 10GE SFPP / 20x 1GE SFP MACsec MIC. [PR1418538](#)
- Certain JNP10008-SF and JNP10016-SF manufactured between July 2018 and March 2019 might have an incorrect core voltage setting. As a workaround, reprogram the core voltage and update the setting in the NVRAM memory. [PR1420864](#)
- With HTTP header enrichment function enabled, the processing of the window scaling option significantly reduces the performance of HTTP sessions from 65 MB/s to less than 40 MB/s, which results in decrease of traffic throughput. The download rate drop is also observed. [PR1420894](#)
- On all platforms running Junos OS, when the file system gets into the Full state and there is not enough spare disk space, it might get into a problematic system condition while committing the configuration. After that, if the consecutive commits are still done in such a problematic condition, commit-check failure logs might be seen eventually. Due to this issue, some processes might not run even if the configuration is present. [PR1423500](#)
- The fast-lookup-filter does not work for the MPC10E-15C-MRATE line card. During installation, the fast-lookup-filter is converted internally to the Dmem filter. [PR1431451](#)
- The **my-mac-check-failed** exception counter display is missing from the CLI output, but the functionality is working as expected. [PR1438761](#)

- If a route is programmed with **Rib.RouteAdd/Update()** with a VXLAN flexible tunnel profile having a destination MAC set to all zeros, reading the route back through **Rib.RouteGet()** results in the destination MAC not being set in the RouteGetReply. [PR1439940](#)
- On the MPC7E, MPC8E, and MPC9E cards, egress stream flush failure and silent dropping of traffic might occur in a rare occasion for a repeatedly flapping link. [PR1441816](#)
- The BGP session that establishes over the GRE tunnel fails when the router receives the BGP packets encapsulated as GRE and uses the firewall filter action to de-encapsulate the GRE header. [PR1443238](#)
- For the EVPN instance with integrated and routing (IRB) interface, the proxy ARP and ARP suppression is enabled by default. With EVPN proxy ARP and ARP suppression enabled and in Junos OS Release 17.2Rx or 17.3Rx ($x \geq 3$, the correlated service release is also affected), the kernel processing on the master Routing Engine might crash due to a software defect on the packet handling. This is a rare issue. [PR1443903](#)
- After fixing PR 1338647, error dropped packets are seen on MQ/XM-based MPC cards, though there is no traffic flowing through the system. [PR1451958](#)
- When using the **replace pattern** command to replace the name in the apply-group, the mgd crashes. [PR1452136](#)
- Timestamp is missed when the show command is edited and run from the CLI command history. [PR1454387](#)
- On a device with MPC10E installed, the firewall filter might be incorrectly updated in the Packet Forwarding Engine when the change (for example, add/delete, deactivate/activate) of firewall filter terms occurs in some scenarios, such as large-scale terms change or changes that occur during the MPC reboot. The incorrect firewall filter might cause the traffic to be silently discarded and leads to an MPC crash. It is a timing issue. [PR1458499](#)
- Commit script does not apply changes in private mode unless you perform a full commit. [PR1465171](#)
- From Junos OS Release 16.2R1 and onwards, if the **commit** command is executed after **commit check**, the daemon (for example, dhcpd and sampled) might not get started even though the related configuration is successfully committed. [PR1468119](#)
- Unable to see generated core files in the **show system coredumps** output when you use the hostname with "." in between. [PR1474118](#)
- A new aggregated Ethernet member interface for a static LAG bundle is installed in the FIB even though the micro- BFD session is down. [PR1474300](#)
- When you reboot the external server, the SNMP values configured within the **/etc/snmp/snmpd.conf** file at the server gets overwritten with the content from the JDM **snmp** configuration section. The trap configuration changes get completely removed. Restarting or stopping and starting JDM does not change the host **/etc/snmp/snmpd.conf** file. Only system reboot of the server occurs. [PR1474349](#)
- There is a behavior change from Junos OS Release 18.4 where the DHCPv6/PD subscribers come over the PPPoE subscriber. The DHCP accounting statistic is not displayed for DHCPv6/PD when brought over PPPoE. The behavior of PPPoE remains the same. [PR1461104](#)

- The **max-drop-flows** statement is not available. [PR1375466](#)
- Revert of RLT to primary can silently discard traffic for around 10 minutes after the primary FPC is online with primary RLT up. [PR1394026](#)
- FPC might crash after performing GRES in a subscriber scenario. [PR1421541](#)
- Random packet drops are observed with flow cache disabled when NIC is mapped to NUMA node 1. [PR1458742](#)
- The MPC2E-NG or MPC3E-NG line card with specific MIC might crash after a high rate of interface flap. [PR1463859](#)
- The GRE tunnel might go down in the scenario with IPv4 and IPv6 IPsec service configured. [PR1470667](#)
- When you reboot the external server, the SNMP values configured within the `/etc/snmp/snmpd.conf` file at the server get overwritten with the content from the JDM SNMP configuration section. The trap configuration changes get completely removed. Restarting or stopping and starting JDM does not change the host `/etc/snmp/snmpd.conf` file. Only system reboot of the server occurs. [PR1474797](#)
- **xqchip_drop_get_q_length** severity of parity error moved from major to minor. [PR1481558](#)
- The SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. [PR1484322](#)
- In node slicing setup after GRES, the RADIUS interim updates might not carry the actual statistics. [PR1494637](#)
- Taking packet-via-dmem capture might cause FPC to crash or wedge. [PR1498422](#)

High Availability (HA) and Resiliency

- On the MX204 and MX10003 Virtual chassis, ISSU might fail with an error message. [PR1480561](#)

Infrastructure

- The following message is seen during FTP: **ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock(No such file or directory)**. [PR1315605](#)
- If the serial number of the PEM starts with **1F1**, the following alarm might be generated: **Minor FPC PEM Temp Sensor Failed** [PR1398128](#)
- The following error message is observed continuously in AD with base configurations: **IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) failed**. [PR1485038](#)

Interfaces and Chassis

- Out-of-sequence packets are seen with the LSQ interface. [PR1258258](#)
- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with a CFM configuration might cause the cfmd process to crash after the upgrade. This is because of the presence of an old version of `/var/db/cfm.db`. [PR1281073](#)
- In MX Series Virtual Chassis, flooding of the following error message can be seen with LACP-enabled aggregated Ethernet interfaces on MPC7, MPC8, and MPC9 cards: **CHASSISD_CONFIG_ACCESS_ERROR: pic_parse_ifname: Check fpc rnage failed**. The error message impacts only DWDM PICs, which does not effect the MPC7, MPC8, and MPC9 cards. Hence, this syslog message can be safely ignored. [PR1349277](#)
- LFM sessions toward scaled peers might flap during the unified ISSU switchover phase. [PR1377761](#)
- Static demux0 logical interfaces do not come up after configuration changes if the underlying interface is an et- (100-Gigabit Ethernet) interface. After the configuration change, the et- interface gets flushed in order to reparse the configuration. During this time, dcd does not create the dependency between demux0 logical interfaces and the underlying et- interface, which results in flushing off the demux0 logical interfaces. This issue is seen only if the underlying interface is an et- interface. As a workaround, restart dcd, reboot the entire Routing Engine, or use the **commit full** command instead of the **commit** command while committing the new configuration. [PR1401026](#)
- If an aggregated Ethernet interface has the VRRP configuration, in the following use cases, member logical interfaces cannot be created after the member physical interface comes up, and the aggregated Ethernet interface will be in the Down state: **fpc restart (request chassis fpc restart slot <>) chassis-control restart (restart chassis-control) reboot both Routing Engine (request system reboot both-routing-engines)**

Therefore, before performing these operations, remove the VRRP configuration from the aggregated Ethernet interface. [PR1429045](#)
- When the **QSFP-100GE-DWDM2** transceiver is installed on the device, the harmless severity level log might be flooded periodically by this transceiver. [PR1453919](#)
- The following message might be seen while committing configuration on MXVC: **registration is being denied**. [PR1431377](#)
- SFBs goes to the **Check** state and the Packet Forwarding Engine raises major **CM_CMERROR_FABRIC_SELFPING** errors in the MX2000 routers, if the SFB is offline while it is in online process. [PR1433522](#)
- Traffic might be dropped because of the next-hop points to ICL even though the local MC-LAG is up. [PR1486919](#)

Junos fusion Enterprise

- The SDPD generates core files at `vfpc_all_eports_deletion_complete vfpc_dampen_fpc_timer_expiry`. [PR1454335](#)

Junos fusion Provider Edge

- V44 or Junos Fusion Environment system , intermediate traffic drop is seen between AD and SD when sflow is enabled on ingress interface. This is not seen always. [PR1450373](#)

Layer 2 Ethernet Services

- On an MX Series platform, if a static demux interface is configured over an underlying interface after the subscriber logs out, then the accounting statistics are not cleared. [PR1383265](#)
- The jdhcpd process might leak memory in the subscriber scenario. [PR1491349](#)

MPLS

- With nonstop active routing (NSR), when the rpd restarts on the master Routing Engine, the rpd on the backup Routing Engine might restart. [PR1282369](#)
- If there are two directly connected BGP peers established over MPLS LSP and the MTU of the IP layer is smaller than the MTU of the MPLS layer, and also if the BGP packets from the host have the DF bit set, the BGP session might keep flapping because of the usage of the wrong TCP-MSS. [PR1493431](#)
- For an SR-TE path with 0 explicit NULL as the innermost label, the SR-TE path does not get installed with label 0. [PR1287354](#)
- The root XML tag in the output has been changed from **rsvp-pop-and-fwd-info** to **rsvp-pop-and-fwd-information** to be consistent with the XML tag convention. [PR1365940](#)
- The following multicast (MVPN) traffic drop can be observed on switchover when vt- is in use: **show routing-instances MVPN-1 instance-type vrf; interface vt-0/0/0.1202 { <<<< multicast; } //snip**
This issue occurs due to the presence of incorrect forwarding states on the backup Routing Engine before the switchover. [PR1434522](#)
- On MX Series platforms, in an MPLS Layer 2 circuit or Layer 2 VPN with FAT (Flow-Aware Transport of Pseudowires) flow labels scenario, the flow label is not pushed when **chained-composite-next-hop ingress l2ckt/l2vpn** is enabled. The issue results in a load-balance problem for the Layer 2 circuit or Layer 2 VPN service. [PR1439453](#)
- When an interface in an MVPN routing instance is changed from a virtual tunnel (VT) interface to a label-switched interface (LSI), the P2MP LSP might get stuck in an incorrect state because no tear down message is created from the LSP egress side. In the end, the MVPN traffic is lost. [PR1454987](#)

- In a large scale P2MP deployment, LSPs might go down randomly across the network due to repeated make-before-break event occurring in the P2MP sub-lsps. [PR1415384](#)
- The rpd process might crash in PCEP for the RSVP-TE scenario. [PR1467278](#)
- The backup Routing Engine might crash if an indirect next-hop is sent by the primary Routing Engine without the associated sgid. [PR1493053](#)

Network Management and Monitoring

- Junos OS sends a cold trap from the new master Routing Engine just after the first GRES. This is because the **cold_start** timestamp file was not present or updated after the reboot. So, for the first GRES, Junos OS sends a cold start trap. [PR1461839](#)

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the following error message: **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient with the system reconverging on the expected state. [PR1054798](#)
- On an MX-VC setup, if the traffic goes through a VCP port and forwards to an egress port to the destination, while the traffic is handled entirely by the same Packet Forwarding Engine, MAC malformation might occur. [PR1491091](#)
- In the MX104 chassis, the **show system buffer** command display all zero. [PR1484689](#)
- An accuracy issue occurs with three-color policers of both types, single rate, and two rate in which the policer rate and burst-size combination of the policer accuracy vary. [PR1307882](#)
- In an Ethernet frame padding with VLAN scenario, if the fragment is a must, the VLAN Ethernet padding is required to minimize the frame size when the Ethernet frame length is less than 68 bytes, and equal to or greater than 64 bytes. But if the VLAN Ethernet padding is configured on a **vlan-tagging** enabled Ethernet interface, the fragment might not work correctly. The MPC error might be observed and the traffic might get lost. [PR1452261](#)
- An MX Series router acting as an egress NG-MVPN PE router with **hot-root-standby** and **sender-based-rpf** features enabled, is supposed to switch to the backup multicast flows from the primary ones if the primary flow rate falls below configured **mvpn hot-root-standby min-rate rate**. The switchover time (when the traffic is silently discarded) is expected to be within 50 milliseconds irrespective of the number of the groups which are undergoing this switchover (assuming that this number is within maximum MoFRR flows supported). Before this PR fix, that switchover time could be more than 50 milliseconds (up to several seconds) if the number of the groups undergoing the switchover at the egress PE simultaneously is greater than ~250 groups. [PR1478981](#)

- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur. Even with enhanced convergence configured, there is no guarantee that subsecond convergence will be achieved. [PR1371493](#)
- In some cases, the status bit of the RPF next hop appears as disabled when it should have been enabled. [PR1404240](#)
- On MX Series routers with MS-MPCs, when both group virtual private network (GVPN) tunnel and MPLS-over-UDP tunnel are used, traffic from a group virtual private network (GVPN) tunnel to MPLS-over-UDP tunnel might fail to get decrypted on the MS-MPC. This causes a complete service loss. [PR1422242](#)
- On all platforms running Junos OS, with NSR enabled, the BGP session with a hold-time of 6 seconds or smaller flaps after the backup Routing Engine is pulled out ungracefully. [PR1428518](#)
- For the bridge domains configured under an EVPN instance, the ARP suppression is enabled by default. This enables the EVPN to proxy the ARP and reduces the flooding of ARP in the EVPN networks. Due to this, the storm-control does not take effect on the ARP packets on the ports under such bridge domain. [PR1438326](#)
- Arrival rates are not seen at the system level when the global-disable FPC is configured. [PR1438367](#)
- A dual Routing Engine Juniper node slicing GNF with no GRES configured and with **system internet-options no-tcp-reset drop-all-tcp** configured might enter a dual backup Routing Engine state upon a manual GNF Routing Engine mastership switchover attempt with the **request chassis routing-engine master [acquire|release|switch]** CLI command from either GNF Routing Engine CLI. [PR1456565](#)
- After performing ISSU with a scaled configuration, high CPU utilization is observed in the MPC 3D 16x10GE card. [PR1461715](#)
- In NTP with the boot server scenario, when the router or switch boots, the NTP daemon sends an 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 is booting, the ntpdate might not work successfully. Then, some cosmetic error messages of time synchronization might be seen, but there is no impact on the time update because the ntp daemon updates the time eventually. [PR1463622](#)
- EVPN traffic loss is observed with the initial configuration before verifying with the **IRB IP next-hop: type-5** with no EVPN inside data center functionality. [PR1466914](#)
- The unicast traffic from IRB interface towards LSI interface might get dropped with the aggregated Ethernet load-balance adaptive or per-packet configured. [PR1458825](#)
- In node slicing setup, MPLS TTL might be set to zero when the packet goes through the aggregated Ethernet interface configured with the circuit cross-connect family. [PR1492639](#)

Routing Protocols

- While interoperating with other vendors in a draft-rosen multicast VPN, by default Junos OS attaches a route target to multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities are prevented from propagating if the BGP route-target filtering is enabled on the device running Junos OS. [PR993870](#)
- If a manually configured rib-group or automatically generated rib-group (through **family inet labeled-unicast resolve-vpn**) is used to copy inet.0 (IP routing table) routes to inet.3 (MPLS routing table), the process rpd might continuously generate soft core files after **protocols bgp path-selection always-compare-med** is configured. [PR1487893](#)
- On the MX2010 Series routers, the BFD session on the IS-IS step up flaps during the ISSU - FRU upgrade stage. [PR1453705](#)
- When configuring an alternate incoming interface for a PIM RPF check using rpf-selection, you might find that the additional groups outside the configured range switches to the alternate incoming interface. [PR1443056](#)
- JTASK_SCHED_SLIP for rpd might be seen on performing restart routing or OSPF protocol disable with scaled BGP routes in an MX104 router. [PR1203979](#)
- Certain BGP traceoption flags (for example, open, update, and keepalive) might result in trace logging of debugging messages that do not fall within the specified traceoption category, which results in some unwanted BGP debug messages being logged to the BGP traceoption file. [PR1252294](#)
- LDP and OSPF are in the Synchronization 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 0.0.0.01Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050Adj count: 1Hello: 10, Dead: 40, ReXmit: 2, Not StubAuth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTCProtection type: NoneTopology default (ID 0) -> Cost: 1050LDP sync state: in sync, for: 00:04:03, reason: IGP interface downconfig holdtime: infinity

According to the current analysis, **IGP interface down** is observed because although LDP notified OSPF that LDP synchronization was achieved, OSPF was not able to take note of the LDP synchronization notification because the OSPF neighbor was not up yet. [PR1256434](#)
- BGP I/O thread is added in Junos OS Release 16.1R1, whereby BGP writes were batched to improve efficiency. This might sometimes lead to some latency in sending BGP updates while reacting to certain network events. [PR1332301](#)
- SCP command with routing option (-JU) is not supported. [PR1364825](#)
- It is possible for a GNF with rosen6 multicast to display stuck KRT queue entries after recovery from a dual Routing Engines reboot at the BSYS. [PR1367849](#)

- The following error messages are observed: **pimd_rtrequest_v4(1133), IS_MASTER_RE: 1, Process: rpd, RTM_ID: 5, error: 17, errmsg: rt exists; ifindex = 340**. These logs are not harmful and have no functional impact. This error message just shows the state of PIM register messages. These logs are already LOG_DEBUG for external builds. You do not need to make any change in any of the components. [PR1371431](#)
- At scale, a GNF with PS over RLT and multiple MPCs might show BFD flap at recovery. [PR1386574](#)
- In a BGP scenario with multipath enabled, if you apply an import or export policy of IPv6 routes with an IPv4 next hop to a BGP neighbor, the rpd might crash continuously. [PR1390428](#)
- If an import policy is applied to a BGP neighbor and the policy has an indirect IPv4 next hop for IPv4 and IPv6 routes (IPv6 routes resolved over IPv4), when the BGP unresolved route is withdrawn, rpd crash might be seen. [PR1391568](#)
- Policy-based label allocation is not supported for an IPv6 prefix. Commit might be successful but configuration does not take effect. There is no functional impact. [PR1395040](#)
- When the MoFRR feature is used in a scaled environment (in terms of number of routes and next hops), the actual convergence of multicast traffic might reach hundreds of milliseconds due to suboptimal handling of MoFRR forwarding states at the Packet Forwarding Engine level. [PR1399457](#)
- Sometimes when a new logical router is configured, logical router core files might be seen on the system if the kernel is reporting low memory (this core file is harmless). [PR1403087](#)
- During NSR initial state replication on a scaled setup, while BGP state replication is still going on, the BGP task replication might get marked as completed. This is because BGP replication is triggered and controlled by the backup Routing Engine. You must check the output of the **show BGP replication** command to confirm whether replication has actually completed. This corner case scenario is valid only on a scaled setup and during initial state synchronization. [PR1404470](#)
- The mcsnoopd core files are generated immediately after the commit change related to a VXLAN-EVPN configuration. [PR1408812](#)
- Autotranslations fail when the static adjacency SID is configured with OSPF as IGP. [PR1414612](#)
- Variability of convergence improvement with the **send-addpath-optimization** statement enabled is observed. [PR1395684](#)
- Packet drop and CPU spike on Routing Engine might be seen in certain conditions if the **labeled-unicast** protection is enabled for a CsC-VRF peer. [PR1456260](#)
- RPKI validation is broken. [PR1464931](#)
- The bbe-smgd process generates core files on the backup Routing Engine. [PR1466118](#)
- RPD crashes due to the flapping of the BGP session. [PR1490079](#)

Subscriber Access Management

- Authd reuses addresses too quickly before jdhcpd can completely clean up the old subscriber, which floods the error log; for example: `jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815.` [PR1402653](#)
- While verifying the deleting services through CoA, when the specified family-type has been deactivated since incorrect number of service sessions are active. Hence, some services remains active even though the family is deleted. [PR1479486](#)
- Subscriber accounting messages retransmissions exist even after configuring accounting retry 0. [PR1405855](#)

User Interface and Configuration

- The **test configuration /config/file** configuration fails to check for a dynamic profile when the subscriber is active. [PR1376689](#)
- Changing nested apply groups does not take effect. [PR1427962](#)
- Previous configuration might still take effect after performing rollback rescue. [PR1489575](#)

VPNs

- The multicast VPN MIB is not properly compiled into the Juniper MIB package bundle. This might cause `mib-jnx-mvpn.txt` to be included as part of the Juniper Enterprise MIB set. [PR1394946](#)
- MVPN traffic loss observed while verifying MULTICAST ROUTE with VT for VPNA. [PR1460480](#)
- In an 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](#)

SEE ALSO

[What's New | 88](#)

[What's Changed | 111](#)

[Known Limitations | 121](#)

[Resolved Issues | 142](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 195](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3 | 142](#)
- [Resolved Issues: 19.1R2 | 155](#)
- [Resolved Issues: 19.1R1 | 181](#)

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

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

Resolved Issues: 19.1R3

Application Layer Gateways

- SIP messages that need to be fragmented might be dropped by SIP ALG. [PR1475031](#)

Class of Service (CoS)

- MX Series device-generated OAM/CFM LTR message are sent with a different priority than the incoming OAM/CFM LTM message. [PR1466473](#)
- Unexpected traffic loss might be observed in certain conditions. [PR1472083](#)
- The MX10008 and MX10016 devices might generate cosd core files after executing the **commit** or **commit check** command if the **policy-map** configuration is set. [PR1475508](#)

EVPN

- Delay factor might send back ARP request/NS to the local segment under EVPN-ETREE leaf role conditions. [PR1459830](#)
- Traffic received from the VTEP gets dropped if the VNI value used for type-5 routes is greater than 65,535. [PR1461860](#)
- The rpd might crash after the EVPN-related configuration is changed. [PR1467309](#)
- The rpd process might crash after committing changes in the EVPN environment. [PR1439537](#)
- MAC address present in the EVPN database are not get installed in the EVPN MAC-table (ethernet-switching table) when the remote MAC is learned from the multihomed PE devices. [PR1463722](#)

Forwarding and Sampling

- The pfded might crash and not be able to come up. [PR1452363](#)
- In some rare scenarios upon FPC or PIC reboot, the Packet Forwarding Engine daemon database might not get updated with the correct location_id for some physical interfaces, then a problem with statistics on some interfaces of a router might be observed. [PR1458143](#)
- The l2ald process might experience memory leak on platforms running Junos OS. [PR1455034](#)
- In a scenario when the VGA on IRB is deactivated, activated, or commit the configuration at the same time on both the PEs of a site, the type 1 ESI or AD route might not be generated locally. [PR1464778](#)
- The following syslog error message is observed: **pfed: rtslib: ERROR received async message with no handler: 28** [PR1458008](#)
- The following logs are observed: **L2ALD_MAC_IP_LIMIT_REACHED_IF: Limit on learned MAC+IP bindings reached for .local.1048605; current count is 1024** [PR1462642](#)
- On MX10008 and MX10016 routers, the **bandwidth-limit** policer cannot be set higher than 100 gigabits. [PR1465093](#)
- 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](#)
- Traffic might not be forwarded into the right queue but might be forwarded into 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** option is present in the filter. [PR1478964](#)

General Routing

- The following error message is reduced from fatal to major: **XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR**. [PR1390333](#)
- Mandatory TLV ttl learnt from LLDP neighbors is not streamed along with other learnt parameters from the neighbors. [PR1459441](#)
- On the MX150 device, there is a log severity level change. [PR1411846](#)
- Original PFEBM task, which is system-critical for internal network performance or resilience runs as a medium priority. [PR1429797](#)
- FIPS:GMIC2-KATS does not run on specific MPC3E-3D-NG scenario and the following error message is observed: **gi2mic_vsc8490_port_init: FIPS mode not set** [PR1418538](#)
- Certain JNP10008-SF and JNP10016-SF that are manufactured between July 2018 and March 2019, might have incorrect core voltage setting. The issue can be corrected by reprogramming the core voltage and updating the setting in the NVRAM memory. [PR1420864](#)
- During host ping with gr- tunnel endpoint, lt- interface termination, gr- interface input, and lt- interface, the output counters comes as host path with addition to the transit counter. [PR1461593](#)

- The rpd might generate core files when being terminated either during a user-initiated restart or when deactivating a logical system. This crash is seen only when rpd is being shut down, so overall impact on the network is minimal. [PR1418192](#)
- jnxFruState shows the 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](#)
- On all the devices running Junos OS, with channelizing configured on FPCs, if a 40-Gbps port that is channelized to 10-Gbps ports already (for example, xe-2/0/16:0) is being channelized to 10-Gbps again, the port might get incorrectly channelized. [PR1423496](#)
- On the MX Series platform, the PPP sessions do not work properly. [PR1428212](#)
- The **Protect core** configured router might send IPFIX sampling packets with the wrong next-hop information. [PR1430244](#)
- The l2cpd process might crash and generate a core file when the interfaces flaps. [PR1431355](#)
- Micro-BFD 3x100 ms flap upon inserting a QSFP transceiver in another port. [PR1435221](#)
- When you reboot or power off the backup Routing Engine, trap message reported is seen. This is a generic design for the TVP platform. [PR1436212](#)
- Unified ISSU is failing from Junos OS Release 19.1R1 legacy images. (Seen only on legacy image) [PR1438144](#)
- The rpd might generate core files during router bootup due to a file pointer issue as there are two code paths that can close the file. [PR1438597](#)
- Incorrect values appears in **JUNIPER-TIMING-NOTFNS-MIB**. [PR1439025](#)
- The rpd creates route pointing to chain composite instead of indirect next hop for a PE-PE directly connected case. [PR1439317](#)
- On Junos OS, if a group is applied at the nonroot level and later a statement from the group is deleted, change bits are not set for the hierarchy where the group is applied. As a result, the respective daemon is not notified for the changes that causes mgd to generate core files. [PR1439805](#)
- New OID is added that calculates the buffer utilization where inactive memory is not considered as free memory. [PR1441680](#)
- The MX Series device might rewrite the HTTPS request with destination port 80. [PR1446085](#)
- The static route for NAT might never come up if you switch over the service interface that has NAT and GR configuration. [PR1446267](#)
- The rpd process might crash when it is terminated immediately after it has been started. [PR1446320](#)
- The bbe-smgd generates core file on the 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](#)
- FPC generates the **Voltage Tolerance Exceeded** major alarm for each IP 2V5 sensors. [PR1451011](#)

- Main chassisd thread at a JNS GNF might experience stalls upon GNF SNMP polling for hardware-related OIDs. [PR1451215](#)
- RMPC-generated core files are found after configuration changes are done on the network for PTP/Synchronous Ethernet. [PR1451950](#)
- Incorrect output is observed in the **show snmp mib walk jnxTimingNotfnsMIB.3** statement. [PR1453436](#)
- On the MX10003 platform, alarm is not sent to syslog. [PR1453533](#)
- On MPC3E-NG cards with 100-Gigabit Ethernet interface in use, if the interface detects Loss of Lock (LOL) on the link without Loss of Signal (LOS), the interface goes down and might not come up again after the link is recovered. [PR1454595](#)
- After a software upgrade, SNMP MIB Walk does not roll or fetch any information. [PR1455667](#)
- The CLI command with **invoke-on** and **display xml rpc** results in unexpected multiple RPC commands. [PR1456578](#)
- Continuous crashing of the **bbe-statsd** process might be seen if any parameter is set to 0 in the **mx_large.xml** file. [PR1457257](#)
- The chassisd process and all FPCs might restart after a Routing Engine switchover. [PR1457657](#)
- If the dynamic assignment of VoIP VLAN is used, the switch might not send correct VoIP VLAN information in the LLDP MED packets after any configuration change and commit. [PR1458559](#)
- The **FPC X major errors** alarm might be raised after committing the PTP configuration change. [PR1458581](#)
- The following error messages might be seen after restarting the chassisd: **create_pseudos: unable to create interface device for pip0 (File exists)**. [PR1459373](#)
- Incomplete output is observed when you run the **show ancp subscriber access-aggregation-circuit-id < access aggregation circuit ID>** statement. [PR1459386](#)
- In an MC-LAG scenario, traffic destined to VRRP virtual MAC gets dropped. [PR1459692](#)
- Traffic is discarded silently upon interface flap after DRD autorecovery. [PR1459698](#)
- AUTO-CORE-PR :cpdd generates core files at
**ServicesMgr::ServicesManager::cpddSmdInterface::processServiceNotifyMsg
,SmdInterface::cbStateSyncServiceNotifyMsgHandler ,statesync_consumer_poll_new_state_cb.**
[PR1459904](#)
- In a subscriber management environment, subscriber statistics reported by CLI commands and RADIUS can be broken if in-service software upgrade (ISSU) is performed from any Junos OS release earlier than Junos OS Release 18.4 to 18.4 or a later build. [PR1459961](#)
- The PTP does not work with destination NAT. [PR1460027](#)
- If **VLAN-offload** is configured on the vMX platform, **input-vlan-map** might not work. [PR1460544](#)
- IPv6 prefix might be hidden when received over an IPv4 BGP session. [PR1460786](#)
- The PTP function might consume kernel CPU for a long time. [PR1461031](#)

- The **bbe-smgd** might generate core files when all RADIUS servers are unreachable. [PR1461340](#)
- In an EVPN scenario, memory leak might be observed when **proxy-macip-advertisement** is configured. [PR1461677](#)
- The repd core file is seen during system bootup. [PR1461796](#)
- There is a memory issue with the BBE statistics collection and management process, **bbe-statsd**, on the backup Routing Engine. [PR1461821](#)
- When both DIP switches and power switch are turned off, the following error message appears:
CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed. [PR1462065](#)
- On the MX204 platform, **radius-acct-interim** statistics are not populated for subscribers. Statistics are properly populated in the **radius-acct-stop** packets. [PR1462325](#)
- An interface might get stuck in the Down state on certain MX Series platforms. [PR1463015](#)
- On an MX Series platform with enhanced subscriber enabled, if you make some changes to a dynamic-profiles filter, the subscribers built on the filter might no longer forward traffic. [PR1463420](#)
- RPC ALG causing mspmand to generate core files when an MX Series device is used as a stateful firewall with the MS-MIC or MS-MPC service cards. [PR1464020](#)
- The IPoE subscriber route installation might fail. [PR1464344](#)
- The **bbe-smgd-core** (0x000000000088488c in **bbe_autoconf_delete_vlan_session_only** (session_id=918) file is observed at
../..../src/junos/usr.sbin/bbe-svcs/smd/plugins/autoconf/bbe_autoconf_plugin.c:3115).
[PR1464371](#)
- The PPP IPv6CP might fail if the **routing-services** command is enabled. [PR1464415](#)
- The CPU utilization on the mgd daemon might get stuck at 100 percent after the netconf session gets interrupted by the flapping interface. [PR1464439](#)
- The MS-MIC might not work when it is used on a specific MPC. [PR1464477](#)
- The PPPoE session gets in the **Terminated** state and the accounting stop for the session is delayed. [PR1464804](#)
- MPC5E or MPC6E line cards might crash due to internal thread overusing the CPU. [PR1464820](#)
- In the MAC-MOVE scenario, the jdhcpd might consume high CPU and no further subscribers can be brought up if there are more than 4000 DHCP relay clients. [PR1465277](#)
- The physical interface of an aggregated Ethernet interface bundle might take time to come up after you enable or disable it. [PR1465302](#)
- ICMP error messages are still not received after you enable the **enable-asymmetric-traffic-processing** command. [PR1466135](#)
- The PPPoE subscribers get stuck due to the PPPoE inline keepalives that do not work properly. [PR1467125](#)

- Layer 2 wholesale is not forwarding all client requests with stacked VLAN. [PR1467468](#)
- The rpd process might stop after several changes to the flow-spec routes. [PR1467838](#)
- Crypto library shim memory utilization performance improvement by using data shim instead of control shim. [PR1467874](#)
- SNMP interface-mib stops working for PPPoE clients. [PR1470664](#)
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card causes major alarms on unrelated remote FPCs. [PR1471372](#)
- When PTP is configured in the hybrid mode, the Synchronous Ethernet frequency drifts. [PR1471502](#)
- On the MX10008 and MX10016 line cards, the ARP suppression that is enabled by default in EVPN does not work. [PR1471679](#)
- Service accounting statistics do not get updated after changes are made to the firewall filters. [PR1472334](#)
- The kernel might crash and vmcore might be observed after committing the configuration change. [PR1472519](#)
- Performing back-to-back rpd restarts might cause rpd to crash. [PR1472643](#)
- SDB goes down very frequently if the **reauthenticate lease-renewal** statement is enabled for DHCP. [PR1473063](#)
- Ingress multicast replication does not work with GRES configuration. [PR1474094](#)
- clksyncd generates core file after GRES. [PR1474987](#)
- The RADIUS accounting updates of service session have wrong data statistics. [PR1475729](#)
- The **bbe-mibd** process crashes as a result of aborting `../src/junos/lib/libjuniper++/memHeap.cc:226`. [PR1476596](#)
- In a NAT-T scenario, the IKE version 2 IPsec tunnel might flap if the tunnel initiator is not behind NAT. [PR1477483](#)
- On MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9, the PFE might be disabled due to a major error. [PR1478028](#)
- FPC memory leak might occur after executing the **show pfe route** command. [PR1478279](#)
- The rpd walks through irrelevant routing tables when you run the **show route protocol l2-learned-host-routing** or **show route protocol rift** command and this causes rpd to crash while processing a **bgp.rtarget.0** route. [PR1481953](#)
- UID might not be released properly in some scenario after the service session deactivates. [PR1188434](#)
- The **show subscriber extensive** command incorrectly displays DNS addresses that are provided to DHCP clients. [PR1457949](#)
- DHCP relay with forward-only might fail to send OFFER messages when the DHCP client is terminated on the logical tunnel interface. [PR1471161](#)

- The dynamic profile for a VPLS-PW pseudowire incorrectly reports the Dynamic Static Subscriber Base Feature license alarm. [PR1473412](#)
 - The statistics of the traffic generated by the Routing Engine on the MX platform is incorrect. [PR1432724](#)
 - In the Junos OS PTP deployment, when configuring the aggregated Ethernet child interface in the protocol PTP and after performing the FPC restart, all of the interfaces on that FPC might be brought to the Administrator down state. [PR1442665](#)
 - The IS-IS setup sends the system host-name instead of system-id in the OC paths in the Isdb or Adjacency xpaths in a periodic streaming and on-change notification. [PR1449837](#)
 - On subscribing to `/interfaces/interface/state/`, it might get subscribed to `/junos/system/linecard/interface/traffic/` internally instead of `/junos/system/linecard/interface/queue/` that might impact traffic. [PR1456275](#)
 - The traffic might get stuck on the MS-MPC or MS-MIC with sessions receiving huge number of affinity packets. [PR1459306](#)
 - FDB does not flush that causes the traffic to be dropped silently in the Ethernet ring scenario. [PR1459446](#)
 - Fabric hardening (FH) is the process of controlling bandwidth degradation to prevent traffic to be silently discarded. When the FH process is processing and if SFB or SCB fails, the FH process gets stuck that causes traffic lost. [PR1461356](#)
- On the specific Junos OS platforms, the Routing Engine switchover might not be triggered when the master CB clock failure is detected. The master CB with the faulty clock can not operate normally and this issue might cause fabric plane failure. [PR1463169](#)
- Need to add the Backport jemalloc profiling CLI support to all Junos OS Releases where jemalloc is present. [PR1463368](#)
 - When the tunnel-services are configured on a PIC, the optics measurements that subscribed through gRPC might not be streamed. [PR1468435](#)
 - The pccd might generate core files and PCEP session might flaps in the PCE initiated or PCE delegated LSP scenario. [PR1472051](#)
 - The tcp-log connections fail to reconnect and gets stuck in the **Reconnect-In-Progress** state. [PR1469575](#)
 - On all the Junos platforms with l2cpd daemon, committing configuration changes that are processed by l2cpd (for example, **flexible-vlan-tagging**, **stacked-vlan-tagging**, **vlan-tagging**, and **family ethernet-switching**) might cause marginally memory leak. [PR1469635](#)
 - A **hierarchical-scheduler** should not be configured on a ps- interface. The corresponding anchor lt-interface must be used for that purpose instead. [PR1470049](#)
 - The following chassis alarm is reported on BSYS: **RE0 to one or many FPCs is via em1: Backup RE**. [PR1472313](#)
 - The subnet information might be corrupted if it is passed by a RADIUS server. [PR1474097](#)

- The QSA adapter Lane 0 port might be also brought down when you disable one of the other lanes. [PR1474231](#)
- In the MX2000 chassis with multiple line cards (XM and EA based) and SFB2 fabric, a longer traffic drops in the order of 30 seconds depending on the number of line-cards in the chassis might be observed. [PR1476505](#)
- Traffic loss might occur due to the LNS subscribers in case the **routing-service** statement is enabled under the dynamic-profile. [PR1476786](#)
- The TCP-log sessions might be in the **Established** state but the logs are not sent out to the syslog server. [PR1478972](#)
- The rpd process might crash when the prefix list address is changed from IPv4 to IPv6. [PR1421076](#)
- On the MX104 Series router with any 2 port license installed on 10-Gigabits interface and phy-timestamping enabled in PTP, PTP might not work. [PR1421811](#)
- The default configuration does not create any logical interfaces and LLDP cannot discover the neighbor for those interfaces, which the logical interface is not configured explicitly in the Junos OS configuration. [PR1436327](#)
- When flapping the existed flood next hop type routes, it might cause rpd crash or consume 100 percent of CPU. This issue might cause the routing protocols sessions or neighbors flap or traffic loss. [PR1441550](#)
- The MVPN traffic might be dropped after performing switchover. [PR1463302](#)
- On the MX104 Series routers, the clksyncd crash might be observed when PTP over an aggregated Ethernet is configured. [PR1471466](#)
- The protocol MTU might not be changed to lt- interface from the default value. [PR1478822](#)
- The MX204 router might restart if the system is configured for the subscriber management for PPPoE over aggregated Ethernet with the routing-service is enabled in the dynamic-profile. [PR1482431](#)
- After kmd restarts, the IPsec SA comes up but the traffic fails for some time. [PR1480692](#)
- On the MPC10 line cards, the logical tunnel interface might not work. [PR1484751](#)
- The following general routing syslog error messages might be generated by the Packet Forwarding Engine: **PFEIFD: Could not decode media address with length 0.** [PR1341610](#)
- The default credentials is supplied in the configuration. [PR1344858](#)
- High CPU utilization of the FXPC process might be observed with the **class-of-service** changes on interfaces. [PR1407098](#)
- Need to disable PTP and display warning when hyper mode is configured. [PR1429527](#)
- IRB over VTEP unicast traffic might get dropped. [PR1436924](#)
- The chassisd process might crash. [PR1460657](#)
- The chassisd process might crash during GRES on the device with SCBE3 or SCB4. [PR1464375](#)

- On the MPC7, MPC8, and MPC9 line cards, WO packet error and FPC major alarm are observed when reassembling the small fragments. [PR1465490](#)
- Services card might restart due to a race condition when DNS filtering is enabled. [PR1466567](#)
- Few of the DHCP INFORM packets specific to particular VLAN might take the incorrect resolve queue. [PR1467182](#)
- On the MX10003 and MX204 routers, the temperature sensor name for PEMs under the **show chassis environment pem** command is incorrect. [PR1468642](#)
- Receipt of specific packets causes service cards to restart when the DNS filtering is configured. [PR1469188](#)
- Need to fix multiple FreeBSD vulnerabilities. [PR1470693](#)
- The pkid process crashes at **bn_i2c (pval=0x1d, cont=0x0, putype=0xffffcce8, it=0xc8c848b8 < BIGNUM_it>)** at **../..../src/crypto/openssl/crypto/asn1/x_bignum.c:127**. [PR1471878](#)
- If static VLAN demux unit is configured the **inner-range** statement, only the first VLAN from the range is used to send PADI. [PR1472444](#)
- Manually configured ERO on NS controller are lost when PCEP session is bounced. [PR1472825](#)
- Services card might restart when DNS filtering is enabled. [PR1474056](#)
- The Packet Forwarding Engine crashes on MPC7, MPC8, and MPC9 line cards upon receipt of large packets requiring fragmentation. [PR1474154](#)
- On the MX104 devices, the chassisd process generates core file when cmerror is triggered. [PR1475396](#)
- The **Power Supply failed** and **Power Supply OK** SNMP traps are not generated. [PR1479133](#)
- The SCBE3 fabric plane goes to the **Check** state in MX Series Virtual Chassis. [PR1479363](#)
- 100G interface might randomly fail to come up after the maintenance operations. [PR1481054](#)
- Traffic impact might be observed when policy-multipath is configured without LDP on the SPRING-TE scenario. [PR1483585](#)
- Kernel crashes (vmcore) upon receipt of a malformed IPv6 packet. [PR1486948](#)
- Incorrect frame length of 132 bytes might be captured in the packet header. [PR1487876](#)

Infrastructure

- FPC might reboot if jlock hog occurs on all Junos VM-based platforms. [PR1439906](#)
- The duplex status of management interface might not be updated in the output of the **show** command. [PR1427233](#)
- The scheduled tasks might not be executed if the **cron** daemon goes down without restart automatically. [PR1463802](#)

- The l2ald and eventd processes hogs hundred percent after issuing the **clear ethernet-switching table** command. [PR1452738](#)
- Memory leak leads to kernel crash (vmcore) due to SNMP polling. [PR1482379](#)

Interfaces and Chassis

- When the logical interface is associated to a routing instance inside, an LR is removed from the routing instance and the logical interface is not added to the default routing instance. [PR1444131](#)
- After performing an interface flap, Cordoba related interface goes down. [PR1475777](#)
- Commit error is not thrown when member link was added to the multiple aggregation group with different interface specific options. [PR1475634](#)
- Continuous VRRP state transition (VRRP master/backup flaps) are seen when one device drops VRRP packets. [PR1446390](#)
- The voltage high alarm might not be cleared when the voltage level comes back to normal for MIC on MPC5. [PR1467712](#)
- When you configure ESI on a physical interface, the traffic drops if you disable the logical interface under the physical interface. [PR1467855](#)
- Interface descriptions might be missing under the **logical-systems** configuration. [PR1449673](#)
- Mismatched MTU value causes the RLT interface to flap. [PR1457460](#)
- Executing commit might become nonresponsive as the dcd process becomes nonresponsive. [PR1470622](#)
- MC-AE interface shown as unknown status when you add the subinterface as a part of the VLAN on the peer MC-AE node. [PR1479012](#)
- VRRP failover delay timer starts with configured interval instead of the default interval. [PR1452748](#)
- The MC-LAG configuration-consistency ICL-configuration might fail after committing some changes. [PR1459201](#)
- The EOAM CFM primary-vid functionality does not work if the **enhanced-cfm-mode** is enabled. [PR1465608](#)
- The Virtual Chassis de-couples and couples during the authd daemon restart. [PR1467298](#)

Junos fusion enterprise

- Loop detection might not work on the extended ports. [PR1460209](#)

J-Web

- The httpd process might run with a high CPU utilization when J-Web is enabled. [PR1483607](#)

Layer 2 Ethernet Services

- The jdncpd process might go into an infinite loop and utilize full CPU. [PR1442222](#)
- The DHCP subscriber might not come online after rebooting the router. [PR1458150](#)

- The metric does not change when configured under the DHCP. [PR1461571](#)
- The ISSU might fail when a subscriber inflight login is happening. [PR1465964](#)
- Telemetry data for **relay/bindings/binding-state-v4relay-binding** and **relay/bindings/binding-state-v4relay-bound** is not correct. [PR1475248](#)
- When a DHCPv6 Relay-Agent is configured upon receipt of a specific DHCPv6 client message, remote code execution might occur. [PR1461448](#)

MPLS

- The FPC might get stuck in the Ready state after applying a configuration change that removes the RSVP and triggers an FPC restart. [PR1359087](#)
- Kernel crash and device restart might occur. [PR1478806](#)
- The rpd process might crash during shutdown. [PR1471191](#)
- On all the Junos OS platforms with BGP PIC configured, if committing some operations where RSVP ingress routes are affected, the rpd process might. [PR1471281](#)
- When changing the protocol LDP preference, all the LDP adjacencies would be bounced, that results in all the LDP targeted sessions to flap. [PR1459301](#)
- Previous configured credibility preference is not considered by CSPF, even though the configuration is deleted or changed to prefer another protocol in TED. [PR1460283](#)
- The device might use the **local-computed** path for the PCE-controlled LSPs after the link or node fails. [PR1465902](#)
- RSVP LSPs might not come up in a scaled network with very high number of LSPs if NSR is used on a transit router. [PR1476773](#)
- RPD utilizes full CPU load and generates core files on the backup Routing Engine. [PR1479249](#)
- The traffic might be silently discarded after the LACP timeouts. [PR1452866](#)
- High CPU utilization for the rpd process might be seen if RSVP is implemented. [PR1490163](#)

Platform and Infrastructure

- With CNH enabled, the MPLS CoS rewrite does not work for 6PE traffic. [PR1436872](#)
- JNH memory leak might be seen when the CFM session over the VPLS LSI interface or VT interface flaps. [PR1468663](#)
- Cm errors alarms on certain MPC line cards are classified as major, which should be minor or nonfatal. If these errors are generated, it might lead to a bad hardware condition and then the Packet Forwarding Engine disable action might be triggered. [PR1449427](#)
- When the REST service configuration is modified, (for example, the REST service is configured and then deleted for multiple times) the system might become unresponsive, even to the SSH and console. This issue has service impact. [PR1461021](#)

- FPC might reboot with vmcore due to memory leak. [PR1449664](#)
- In an EVPN-VXLAN scenario, sometimes host-generated packets are dropped as they hit **reject route** in the Packet Forwarding Engine. [PR1451559](#)
- The MPC might drop packets after you enable the firewall fast lookup filter. [PR1454257](#)
- The DDoS protection does not stop logging when remote tracing is enabled. [PR1459605](#)
- Need to make the **version-03** CLI configuration flag optional. [PR1462186](#)
- On the MX204 device, GRE with sampling causes the following Packet Forwarding Engine error: **MQSS(0): MALLOC: Underflow error during reference count read - Overflow 1, Underflow 1, HMCIF 0, Address 0x8d62e0** [PR1463718](#)
- EVPN-VXLAN tpye-5 tunnel might not work properly on the MX Series routers. [PR1466602](#)
- On the MX Series Virtual Chassis setup, the Layer 2 traffic over aggregated Ethernet interfaces sent from one member to another is corrupted. [PR1467764](#)
- The switch might not be able to learn MAC address with **dot1x** and **interface-mac-limit** configured. [PR1470424](#)
- Kernel crash (vmcore) or FPC crash is observed due to mbuf leak. [PR1468183](#)
- Memory leaks in the Packet Forwarding Engine due to the flapping of the 802.1X authenticator port interface. [PR1480706](#)

Routing Policy and Firewall Filters

- The rib-group might not process the exported route correctly. [PR1450123](#)
- Routes resolution might be inconsistent if any route resolving over the multipath route. [PR1453439](#)

Routing Protocols

- Route churn might be seen after changing the **maximum-prefixes** configuration from value A to value B. [PR1423647](#)
- If the same neighbor is configured under different RIP groups, the commit check fails to capture this invalid configuration and commit can be done successfully. However, the rpd process crashes. [PR1485009](#)
- High CPU utilization might be observed when the outgoing BGP updates are sent slowly. [PR1487691](#)
- The rpd might crash after configuring **independent-domain** under the master routing-instance. [PR1469317](#)
- The rpd process might crash with the BGP multipath and route withdrawing occasionally. [PR1481589](#)
- The rpd process crash might be observed if the multipath BGP routes leak from one routing instance to another routing table. [PR1437837](#)
- The CPU utilization on rpd spins at 100 percent after the same external BGP route is learned in different VRF tables. [PR1442902](#)
- The rpd might crash with SRTE configuration change. [PR1442952](#)

- The rpd might crash after configuring the OSPF NSSA area-range and summaries. [PR1444728](#)
- The BGP routes might fail to be installed in the routing instance if the **from next-hop** policy match condition is used in the VRF import policy. [PR1449458](#)
- TI-LFA backup paths for adjacency SIDs is broken in OSPF. [PR1452118](#)
- SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. [PR1454177](#)
- The rpd scheduler slip for BGP GR might be up to 120 seconds after the peer goes down. [PR1454198](#)
- MoFRR with MLDP inband signal does not work. [PR1454199](#)
- The rpd memory might leak in a certain MSDP scenario. [PR1454244](#)
- The rpd process might crash when multipath is in use. [PR1454951](#)
- Prefix SID conflict might be observed in an IS-IS setup. [PR1455994](#)
- The rpd might crash when the OSPF router ID gets changed for NSSA with **area-range** configured. [PR1459080](#)
- The rpd memory leak might be observed on the backup Routing Engine due to BGP flap. [PR1459384](#)
- The rpd scheduler slips might be seen on an RPKI route validation-enabled BGP peering router in a scaled setup. [PR1461602](#)
- Install all possible next hops for OSPF network LSAs. [PR1463535](#)
- The IS-IS IPv6 routes might flap when there is an unrelated commit under the **protocol** configuration. [PR1463650](#)
- BGP peers might flap if the parameter of hold-time sets is small. [PR1466709](#)
- The BFD client session might flap when you remove the BFD configuration from the peer end (from other vendor) of the BFD session. [PR1470603](#)
- The rpd might crash when both the **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 a cluster from a BGP group might cause prolonged convergence time. [PR1473351](#)
- With IS-IS configured and in a very rare case, memory corruption might occur that might cause the rpd process to crash continuously. [PR1455432](#)
- The rpd leaks RPD memory in RT_NEXTHOPS_TEMPLATE. [PR1463112](#)
- The rpd process crashes due to specific BGP update packets. [PR1481641](#)
- The rpd process might be crashed after the BGP peer flaps. [PR1482551](#)

Services Applications

- The jl2tpd process might crash during the restart procedure. [PR1461335](#)

- MX Series L2TP fails to forward the **agentCircuitId** and **agentRemoteId** AVP toward the LNS. [PR1472775](#)
- The kmd might crash due to an incorrect IKE SA establishment after the remote peer's NAT mapping address is changed. [PR1477181](#)

Subscriber Access Management

- DHCPv6 subscribers might get stuck in a state after the authd process crashes. [PR1460578](#)
- The subscriber address allocation might fail after deleting the pool link in the middle of the chain. [PR1465253](#)
- No volume attribute is in the Accounting Stop for Service session when the Activated Services session is configured. [PR1470434](#)
- The subinterfaces might be missing in the NAS port ID. [PR1472045](#)
- The authd might crash after ISSU from releases before Junos OS Release 18.3 to releases after Junos OS Release 18.4. [PR1473159](#)
- Some address-relevant fields are missing when you execute the **test aaa ppp** command. [PR1474180](#)
- The CoA request might not be processed if it includes the **proxy-state** attribute. [PR1479697](#)
- The **mac-address** CLI option are not hidden under the **access profile radius options calling-station-id-format** statement. [PR1480119](#)
- NAS-Port-ID includes sub-interface in the Radius messages for the aggregated Ethernet interface. [PR1484351](#)

VPNs

- The P1 configuration delete message is not sent on loading the baseline configuration if there has been a prior change in the VPN configuration. [PR1432434](#)
- The Layer 2 circuit connections might get stuck in the **OL** state after changing the Layer 2 circuit community and flapping the primary LSP path. [PR1464194](#)
- The rpd process might crash due to memory leak in the **MVPN RPF Src PE** block. [PR1460625](#)
- The rpd process generates core file at **rtbit_reset** and **rte_tgtexport_rth**. [PR1379621](#)
- In an MVPN scenario with ingress replication selective provider tunnel used, if the **link-protection** command is added or deleted from the LSP for MVPN, the rpd process might crash. [PR1469028](#)

Resolved Issues: 19.1R2

Class of Service (CoS)

- Traffic drop occurs when deleting MPLS family or disabling interface that has non-default EXP rewrite-rules. [PR1408817](#)
- The host-inbound packets might be dropped if configuring host-outbound FC. [PR1428144](#)

- The dfwd crash can be seen with forwarding-class configuration in policers. [PR1436894](#)

EVPN

- The RA packets may be sent out without using the configured virtual gateway address. [PR1384574](#)
- [EVPN/VXLAN] VTEP tunnel doesn't get deleted when EVPN peer goes down. [PR1390965](#)
- The process rpd crash might be observed with EVPN type-3 route. churn [PR1394803](#)
- Traffic drop might be seen when the core-facing link comes up in EVPN-VXLAN scenario. [PR1408840](#)
- The next hop is not cleaned up properly when one of the multihomed CE-PE links goes down. [PR1412051](#)
- EVPN-MPLS Single Active: [EVPN/7] /32 host route always appears on non-default PE device if CNH is ON, **remote-ip-host-routes** has no effect [PR1419466](#)
- Rpd crash occurs on backup Routing-Engine after you enable nonstop-routing with EVPN. [PR1425687](#)
- The device might proxy the ARP Probe packets in an EVPN environment [PR1427109](#)
- The CE interface IP address is missed in mac-ip-table of the EVPN database [PR1428581](#)
- Incorrect MAC count is seen with **show evpn/bridge statistics**. [PR1432293](#)
- Stale MAC addresses are present in the bridge mac-table in EVPN/MPLS scenario. [PR1432702](#)
- Asynchronous state between ARP table and Ethernet switching table happens if EVPN ESI link flaps multiple times [PR1435306](#)
- IRB logical interface is not up when local L2 member is down and IM NH is present. [PR1436207](#)
- Configuring ESI on a single-homed 25G port might not work [PR1438227](#)
- The specific source-ports of UDP packet are dropped on EVPN/VXLAN setup [PR1441047](#)
- Restarting l2-learning might cause some remote MAC addresses to move into forwarding **dead** state. [PR1441565](#)
- Traffic drop might be seen in EVPN Layer 3 gateway scenario. [PR1442319](#)
- The core-isolation feature does not work after you set or delete the **no-core-isolation** command on MX Series Devices. [PR1442973](#)
- The EVPN type 2 routes might not be advertised properly in logical systems. [PR1443798](#)
- The localhost 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 and this results in packet loss [PR1455973](#)
- Delay factor might send back ARP request/NS to local segment under EVPN-ETREE leaf role conditions. [PR1459830](#)

- In EVPN scenario memory Leak might be observed when **proxy-macip-advertisement** is configured. [PR1461677](#)
- Traffic received from vtep gets dropped if the VNI value used for type-5 routes is greater than 65535. [PR1461860](#)

Forwarding and Sampling

- In some later releases firewall filter action **decapsulate gre** cannot decapsulate IP-over-IP and IPv6-over-IP traffic. [PR1398888](#)
- The SRRD might crash when memory corruption occurs. [PR1414568](#)
- EVPN enhancement for MAC flush mechanism is needed in Junos OS. [PR1421018](#)
- Junos 19.1: Firewall filter and policers not working correctly. [PR1424183](#)
- rt-delay-threshold can be set below 1 second - but rt-marker-interval is limited to 1 second. [PR1425544](#)
- The device is in amnesiac mode after ISSU with **mgd: error: configuration check-out failed** error generated. [PR1432664](#)
- Enable interface with input/output vlan-maps to be added to a routing instance configured with a vlan-id/vlan-tags (instance type virtual-switch/vpls) [PR1433542](#)
- The high CPU utilization of l2ald is seen after replacing EVPN configuration. [PR1446568](#)
- [MX204] input/output counters of AE bundle/member links configured on non-default logical systems are not updated. [PR1446762](#)
- ARP packets are getting dropped by Packet Forwarding Engine after you restart chassis-control in MX Series devices. [PR1450928](#)
- Commit error and dfwd core files might be observed when applying a firewall filter with action **then traffic-class** or **then dscp**. [PR1452435](#)

General Routing

- MX Series Virtual Chassis: suboptimal aggregate Ethernet load balancing occurs when an Aggregate Ethernet bundle is part of an ECMP path. [PR1255542](#)
- BGP IPv4 PIC: Packet Forwarding Engine selector gets stuck in rerouted state on unilist next hop after primary AE link is activated or deactivated. [PR1354786](#)
- Traffic might be blocked on MX Series device with MS-MPC/MS-MIC. [PR1358019](#)
- Interface with Tri Rate Copper SFP(P/N:740-01311) in "MIC 3D 20x 1GE(LAN)-E,SFP" will stop forwarding traffic after ISSU upgrade. [PR1379398](#)
- FPC errors might be seen in subscriber scenario. [PR1380566](#)
- The unicast traffic from IRB interface toward LSI might be dropped due to Packet Forwarding Engine mismatching at egress processing. [PR1381580](#)

- Interface filter statistics are not showing the input packet count/rejects, and **show pfe statistics traffic** does not report any normal discard. [PR1383579](#)
- Subscriber connection setup is 30 percent lower than expected. [PR1384722](#)
- The rpd process might end up with a stuck krt queue in VRF scenario. [PR1386475](#)
- Migrate from syslog API to Errmsg API - VMhost messages seen in Junos OS. [PR1387099](#)
- BBE SMGD generates core files if MTU is changed while subscribers are logged in on the physical interface. [PR1389611](#)
- The **high-cos-queue-threshold** range is changed [uint 0 .. 90;]. [PR1390424](#)
- The BNG might not respond with PADO and create any demux interface when PPPoE PADI packet is received. [PR1390989](#)
- FPC might reboot on VMX in subscriber scenario. [PR1393660](#)
- The FPC cards might not come up while performing ISSU on MX10003. [PR1393940](#)
- IDS aggregate configuration command should not be considered for the installation of the IDS dynamic filter. [PR1395316](#)
- Layer 3 gateway did not update ARP entries if IP or MAC addresses quickly move from one router to another router in EVPN-VXLAN environment [PR1395685](#)
- **VMHost RE 0 Secure BIOS Version Mismatch** and **VMHost RE 1 Secure Boot Disabled** alarms are seen. [PR1397030](#)
- The service PIC might crash while changing CGNAT mode. [PR1397294](#)
- The PPPoE subscribers are unable to reconnect after FPC reboot. [PR1397628](#)
- Confirmation message is missing when issuing **request vmhost reboot re**. [PR1397912](#)
- The CLI command **show system firmware** gets hidden on MX Series platforms. [PR1398022](#)
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. [PR1399683](#)
- The na-grpcd log file is not rotated and keeps growing until Routing Engine is out of disk space. [PR1401817](#)
- Continuous kernel crashes might be observed in backup Routing Engine or Virtual Chassis backup router. [PR1404038](#)
- With MS-MPC and MS-MIC service cards SYSLOG messages for port block interim might show 0.0.0.0 for the private IP addresses and PBA release messages might show IP that has undergone NAT as the private IP. [PR1404089](#)
- Incorrect display of assigned prefixes to a subscriber in the output of **show interface < dynamic demux interface>**. [PR1404369](#)
- Voltage read failed for rail LTC3887-EA1-VDD0V9R2-CH0. [PR1405787](#)

- FPC crash might be seen when adding a leg to an AE bundle or FPC restarts in subscriber scenario. [PR1405876](#)
- The FPC crash might be observed in MS-MPC high availability environment. [PR1405917](#)
- The rpd might crash due to a race condition with the combination of community actions done at both a BGP import policy and a forwarding-table policy [PR1406357](#)
- Change the default parameters for resource-monitor rtt-parameters. [PR1407021](#)
- FPC might crash during the subscriber-related stress tests. [PR1407285](#)
- The rpd might crash when a commit check is executed on LDP trace options filtering. [PR1407367](#)
- FPC crash and slow convergence upon HMC Fatal error condition when inline J-Flow is used. [PR1407506](#)
- Openconfig-network-instance:network-instances support for IS-IS must be hidden unless supported. [PR1408151](#)
- The ToS/DSCP and TTL fields might not be copied into the outer IP header in Group VPN scenario. [PR1408168](#)
- The misconfiguration of dynamic profile might cause the login issues of the subsequent subscribers. [PR1409398](#)
- MX-MPC2-3D-EQ and MPC-3D-16XGE-SFPP will now show "Exhaust A" temperature, rather than Intake temperature. [PR1409406](#)
- The CPU might be hogged by jsd process in JET scenario [PR1409639](#)
- Nonexistent subscribers might appear in **show system resource-monitor subscribers-limit chassis extensive** output. [PR1409767](#)
- Packets might be dropped if the traffic is forwarded through an LT interface. [PR1410970](#)
- Slow SNMP on entityMIB during subscribers load test. [PR1411062](#)
- A steady increase of the Packet Forwarding Engine heap memory utilization might happen when PPPoE subscribers are flapping. [PR1411389](#)
- Parity error might cause FPC alarm. [PR1411610](#)
- The **file copy** command might not work if the **routing-instance** option is not specified. [PR1412033](#)
- The spfe on satellite device in Junos fusion setup might crash and it could cause the satellite device to get offline. [PR1412279](#)
- Junos OS PCC might reject PCUpdate/PCCreate message if there is a metric type other than type 2. [PR1412659](#)
- PPPoE subscribers might not be able to log in after unified ISSU. [PR1413004](#)
- The rpd memory leak might be seen due to the incorrect processing of a transient event. [PR1413224](#)
- Need to reduce max flow table size when using flex-flow-sizing. [PR1413513](#)
- DHCP subscribers over HAG can result in core file generation. [PR1413862](#)

- The services load balance might not be effective for AMS if the hash key under the forwarding-options hierarchy is configured. [PR1414109](#)
- FPC crash might be observed if it reaches heap utilization limit. [PR1414145](#)
- DHCP/DHCPv6 subscribers might fail to establish sessions on PowerPC based MX Series platforms. [PR1414333](#)
- Anomaly in LED behavior occurs after rebooting the directly connected device. [PR1414532](#)
- NPC might not apply configured resource-monitor thresholds after NPC. restart [PR1414650](#)
- Firewall filters are not getting programmed into Packet Forwarding Engine. [PR1414706](#)
- The user might not enter configure mode because of mgd is in lockf status. [PR1415042](#)
- **ICMP MTU exceeded** error generated from Packet Forwarding Engine does not reach the expected source. [PR1415130](#)
- Port speed change and scaled AE configuration can lead to MQSS errors and subsequent card crash. [PR1415183](#)
- MTU issue might cause PS interface to flap during dcd restart or GRES switchover [PR1415207](#)
- PCE-initiated LSPs get deleted because of incorrect timer timeout. [PR1415224](#)
- The IRB interface might flap after committing a configuration change on any interface [PR1415284](#)
- Jdhcpd core files is observed after the active lease-query configurations are deleted. [PR1415990](#)
- BMP type 1 message with extra 24 bytes occurs at end of the message. [PR1416301](#)
- Some IPsec tunnels might fail to pass traffic after GRES on MX Series platform. [PR1417170](#)
- The ECMP fast reroute protection feature might not work on MX5, MX10, MX40, MX80, and MX104. [PR1417186](#)
- The IPv6 neighbor might become unreachable after the primary link goes down in VPLS multihoming scenario [PR1417209](#)
- An IPv4 packet with a zero checksum might not be translated to IPv6 packet properly in a NAT64 scenario. [PR1417215](#)
- An invalid XML reply containing a duplicate tag might be seen when requesting **get-arp-table-information** through NETCONF. [PR1417269](#)
- The JSU package installation might fail. [PR1417345](#)
- Some subscribers might be offline when doing GRES or daemon restart [PR1417574](#)
- Zero tunnel statistics are shown for the soft-gre tunnel. [PR1417666](#)
- The BGP session might flap after Routing Engine switchover. [PR1417966](#)
- CGNAT with MS-MPC card doesn't account for AP-P out of port errors or generate a syslog message when this condition is met [PR1418128](#)

- There is no SNMP trap message generated for jnxHardDiskMissing/jnxHardDiskFailed on Summit MX [PR1418461](#)
- Adding two or more ps interfaces might cause traffic drop in l2circuit scenario [PR1418610](#)
- The lsp-cleanup-timer is not being honored when it is configured to be greater than 2,147,483,647. [PR1418937](#)
- The PPPoE negotiation of subscriber connection might fail when 65535 is assigned as session-id [PR1418960](#)
- RX alarms are not set as according to the threshold value configured for the DCO Tunable Optics. [PR1419204](#)
- A PPP session under negotiation might be terminated if another PPPoE client bears the same session ID. [PR1419500](#)
- CPU usage on service PIC might spike while forming an IPsec tunnel in a DEP/NAT-T scenario. [PR1419541](#)
- A new tunnel could not be established after changing the NAT mapping IP address until the IPsec SA clear command is run. [PR1419542](#)
- The message **rtsock_peer_unconsumed_obj_free_int: unable to remove node from list** is logged extensively. [PR1419647](#)
- bbe-mibd memory leak causes the daemon to crash during live subscribers and SNMP OIDs query. [PR1419756](#)
- The IPsec tunnel might get down when the platforms running Junos OS and the peer both act as the initiator and try to bring an IPsec tunnel up at the same time. [PR1420293](#)
- The **show chassis power** output status doesn't seem right and there are also similar error messages in the syslog after you turn power off or on. [PR1420571](#)
- SPC3 Storage and hard disc error log messages. [PR1420800](#)
- PTP phase is aligned but TE and cTE are not good. [PR1420809](#)
- The FPC CPU might be hogged if channelized interfaces are configured. [PR1420983](#)
- MX LNS might fail to forward the traffic on the subscriber access route. [PR1421314](#)
- Failed to reload keyadmin database for **/var/etc/keyadmin.conf**. [PR1421539](#)
- **Bbemg_smgd_lock_cli_instance_db** should not be logged as error messages. [PR1421589](#)
- VCP port reports MTU value 9152 in the ICMP MTU exceeded message while the VCP port MTU is set to 9148. [PR1421629](#)
- The ps access interface is not marked as ccc down on standby/nondesignated PE devices. [PR1421648](#)
- After a control plane event, a few ipsec tunnels failed to send traffic through the tunnel. [PR1421843](#)
- RPM syslogs are not getting generated after deactivating aggregate interface. [PR1421934](#)

- The changed value of **remote-gateway** does not take effect when the router acts as an initiator of IPsec-VPN tunnel. [PR1421977](#)
- RSI bloat occurs due to vmhost-based log collection. [PR1422354](#)
- Packet Forwarding Engine wedge might be observed after performing the command **show forwarding-options load-balance** [PR1422464](#)
- The XML output might be not hierarchically structured if you issue the **show security group-vpn member ipsec statistics** command. [PR1422496](#)
- The CoS IEEE-802.1 classifier might not get applied when it is configured with service activation on underlying-interface [PR1422542](#)
- Incorrect burst-size is seen when the traffic-control-profile is applied to a ps (or pseudowire) interface. This causes unexpected behavior and the queues are not able to process the expected traffic. [PR1422549](#)
- The allocation of MAC address might fall out of the MAC address pool on MX204 platform [PR1422679](#)
- SFP-T/SX/LX is not working with QSA adapter in on MX10003. [PR1422808](#)
- The **show system subscriber-management summary** command should include a failure reason for standby disconnect when primary and backup Routing Engine memories mismatch. [PR1422976](#)
- A stuck lock in shared memory might prevent subscribers from logging in again after daemon crash. [PR1424607](#)
- Incorrect PIC mode on MX204 MX1RU when PIC mode is changed to default mode. [PR1423215](#)
- While committing a huge configuration customer is seeing the error **error: mustd trace init failed.** [PR1423229](#)
- The **set forwarding-options enhanced-hash-key symmetric** is not effective on MX10003. [PR1423288](#)
- IP packet drop might be seen in Layer 2 circuit scenario. [PR1423628](#)
- Traffic is dropped after FPC reboot with AE member links deactivated by remote device. [PR1423707](#)
- MPC10E-15C-MRATE: crash seen at **Ktree alloc (jnh_dfw_instance_add (filter_index=< optimized out>) at ../../../../src/pfe/common/applications/dfw/dfw_iff.c:1030** with inline and scale prefix filter. [PR1423709](#)
- On MX204 Optics 'SFP-1GE-FE-E-T' I2C read errors are seen when an SFP-T is inserted into a disabled state port. [PR1423858](#)
- A PTP asymmetry change needs PTP bouncing. [PR1423860](#)
- The bbe-smgd process might crash after executing the command **show system subscriber-management route prefix <>**. [PR1424054](#)
- The port configured for 1-Gbps speed flaps after Routing Engine switchover. [PR1424120](#)
- The interface configured with 1G speed on JNP10K-LC2101 cannot come up [PR1424125](#)

- The system does not reboot , even though **disk-failure-action reboot** or **disk-failure-action halt** is configured. [PR1424187](#)
- Continuous disk error logs appear on VCP virtual console (requesting switchover due to disk failure on ada1). [PR1424771](#)
- The rpd keeps crashing after changing configuration [PR1424819](#)
- The jdhcpd might consume 100 percent CPU and crash if **dhcp-security** is configured. [PR1425206](#)
- Interface with FEC disabled might flap after Routing Engine mastership switchover. [PR1425211](#)
- The rpd will crash continuously if MD5 authentication on any protocols is used along with master-password [PR1425231](#)
- Soft GRE tunnel route is lost after reboot/GRES or upgrade in WAG scenario [PR1425237](#)
- The mspmand process might crash and restart with a mspmand core file created after doing a commit change to deactivate and activate service-set [PR1425405](#)
- The following log message is seen continuously on MX204 router: **fru_is_present: out of range slot 0 for.** [PR1425411](#)
- Getting Unisphere-UpStream-Calc-Rate as 0 while verifying L2BSA RADIUS accounting stop packets after performing GRES. [PR1425512](#)
- All interfaces creation failed after NSSU. [PR1425716](#)
- MPC reboot or Routing Engine mastership switchover might occur on MX204/MX10003. [PR1426120](#)
- Logical Interfaces Targeting: 18,000 phantom distributed interfaces are displayed for AE interface with the targeted distribution enabled on it, when there are no active subscribers. [PR1426157](#)
- Interfaces might go come to down after device reboots. [PR1426349](#)
- PEMs lose DC output power load sharing after PEM switch off and on operation on MX Series platforms. [PR1426350](#)
- Some CFM and BFD sessions might flap while collecting MPLS statistics [PR1426727](#)
- The **show lldp neighbors interface** command does not display all interface information. [PR1426793](#)
- The decoding of telemetry data at collector might not be correct if you configure. [PR1426871](#)
- Traffic loss might be seen when multiple IPsec tunnels are established with the remote peer. [PR1426975](#)
- Traffic might not flow through MACsec interface even after an unsupported cipher-suite is removed. [PR1427294](#)
- ENTITY MIB has incorrect containedIn values for some fixed MPCs with built-in PICs. [PR1427305](#)
- Rebooting or halting Virtual Chassis member might cause 30 seconds of down time on RTG link. [PR1427500](#)
- When broadband edge PPPoE and DHCP subscribers coming up over Junos fusion satellite ports are active, **commit full** and **commit synchronizaton full** commands fail. [PR1427647](#)

- When installing YANG package without **proxy-xml** command, the CLI environment might not work well. [PR1427726](#)
- The subscriber IP route might got stuck in bbe-smgd if the subscriber IP address is the same as the local IP address. [PR1428428](#)
- In correct display of MAC/MAC+IP and count values occurs, after setting **global-mac-limit** and **global-mac-ip-limit**. [PR1428572](#)
- The PTSP subscribers are stuck in the configured state after being rejected by the RADIUS server. [PR1428688](#)
- Incorrect IGMP statistics are seen for dynamic PPP interfaces. [PR1428822](#)
- Fabric drops might be seen on MX10003 platform when two FPCs come online together. [PR1428854](#)
- Incorrect IGMP interface counter for dynamic PPP interfaces. [PR1429018](#)
- The emitted XML is INVALID error is thrown for **show virtual-network-functions**. [PR1429090](#)
- L2TP subscriber and MPLS pseudowire subscriber volume accounting statistics value remains unchanged post unified ISSU. [PR1429692](#)
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs. [PR1429821](#)
- The AE interface does not come up after rebooting the FPC/device even though the physical member link is up. [PR1429917](#)
- Configuration is prevented from being applied on MX Series device in subscriber scenario. [PR1430360](#)
- Performance degradation for about 20 seconds occurs after the fabric board on MX10008/100016 is taken offline. [PR1430739](#)
- Disabling DAC QSFP port might not work on MX204/10003 or EX9251. [PR1430921](#)
- Inline LSQ might not work when it is configured on the same FPC where MIC-3D-16CHE1-T1 is slotted. [PR1431069](#)
- Error might be observed when using a script to load configuration. [PR1431198](#)
- The destination unreachable counter was counting up without receiving traffic. [PR1431384](#)
- During the stress tests, bbe-smgd process might crash on backup Routing Engine when performing GRES. [PR1431455](#)
- The bbe-smgd might crash if subscribers are trying to log in or out while a configuration commit is occurring at the same time. [PR1431459](#)
- Subscribers coming from new physical interfaces might not log in in due to 512 entries limit in the subscriber-limit table [PR1431566](#)
- **SIB Link Error** detected on a specific Packet Forwarding Engine might cause complete service impact. [PR1431592](#)

- Allow installation of three identical framed routes in the same routing-instance. [PR1431891](#)
- MX10003 - A PEM not present alarm is raised when the minimum required PEM exists in the system. [PR1431926](#)
- Dual stack subscriber accounting statistics are not baselined when one stack logs out. [PR1432163](#)
- Traffic might be sent on the standby link of AE bundle and get lost with LACP fast-failover enabled [PR1432449](#)
- Change to in-use parameterized filter prefix-list could result in bbe-smgd core file on the backup Routing Engine. [PR1432655](#)
- Traffic will be dropped if **sa-multicast** is in the configuration. [PR1433306](#)
- RSI and RSI brief should not include **show route forwarding-table** when Tomcat enabled. [PR1433440](#)
- Collected service statistics are all zero after ISSU for MPC2. [PR1433589](#)
- Lawful intercept for subscriber traffic is not programmed in Packet Forwarding Engine if it's activated by Access-Accept. [PR1433911](#)
- URL case-sensitivity support is needed. [PR1434004](#)
- Incorrect PLUGGABLE ID 17 on MX10003-LC2103. [PR1434183](#)
- RPD generates a core file during the route flash when the policy is removed. [PR1434243](#)
- The repd process might crash after booting first time with a newly installed Junos release. [PR1434363](#)
- Packet Forwarding Engine memory leak might be seen if MLPPP links are flapped. [PR1434980](#)
- MicroBFD 3x100ms flap upon inserting a QSFP to other port. [PR1435221](#)
- DHCPv6 advertise to client might use incorrect destination MAC address. [PR1435694](#)
- Total number of packets mirrored, after DTCP trigger add and DTCP enable, is not in the expected range while verifying traffic on mirror port after DTCP drop policy enable. [PR1435736](#)
- MPC7/8/9/MX10003 MPC/EX9200-12QS/EX9200-40XS line card might crash in a scaling setup. [PR1435744](#)
- The mc-ae interface might get stuck in waiting state in dual mc-ae scenario. [PR1435874](#)
- The local route in the secondary routing table that gets stuck in the KRT. [PR1436080](#)
- ifHCInOctets counter on AE interface going to zero value when SNMP MIB walk is executed. [PR1436201](#)
- A few static PPP subscribers get stuck in initialization state permanently and the following error message is seen **Failed to create client session, err=SDB data corrupted**. [PR1436350](#)
- Subscriber interim statistics might be reset to zero in MX Series Virtual Chassis setup after GRES. [PR1436419](#)
- Router is not reachable after downgrade from Junos OS Release 18.2-20190513.0 to 18.2R2.6. [PR1436832](#)

- MPC10E-15C-MRATE: Micro BFD sessions do not come up in centralized mode. [PR1436937](#)
- Schema XSDs are missing objects/commands in Junos OS Release 19.1R1. [PR1437469](#)
- The CPU utilization on a daemon might remain around 100 percent or the backup Routing Engine might crash in race conditions. [PR1437762](#)
- LNS router might send the router-advertisement packet with NULL source link-layer option field. [PR1437847](#)
- The chassisd might crash after enabling hash-key. [PR1437855](#)
- (seen only on legacy image) Unified ISSU is failing from Junos OS Release 19.1R1 legacy images. [PR1438144](#)
- Subscriber flows might not be synchronized between AE members on MX Series Virtual Chassis platforms. [PR1438621](#)
- CGNAT logs are not received by the syslog server over TCP-based-syslog when data traffic is sent at 10000 sessions/second. [PR1438928](#)
- Command **show jdaf service cmd statistics / clients** is not available on Evolved. [PR1439118](#)
- FPC on Virtual Chassis backup router might reboot in MX Series Virtual Chassis scenario. [PR1439170](#)
- The **vlan all interface all** combination is not working as expected under VSTP. [PR1439583](#)
- The bbe-smgd core files are generated after restart. [PR1439905](#)
- CoS related errors are seen and subscribers could not get service. [PR1440381](#)
- CPU might hang or interface might be stuck down on particular 100G port on MX, EX, and PTX. [PR1440526](#)
- FPC may stuck in 100% CPU utilization due to continuous work of bulk manager thread. [PR1440676](#)
- DHCP offer packets towards IRB over LT interface getting dropped in DHCP relay environment. [PR1440696](#)
- The Layer 2 dynamic VLAN might be missed when an interface is added or removed for an ae interface. [PR1440872](#)
- For a route received through EBGp the AIGP value might not be considered as expected. [PR1441438](#)
- LINX: SNMP trap comes twice for FRU removal in MX10000- 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 be dropped in VPLS scenario. [PR1442121](#)
- The chassisd is unable to power off a faulty FPC after Routing Engine switchover which 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 "enhanced-ip" or "enhanced-ethernet" mode with DCU (destination-class-usage) accounting enabled, MS-DPC might drop all traffic that should egress through ae 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. Now severity of those logs are corrected as INFO. [PR1442598](#)
- DHCPv6 Client might fail to get an IP address. [PR1442867](#)
- The bbe-smgd might crash on MX Series platforms. [PR1443109](#)
- 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: Larger than MTU packets of GRE data get dropped when sampling is enabled on the egress interface. [PR1444186](#)
- High CPU utilization might be observed for **eventd** along with error logs. [PR1444462](#)
- Inline-keepalive might stop working for LNS subscribers if the **routing-services** command is enabled. [PR1444696](#)
- Access route might stuck in bbe-smgd and RPD not cleared. [PR1445155](#)
- The CPCDD process continuously generates core files and the process stops, 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 because of device reboot. [PR1445428](#)
- The 1G interface on MX204 might stay down after the device is rebooted. [PR1445508](#)
- The l2ald might crash when FPC is restarted. [PR1445720](#)
- The mspmand process might crash if URL filtering is configured and one blacklisted domain name is a sub-string of another blacklisted domain name in URL filter database file. [PR1445751](#)
- The process jdncpd might crash after issuing the **show access-security router-advertisement-guard** command. [PR1446034](#)
- The static route for NAT might never come up if switchover the service interface which has NAT and GR configuration. [PR1446267](#)
- Accurate statistics might not include packets forwarded during the last two seconds before subscriber termination. [PR1446546](#)
- NAT service-set in certain scale might fail to get programmed. [PR1446931](#)
- CST: **ISSU:core-RMPC3.gz.core.0** is seen and ISSU-failure seen for MPC5. [PR1446993](#)

- The jflow version 5 stops working after changing "input rate" value. [PR1446996](#)
- The **no-control-word** creates a traffic black hole when used with Redundant LT (or rlt interface) for PWHT (or ps0 interface). [PR1447917](#)
- The rpd process might crash if BGP is activated/deactivated multiple times. [PR1448325](#)
- DCD CPU spike is seen after a Junos upgrade from Junos OS Release 14.2 to 16.1. [PR1448858](#)
- PR-1444575-fix-test: FPC rebooted during off-lining PIC-0. [PR1449067](#)
- The DHCP relay feature might not work as expected with **helpers bootp** configured. [PR1449201](#)
- Increase in the maximum value of **delegation-cleanup-timeout**. [PR1449468](#)
- Need to provide more meaningful error message, while doing commit on JDM without exchanging the SSH keys. [PR1449871](#)
- **No localhost ifl for rtt 65535** can be seen on MX Series running junos enhanced subscriber management feature. [PR1450057](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)
- VLAN configuration change with l2ald restart might cause Kernel sync issues and impact forwarding. [PR1450832](#)
- Configuring a new burst-size under traffic-control-profile is not taking effect. [PR1451033](#)
- IPSec[SNMP]: Snmp query for IPSec Decrypted/Encrypted packets does not fetch right values; observing **KMD_SNMP_FATAL_ERROR**. [PR1451324](#)
- RMPC core files are found found after configuration changes done on the network for PTP/Clock Synchronization. [PR1451950](#)
- MX10003: MACsec framing errors are seen when ever sequence number exceed 2 power 32 with XPN (Extended Packet Numbering). [PR1452851](#)
- PTP might go out of sync due to l2ald hwdb access failure. [PR1453531](#)
- Alarm was not sent to syslog on MX10003 platform. [PR1453533](#)
- ANCP subscriber information is lost after daemon restart. [PR1453837](#)
- The FPC might crash when the severity of error is modified. [PR1453871](#)
- Radius Interim accounting statistics are not populated on the MX204. [PR1454541](#)
- The access request for L2BSA port up may not be retransmitted if the radius server used to be unreachable. [PR1454975](#)
- JNS/GNF: CRAFTD syslogs fatal errors along with junk characters upon its startup and exits after four startup attempts. [PR1454985](#)
- 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](#)

- There is high temperature from **show chassis environment** output after MPC4E insert to slot 5. [PR1456457](#)
- The CLI command with **invoke-on** and **display xml rpc** results in unexpected multiple RPC commands. [PR1456578](#)
- 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 subscriber routes are not cleared from backup Routing Engine when session is aborted. [PR1458369](#)
- Subscribers unable to login due to NACK from MCAST after 2million + mcast subscribers log in. [PR1458419](#)
- The error messages with **create_pseudos: unable to create interface device for pip0 (File exists)** might be seen after restarting chassisd. [PR1459373](#)
- Incomplete output of **show ancp subscriber access-aggregation-circuit-id < access aggregation circuit ID>**. [PR1459386](#)
- Telemetry streaming of mandatory TLV 'ttl' learnt from LLDP neighbor is missing. [PR1459441](#)
- FDB is not flushing cause silent drop in traffic in ethernet ring scenario. [PR1459446](#)
- In MC-LAG scenario traffic destined to VRRP virtual MAC gets dropped. [PR1459692](#)
- AUTO-CORE-PR :CPCDD core found @
ServicesMgr::ServicesManager::cpcddSmdInterface::processServiceNotifyMsg
,SmdInterface::cbStateSyncServiceNotifyMsgHandler ,statesync_consumer_poll_new_state_cb.
[PR1459904](#)
- The PPTP does not work with destination NAT. [PR1460027](#)
- repd core file is seen during system boot up. [PR1461796](#)
- The BBE statistics collection and management process, **bbe-statsd** memory issue on backup Routing Engine. [PR1461821](#)
- The **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed** when both DIP switches and power switch are turned off. [PR1462065](#)
- The MPC2E-NG/MPC3E-NG card with specific MIC might crash after a high rate of interface flaps. [PR1463859](#)
- The PPPoE session gets in the **Terminated** state and the accounting stop for the session which is delayed. [PR1464804](#)
- MPC5E or MPC6E line cards might crash due to internal thread overusing the CPU. [PR1464820](#)

Infrastructure

- SNMP OID IFOutDiscards are not updated when drops increasing. [PR1411303](#)
- The traffic to the NLB server might not be forwarded if the NLB cluster works on multicast mode. [PR1411549](#)
- Junos OS: MX Series: An MPC10 Denial of Service (DoS) due to OSPF states transitioning to Down, causes traffic to stop forwarding through the device. [PR1418955](#)
- Increase in Junos OS image size for Junos OS Release 19.1R1. [PR1423139](#)
- The duplex status of management interface might not be updated in the output of **show** command. [PR1427233](#)
- The operations on console might not work if the statement **system ports console log-out-on-disconnect** is configured. [PR1433224](#)

Interfaces and Chassis

- Changing the value of **mac-table-size** to default might lead all FPC to reboot. [PR1386768](#)
- NPC crash @rt_nh_install (rnh=0x618123d8, rnh_src=0x0, rt=< optimized out>, p_rtt=0x74f886c0) at ../../../../src/pfe/common/pfe-arch/trinity/applications/route/rt_nh.c:631. [PR1396540](#)
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. [PR1402606](#)
- Unrelated ae interfaces might go down if committing configuration changes. [PR1409535](#)
- MX Virtual Chassis unified ISSU is not supported when Redundant LT (RLT) is configured. [PR1411729](#)
- Family inet of the unnumbered interface might be getting deleted when deleting one of the IPs of the binding interface. [PR1412534](#)
- Inline Periodic packet management (PPM) adjacency (rx) session might be programmed with the incorrect packet template. [PR1417707](#)
- Monitor ethernet loss-measurement command returns Invalid ETH-LM request for unsupported outgoing IFL. [PR1420514](#)
- Invalid speed value on an interface might cause other interface configuration loss. [PR1421857](#)
- Syslog message : /kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex upon LFM related configuration commit on ae interfaces. [PR1423586](#)
- The demux interfaces will be down after changing the MTU of the underlying et interface. [PR1424770](#)
- The cfmd might crash on DPCE. [PR1424912](#)
- The IFLs in EVPN routing instances might flap after committing configurations. [PR1425339](#)
- The statement **flexible-queuing-mode** is not working on FPCs of Virtual Chassis member 1. [PR1425414](#)
- Upgrade from pre Junos OS Release 17.4R1 release results in cfmd core files. [PR1425804](#)
- CFM message is flooding. [PR1427868](#)

- The vrrpd process might crash after deleting VRRP sessions for several times. [PR1429906](#)
- The NCP session might be brought down after IPCP Configure-Reject is sent. [PR1431038](#)
- VRRP mastership might flap when the tracked route is deleted or the tracked interface goes down. [PR1432361](#)
- jppd No termination Ack for a LACP Termination request RFC 1661. [PR1433489](#)
- Discrepancy of bytes and packets count in Routing Engine CLI for traffic and transit stats for ZT. [PR1435416](#)
- Mixed link-speed ae bundle could not add new sub-interface successfully. [PR1437929](#)
- Targeted-distribution for static demux interface over aggregate ether interface does not take correct lacp link status into consideration when choosing primary and backup links. [PR1439257](#)
- The cfmd process might crash after a restart on Junos OS Release 17.1R1 and above. [PR1443353](#)
- Enhancement of add or delete a single VLAN in **vlan-id-list** under interface family bridge. [PR1443536](#)
- The OAM CCM messages are sent with a single tagged VLAN even when configuring with two VLANs. [PR1445926](#)
- MX Series Virtual Chassis on MX10003: Unable to connect to newly installed Routing Engine from other Routing Engine's in MX Series Virtual Chassis. [PR1446418](#)
- Initiating a Routing Engine switchover on VRRP backup router through a CLI command might cause VRRP state for ae bundle interfaces transitions to 'Master' state even configured with **protocols vrrp delegate-processing ae-irb** statement, then very shortly afterward to backup again. [PR1447028](#)
- L2ALD failed to update composite NH. [PR1447693](#)
- The ifinfo daemon might crash on the execution of **show interface extensive** command. [PR1448090](#)
- While master Routing Engine failure and system switches to backup Routing Engine, some VRRP sessions ppm transmissions state might be stuck in **Distributed: AWAITING**. [PR1450652](#)
- LACP daemon crashed continuously. [PR1450978](#)
- [PDT][CFM] CUC-1751: Some CFM UP MEP sessions do not come up in scaled scenarios over Layer2VPN circuits on Lag interfaces. [PR1454187](#)
- The VRRP traffic loss is longer than one 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 using J-Web. [PR1446081](#)

Junos fusion for provider edge

- Incorrect power values for extended optical ports. [PR1412781](#)

- 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 to it. [PR1441869](#)

Layer 2 Features

- VPLS : LSI interface are not created and remote MACs are not learned with `RPD_KRT_Q_RETRIES: ifl iff add: Device busy`. [PR1295664](#)
- QinQ might be malfunctioning if `vlan-id-lists` are configured. [PR1395312](#)
- The `rpd` crashes after `iw0` interface is configured under a VPLS instance. [PR1406472](#)
- Traffic loss might be seen over LDP-VPLS scenario. [PR1415522](#)
- Broadcast traffics might be discarded in a VPLS local-switching scenario. [PR1416228](#)
- Commit error will be seen but the commit is processed if adding more than o. [PR1420082](#)
- VPLS neighbors might stay in down state after configuration changes in `vlan-id`. [PR1428862](#)
- After disabling and enabling the aggregate interface, the next-hop of CE facing aggregate interface might be in an incorrect state. [PR1436714](#)

Layer 2 Ethernet Services

- LACP PDU might be looped towards peer MC-LAG nodes. [PR1379022](#)
- On EVPN setups, incorrect destination MAC addresses starting with 45 might show up when using the `show arp hostname` command. [PR1392575](#)
- Junos OS: MX Series: `jdhcpd` crash when receiving a specific crafted DHCP response message (CVE-2019-0063). [PR1415242](#)
- `jdhcpd` becomes aware about some of the existing configuration only after `commit full` or `jdhcpd` restart. [PR1419437](#)
- Change the `nd6` nexthops to reject NH once l2 interfaces gets disassociated with ipv6 entries. [PR1419809](#)
- The `jdhcpd` process might consistently run at 100% CPU and not provide service if the `delay-offer` is configured for DHCP local server. [PR1419816](#)
- BBE: DHCP subscribers on non-default routing instance went down after unified ISSU. [PR1420982](#)
- `jdhcpd` daemon might crash during continuous stress test. [PR1421569](#)
- MX: LACP: The error message `fpc3 user.err aftd-trio: [bt] #1 JnhHandle::` been logged. [PR1424106](#)
- The DHCP DECLINE packets are not forwarded to DHCP server when `forward-only` is set within `dhcp-reply`. [PR1429456](#)
- DHCP request might get dropped in DHCP relay scenario. [PR1435039](#)
- The `jdhcpd` process might go into infinite loop and cause 100% CPU usage. [PR1442222](#)
- The `dhcp-relay` statement might not work on MX10008 platforms. [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 packet might not be processed correctly if DHCP option 82 is configured. [PR1459925](#)

MPLS

- Stale LSPs might exist if primary LSP goes down immediately after bypass LSP. [PR1242558](#)
- RPD might restart after a MPLS LSP flap if **no-cspf** and **fast-reroute** are configured in an LSR ingress router. [PR1368177](#)
- DSCP bit Marking of LSP self-ping is not compliant with rfc7746. [PR1371486](#)
- The rpd process might crash when executing **traceroute mpls bgp**. [PR1399484](#)
- LDP tunneling configuration triggers huge scheduler slips, causing IGP flapping. [PR1410827](#)
- The rpd might crash in BGP-LU with egress-protection while committing configuration changes. [PR1412829](#)
- The rpd might crash after **spring-te** is deactivated. [PR1414323](#)
- Rpd memory might leak when RSVP LSP is cleared or re-signaled. [PR1415774](#)
- Services dependent on LDP might be impacted if committing any configuration changes. [PR1416032](#)
- Traffic blackhole might be seen due to a long LSP switchover duration in RSVP-signaled LSP scenario. [PR1416487](#)
- LDP route might be missing in inet.3 when enabling sr-mapping-client on LDP-SR stitching node. [PR1416516](#)
- RSVP LSP might get stuck in down state in OSPF Multiarea topology. [PR1417931](#)
- Traffic might be dropped due to LDP label corruption after Routing Engine switchover. [PR1420103](#)
- Incorrect length for Sub-TLV 34 (RFC 8287) in MPLS Echo Request. [PR1422093](#)
- LDP might not update the LDP ingress route metric when inet.3 route flash happens before inet.0. [PR1422645](#)
- The dynamic bypass RSVP LSP tears down when being used to protect LDP LSP. [PR1425824](#)
- mpls ping sweep stops working and CLI gets unresponsive. [PR1426016](#)
- MPLS LSP auto-bandwidth statistics miscalculations might lead to a high bandwidth reservation. [PR1427414](#)
- M/MX: continuous rpd core @ `l2ckt_alloc_label`, `l2ckt_standby_assign_label`, `l2ckt_intf_change_process` in the new backup during GRES in MX2010 box. [PR1427539](#)

- Traffic loss might be observed after changing configuration under **protocols mpls** in **ldp-tunneling** scenario. [PR1428081](#)
- The LDP might withdraw a label for an FEC once the IGP route is inactive in inet.0. [PR1428843](#)
- When MBB for P2MP LSP fails, it is stuck in the old path. [PR1429114](#)
- MPLS ingress LSPs for LDP link protection are not coming up after the disabling/enabling of MPLS. [PR1432138](#)
- SRLG entry shows Unknown after removing it from configuration in **show mpls lsp extensive** output or **show mpls srlg**. Shows **Unknown-0xXX** (XX will vary). [PR1433287](#)
- The P2MP LSP branch traffic might be dropped for a while when the Sender Provider Edge is performing switchover. [PR1435014](#)
- Traffic loss might be seen after LDP session flaps rapidly. [PR1436119](#)
- The rpd might crash after executing **ping mpls ldp**. [PR1436373](#)
- The LDP route and LDP output label are not showing in the **inet.3** table and LDP database respectively if enable OSPF rib-group. [PR1442135](#)
- LINX:lsi intf/Layer2 Virtual Chassis goes down on one router in VPLS domain through the mpls path is still available in inet.3. The 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 Junos pre-16.1 and 16.1+ nodes. [PR1443811](#)
- 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 down state. [PR1445024](#)
- Traffic might silently drop if two consecutive PLRs along the LSP perform local repair simultaneously under certain **mis-configured** conditions. [PR1445994](#)
- The transit packets might be dropped if an LSP is added or changed on MX/PTX device. [PR1447170](#)
- Traffic drop might be seen after traceoption configuration committed in RSVP P2MP scenario. [PR1447480](#)
- The LDP route timer is reset when committing unrelated configuration changes. [PR1451157](#)
- Previous configured credibility preference it is not considered by CSPF despite the configuration is deleted or changed to prefer another protocol in TED. [PR1460283](#)
- RPD core files and high CPU usage is seen on MX104. [PR1460292](#)

Network Address Translation (NAT)

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

Network Management and Monitoring

- The snmp query might not get data in scaled l2circuits environment. [PR1413352](#)
- Syslog match filtering does not work if single line of `/etc/syslog.conf` is over 2048 bytes. [PR1418705](#)
- MX10000 reports jail socket errors. [PR1442176](#)
- hrProcessorFrwID will be set to 0.0 of type "OBJECT IDENTIFIER" to fix the NMS warnings as it is using Integer value not OCTET STRING. [PR1446675](#)

Platform and Infrastructure

- All FPC cards might restart after Layer 3 VPN routes churn. [PR1398502](#)
- Class-of-service configuration changes might lead to traffic drop on cascade port in Junos fusion setup. [PR1408159](#)
- Traffic is getting dropped when there is a combination of DPC/MX-FPC card and MPC card on egress PE router in Layer 3 VPN. [PR1409523](#)
- DDoS violation for lldp, mvrp, provider mvrp and dot1x is incorrectly reported as LACP DDoS violation. [PR1409626](#)
- The VLAN tag is incorrectly inserted on the access interface if the packet is sent from an IRB interface. [PR1411456](#)
- Error logs might be observed after performing unified ISSU. [PR1412463](#)
- The MPC might crash when one MIC is pulled out during this MIC is booting up. [PR1414816](#)
- Distributed multicast forwarding to the subscriber interface might not work. [PR1416415](#)
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10 (based system). [PR1417321](#)
- op URL command can't run a script with libs from `/config/scripts`. [PR1420976](#)
- The ARP request might not be replied although `proxy-arp` is configured. [PR1422148](#)
- The slax scripts triggered by event options might be stuck forever. [PR1422939](#)
- `show jnh trap-info` with incorrect LU instance caused a crash and core file on FPC. [PR1423508](#)
- The native VLAN ID of packets might fail to be removed when leaving out. [PR1424174](#)
- The policer bandwidth might be wrong for the aggregate interface after activating the `shared-bandwidth-policer`. [PR1427936](#)
- With CNH for 6PE, MPLS EXP rewrite rule for non-VPN IPv4 over MPLS traffic might not work. [PR1430878](#)
- Pre-fragmented ICMP IPv4 packets might fail to arrive at the destination. [PR1432506](#)

- The FPC might crash when the firewalls filter manager deals with the firewall filters. [PR1433034](#)
- Enable sensor `/junos/system/linecard/qmon/` causing continuous `ppe_error_interrupt` errors. [PR1434198](#)
- Traffic from the same physical interface cannot be forwarded. [PR1434933](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- The IPv4 packet larger than `mtu-v6` might be dropped by the MAP-E BR device. [PR1435362](#)
- MAP-E encapsulation or decapsulation with specific parameter might work incorrectly. [PR1435697](#)
- The `/var/db/scripts` directory might be deleted after executing `request system zeroize`. [PR1436773](#)
- The BGP session might flap after Routing Engine switchover done simultaneously on both boxes of BGP peer in scaled BGP session setup. [PR1437257](#)
- The next-hop mac address in the output from `show route forwarding-table` command might be incorrect. [PR1437302](#)
- The multicast traffic is dropped while multicast ingress replication is configured with `local-latency-fairness`. [PR1438180](#)
- A certain combination of allow-commands and deny-commands do not work properly after Junos OS Release 18.4R1. [PR1438269](#)
- The inner IPv4 packet might get fragmented using the same size as `mtu-v6` setting which is used for the MAP-E software tunnel in MAP-E configuration. [PR1440286](#)
- When host bound packet received in MAP-E BR router, service interface statistics counter shows incorrect number of bytes. [PR1443204](#)
- Packets drop due to missing destination MAC in the Packet Forwarding Engine. [PR1445191](#)
- Python op scripts executed as user `nobody` if started from NETCONF session, not as a logged in user, resulting in failing PyEZ connection to the device. [PR1445917](#)
- Some hosts behind unnumbered interface are unreachable after the router/FPC restarts. [PR1449615](#)
- FPC might reboot with vmcore due to memory leak. [PR1449664](#)
- REST API process will get non-responsive when a number of request is coming with a high rate. [PR1449987](#)
- The Routing Engine originated IPv6 packets might be dropped when interface-group rule is configured under IPv6 filter. [PR1453649](#)

Routing Policy and Firewall Filters

- Configuration commit operation after policy change causes rpd crash. [PR1357802](#)
- MX Series: CLI knob `as-path-expand last-as` commit failure. [PR1388159](#)
- The `route-filter-list` with non-continuous match might not work as expected after being updated. [PR1419731](#)

- Policy matching RD changes next-hop of the routes which do not carry RD. [PR1433615](#)
- Routes resolution might be inconsistent if any route resolving over the multipath route. [PR1453439](#)

Routing Protocols

- Junos BGP Established state is not shown in **show bgp summary** if only master routing instance is present. [PR600308](#)
- RPD crashes due to assert in **bgp_io_write_user_handler_int()**. [PR1351639](#)
- Qualified next hop of static route might not be withdrawn when BFD is down. [PR1367424](#)
- Routing Engine-based micro BFD packets do not go out with configured source IP when the interface is in logical-system. [PR1370463](#)
- The rpd might crash under a rare condition if GR helper mode is triggered. [PR1382892](#)
- BGP sessions might keep flapping on backup Routing Engine if **proxy-macip-advertisement** is configured on IRB interface for EVPN-VXLAN. [PR1387720](#)
- In rare cases rpd might crash after Routing Engine switchover when BGP multipath and L3VPN vrf-table-label are configured. [PR1389337](#)
- Processing a large scale as-path regex will cause the flap of the route protocols. [PR1396344](#)
- There might be unexpected packets drop in MoFRR scenario if active RPF path is disabled. [PR1401802](#)
- IGMP join through PPPOE sub not propagated to upstream PIM. [PR1407202](#)
- BFD link-failure detection of the broken path will be delayed when IGP link-state update is received from the same peer through an alternative path. [PR1410021](#)
- BGP might get stuck in an Idle state when the peer triggers a GR restart event. [PR1412538](#)
- The Layer 3 VPN link protection doesn't work after flapping the CE facing interface. [PR1412667](#)
- TI-LFA cannot find backup path when ISIS OverLoad bit is set on computing node. [PR1412923](#)
- SID label operation might be performed incorrectly in OSPF SPRING environment. [PR1413292](#)
- The unexpected AS prepending action for AS path might be seen after the **no-attrset** statement is configured or deleted with **vrf-import/vrf-export** configuration. [PR1413686](#)
- Dynamic routing protocol flapping with vmhost Routing Engine switchover on NG-RE. [PR1415077](#)
- The IS-IS-SR route sent by the mapping server might be broken for ECMP. [PR1415599](#)
- Route info might be inconsistent between RIB and OSPF database when using OSPF LFA feature. [PR1416720](#)
- Junos OS: OpenSSL Security Advisory [26 Feb 2019]. [PR1419533](#)
- A memory leak in rpd might be seen if source packet routing is enabled for IS-IS protocol. [PR1419800](#)
- BFD crash after GRES was done @ `__assert (func=0x831a40e "bfdd_link_session", file=0x831a24a "../src/junos/usr.sbin/bfdd/bfdd_session.c"` [PR1420694](#)

- IPv6 IS-IS routes might be deleted and not be reinstalled when MTU is changed under the IFL level for family inet6. [PR1420776](#)
- Route churn might be seen after changing **maximum-prefixes** configuration from value A to value B. [PR1423647](#)
- The rpd might crash if **no-propagate-ttl** is configured in BGP multipath scenario. [PR1425173](#)
- The multicast traffic might be dropped when proxy mode is used for igmp-snooping. [PR1425621](#)
- The rpd might crash in PIM scenario with **auto-rp** enabled. [PR1426711](#)
- The rpd might crash while removing multicast routes that do not have an associated (S,G) state or activating the **accept-remote-source** statement on PIM upstream interface. [PR1426921](#)
- The rpd might crash while handling the withdrawal of an imported VRF route. [PR1427147](#)
- MVPN traffic might be lost for around 30 seconds during Routing Engine switchover. [PR1427720](#)
- The rpd would generate core files due to improper handling of Graceful Restart stale routes. [PR1427987](#)
- RPD might crash with ospf overload configuration. [PR1429765](#)
- The next-hop of IPv6 route remains empty when a new IS-IS link comes up. [PR1430581](#)
- BGP knob **multipath multiple-as** does not work in specific scenario. [PR1430899](#)
- IPv6 aggregate routes are hidden. [PR1431227](#)
- Unsupported configuration (EPE with dynamic-next-hop GRE tunnels) continuously causing RPD to generate core files. [PR1431536](#)
- The **show isis adjacency extensive** output is missing state transition details. [PR1432398](#)
- In BFD and GR enabled scenario, BFD DOWN packets are not being sent immediately after BFD failure. [PR1432440](#)
- Per-Prefix LFA might not work as expected where the last hop needs to be protected on the penultimate node. [PR1432615](#)
- PIM-SM join message might be delayed with MSDP enabled. [PR1433625](#)
- With SR enabled 6PE next-hop is not installed. [PR1435298](#)
- The rpd might crash during the best path changes in BGP-L3VPN with **multipath** and **no-vrf-propagate-ttl** enabled. [PR1436465](#)
- BGP route next-hop can be incorrect in some scenarios with PIC edge configuration. [PR1437108](#)
- Removing SSH Protocol version 1 from configuration. [PR1440476](#)
- RIP routes are discarded by Juniper devices when the next-hop field in the RIPv2 response packet contains a subnet Broadcast address. [PR1441452](#)
- The rpd process might crash in inter-AS option B Layer 3 VPN scenario if CNHs is used. [PR1442291](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)

- The rpd might crash in OSPF scenario due to invalid memory access. [PR1445078](#)
- 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 might crash when the policy applied to the MoFRR is deleted. [PR1446472](#)
- The rpd CPU utilization gets 100% 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](#)
- Layer 3 VPN PE-CE link protection exhibits unexpected behavior on MX2000 platforms. [PR1447601](#)
- Junos BFD sessions with authentication flaps after a certain time. [PR1448649](#)
- Intra-router PPMD[RE] to PPMAN[FPC] connection could be closed if the session timeout is greater than 3 seconds in either direction. [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](#)
- The rpd memory might leak in a certain MSDP scenario. [PR1454244](#)
- The rpd might crash continuously due to memory corruption in IS-IS setup. [PR1455432](#)
- Prefix SID conflict might be observed in ISIS. [PR1455994](#)
- Routing-process is crashing when OSPF router-id get changed for NSSA area. [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](#)

Services Applications

- ms- used for IPSEC PIC is listed in show services ha detail as standby. This is a cosmetic issue. [PR1383898](#)
- **SPD_CONN_OPEN_FAILURE: spd_svc_set_summary_query:** unable to open connection to si-0/0/0 (No route to host). [PR1397259](#)
- [technology/subscriber_services/jl2tpd] [all] RPT BBE Regressions : ERA Value does not match with configured values while verifying new ERA settings and they are reflected in message logs. [PR1410783](#)
- jpppd generates core files on LNS. [PR1414092](#)
- L2TP LAC might fail to tunnel static pp0 subscriber to the desired LNS. [PR1416016](#)
- IPsec SA might not come up when the Local gateway address is a VIP for a VRRP configured interface. [PR1422171](#)
- In subscriber with L2TP scenario, subscribers are stuck in INIT state forever. [PR1425919](#)
- Some problems might be seen if client negotiates LCP with no **ppp-options** to LAC. [PR1426164](#)

- The kmd process might crash when DPD timeout for some IKEv2 SAs happens. [PR1434521](#)
- Traffic might be dropped in IPsec VPN scenario when the VPN peer is behind a NAT device. [PR1435182](#)
- The output of **show subscriber user-name** on LTS shows only one session instead of two. [PR1446572](#)
- The jl2tpd process might crash during the restart procedure. [PR1461335](#)

Software Installation and Upgrade

- JSU might be deactivated from FPC in case of power cycle. [PR1429392](#)

Subscriber Access Management

- Authd telemetry: Linked pool head attribute is incorrect for single pools. [PR1413293](#)
- The subscriber service profile might be unable to be changed by RAR message in PCRF/Gx-Plus scenario. [PR1417987](#)
- CoA-NACK is not sent when performing negative COA Request tests by sending incorrect **session-id**. [PR1418144](#)
- Subscribers might not be able to re-login in Gx-plus provisioning scenario. [PR1418579](#)
- PPPoE session might be disconnected when LI attributes are received in access-accept with invalid data. [PR1418601](#)
- Address allocation issue with linked pools when using linked-pool-aggregation. [PR1426244](#)
- RADIUS authentication server might always be marked with DEAD. [PR1429528](#)
- Subscriber filtering for General Authentication Services traceoptions will report debug messages for other users. [PR1431614](#)
- Incorrect Acct-Session-Time : Acct-Session-Time is not zero, though no Start event occurred. [PR1433251](#)
- The output of **test aaa ppp** is missing **<radius-server-data>** tag. [PR1444438](#)
- On MX platforms a false error might be received for SAE policy activation or deactivation failure. [PR1447632](#)
- Subscribers Login fails when PCRF Server is unreachable. [PR1449064](#)
- DHCPv6 subscribers might be stuck in a state after the authd process crash. [PR1460578](#)
- Problem with **linked-pool-aggregation** after attempting to delete a pool in middle of the chain. [PR1465253](#)

User Interface and Configuration

- Junos fusion: **show chassis hardware satellite** command is not available on Junos OS 17.3 versions. [PR1388252](#)
- Junos OS: Insecure management daemon (MGD) configuration might allow local privilege escalation (CVE-2019-0061). [PR1406219](#)

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 multicast traffic drop might be seen when **static-umh** is configured in NGMVPN scenario. [PR1414418](#)
- The rpd might crash in rosen MVPN scenario when a same provider tunnel source address is being used for both IPv4 and IPv6. [PR1416243](#)
- The deletion of (S,G) entry might be skipped after the PIM join timeout. [PR1417344](#)
- The rpd crash might be seen if layer 2 circuit or local-switching connections flap continuously. [PR1418870](#)
- The rpd process might crash in rare conditions when Extranet NG-MVPN is configured. [PR1419891](#)
- Permanent traffic loss is seen on NGMVPN selective tunnels after Routing Engine switchover (one-time). [PR1420006](#)
- MPLS LSP ping over I2circuit might not work when **flow-label** is enabled. [PR1421609](#)
- The rpd process might crash and generates core files during mpls ping command on I2circuit. [PR1425828](#)
- MVPN using PIM Dense mode does not prune the OIF when PIM prune is received. [PR1425876](#)
- The resumed multicast traffic for certain groups might be stopped in overlapping MVPN scenario. [PR1441099](#)
- Memory leak might happen if PIM messages 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](#)

Resolved Issues: 19.1R1

Application Layer Gateways (ALGs)

- DNS requests with EDNS options might be dropped by DNS ALG. [PR1379433](#)
- On MX Series routers with FTP ALG enabled, if there are more than one FTPS connection between a pair of FTP client and server, the closure of one connection might affect the other connections between the pair of FTP client and server. There might be traffic impact. It is a rare timing issue. [PR1483834](#)

Authentication and Access Control

- MAC move might occur in a DHCP security scenario. [PR1369785](#)
- The dot1xd might crash when dot1xd receives incorrect reply length from the authd. [PR1372421](#)
- IPv4/IPv6 DHCP security client entries will be recorded on TRUSTED ports as well. [PR1390676](#)
- Push-to-JIMS now supports pushing the authenticated entry to all online JIMS servers. [PR1407371](#)

Class of Service (CoS)

- The cosd process might crash during committing configuration change through netconf. [PR1403147](#)

Flow-Based and Packet-Based Processing

- Issues occur with fragmentation and ALG support for Power Mode IPsec. [PR1397742](#)

EVPN

- EVPN type-5 route might be lost if **chained-composite-next-hop** statement is configured. [PR1362222](#)
- Packet drop is seen in EVPN stitching with IRB configured. [PR1363935](#)
- The EVPN implementation does not follow RFC-7432. [PR1367766](#)
- EVPN A/A multihomed PE device occasionally prefers to route to a directly connected prefix using LSPs toward the multihomed peer instead of going directly out of the IRB interface (which is up). [PR1376784](#)
- The RA packets might be sent out without using the configured virtual gateway address. [PR1384574](#)
- EVPN-VXLAN: Flood is not programmed for VTEP interfaces for more than 5 minutes after BGP bounce. [PR1396597](#)
- IPv6 link-local address for virtual-gateway address is marked as duplicate in EVPN. [PR1397925](#)
- When committing a configuration for a VLAN adding to an EVPN instance and an aggregated Ethernet interface, respectively, the newly added VLAN interface count might be zero (0) in that bridge domain. [PR1399371](#)
- EVPN type 2 MAC+IP route is stuck when the route advertisement has two MPLS labels and withdrawal has one label. [PR1399726](#)
- The rpd core file is generated upon Routing Engine switchover with scaled EVPN configuration. [PR1401669](#)
- The rpd crashes due to memory corruption in EVPN. [PR1404351](#)
- EVPN database and bridge mac-table are out of sync after core link flap. [PR1404857](#)
- The rpd might crash on a leaf node when handling withdrawal of the remote or local MAC address in an EVPN-VXLAN scenario. [PR1405681](#)
- The rpd might crash after NSR switchover in a EVPN scenario. [PR1408749](#)

Forwarding and Sampling

- In EVPN A/A scenario with MX or EX acting as PE device, flood next-hops to handle BUM traffic might not get created or miss certain branches when the configuration is performed in a particular sequence. [PR1377749](#)
- LTS subscriber statistics are reported to RADIUS. [PR1383354](#)
- Adjusting **mac-table-size** configuration might cause an l2ald crash. [PR1383665](#)
- The LSI binding for the IPv6 neighbor is missing. [PR1388454](#)

- The filter counter is not written to the accounting file when accounting is enabled on the bridge firewall filter. [PR1392550](#)
- The l2ald process might crash when doing **commit check** for some specific configurations. [PR1395368](#)

General Routing

- We advise migrating from syslog API to Errmsg API: `/src/junos/usr/sbin/mspsmd`. [PR1284654](#)
- MX150: Cannot copy files from the USB flash to Junos OS Virtual Machine. [PR1333201](#)
- Large-scale users logging in and logging out might cause a mgd memory leak. [PR1352504](#)
- Traffic loss might be seen on the new master after interface flap is followed by Routing Engine switchover in a VRRP scenario. [PR1353583](#)
- The packets might be dropped when they go through the MX104 built-in interface. [PR1356657](#)
- The **show chassis ethernet-switch** command output is different on MX10008 routers. [PR1358853](#)
- MX Series BNG does not generate the ESMC/SSM quality level failed SNMP trap alarm. [PR1361430](#)
- The inline J-Flow sampling configuration might cause an FPC crash on MX Series routers. [PR1362887](#)
- MX Series Virtual Chassis: The request to record the VCCP heartbeat state changes in syslog by default. [PR1363565](#)
- FPM board status is missing in the SNMP MIB walk result. [PR1364246](#)
- The netproxy service client component fails to restart after issuing the **request vmhost reboot** command. [PR1365664](#)
- The following errors are seen in the syslog: **LOG : Err] Failed to allocate 2 jnh-dwords for encap_ptr(ether-da)!,LOG: Err] gen_encap_common: jnh-alloc failed! 8** [PR1366811](#)
- When you configure vrrp delegate-processing with tomcat enabled, the Packet Forwarding Engine dropped vrrp packets and count sw error. [PR1369503](#)
- MPC5E restarted at `trinity_pio_io_func, pio_read_u32, xqchip_read_u32, xqchip_issu_disable_q_stats, qchip_issu_disable_q_stats, issu_asic_prepare (pfe_idx=0 '\000')` at `../..../src/pfe/common/applications/issu/jam/issu_jam_npc_pfe.c:65` [PR1369635](#)
- Image installation on SD fails with the following error: **Unable to read reply from software add command to re1; error 1.** [PR1372877](#)
- Core file is generated in ifinfo at `pif_af_fe_info pif_af_ifd` when displaying the af interface information. [PR1373436](#)
- LDP convergence delay might be seen after a IGP metric change with the **bgp-igp-both-ribs** statement configured. [PR1373855](#)
- The filter service might fail to get installed for the subscriber in a scaled BBE scenario. [PR1374248](#)
- A few L2BSA subscribers might be stuck in init, terminating, or terminated state after the previous logout. [PR1375070](#)

- SFB and PDM/PSU-related information is missing in jnxBoxAnatomy MIB on high-end MX Series routers. (MX2010/2020). [PR1375242](#)
- The bbe-smgd core file might be seen after doing GRES. [PR1376045](#)
- MS-MPC might have performance degradation under scaled fragmented packets. [PR1376060](#)
- Interface optic output power is not zero when the port has been disabled. [PR1376574](#)
- CI: Not generating Power Supply failed trap. [PR1376612](#)
- After NAT64 router (with MS-MPC) translates an IPv6 fragment to an IPv4 fragment, router is not inserting the correct value in the identification field of the IPv4 header. [PR1378818](#)
- The bbe-smgd process generates repeated core files and stops running as a result of long-term session database shared memory corruption. [PR1388867](#)
- Traffic might be discarded without notification when CoS configuration is changed on a PS interface. [PR1379530](#)
- Protocol adjacency might flap and FPC might reboot if jlock hog happens. [PR1379657](#)
- MSQQ error logs and potential MPC traffic impact are seen when the physical interface link goes down. [PR1380183](#)
- The pfe_disable action should also disable the logical interfaces belonging to the affected Packet Forwarding Engine. [PR1380784](#)
- Encryption and decryption is not happening, because the Packet Forwarding Engine discards it while testing that group-vpn member was established using the authentication-method preshared key ascii-text. [PR1381316](#)
- Traffic might be discarded without notification that is caused by FPC offline in MC-LAG scenario. [PR1381446](#)
- In MX3ru for Junos OS Release 18.3R1, unified ISSU will fail if QSA is plugged in. [PR1382126](#)
- The MPC6E might crash while fetching PMC device states. [PR1382182](#)
- High CPU utilization is seen for chassisd on bsys, ~20 percent st steady state. [PR1383335](#)
- The configuration configured through the NETCONF session might fail. [PR1383567](#)
- MBFD flaps because clksync congests the scheduler for 100 ms. [PR1384473](#)
- The rpd generates a core file at `krt_table_rtbit_q_handler`, `krt_q_flush` (startp=0xccca2c500, endp=0xccca2e9d0, isflash=0, todo=0x7fffffff204), `rtbit_free` (rtbh=0x4145540). [PR1385005](#)
- The MPLS packets with more than eight labels will not be processed by J-Flow. [PR1385790](#)
- The vFPC CPU is running very high on vMX. [PR1385853](#)
- The device with more than five IP addresses configured in the DHCP server-group goes into amnesiac mode after reboot. [PR1385902](#)
- In subscriber management environment DHCP Subscriber might get stuck in terminated state. [PR1386662](#)

- In case an LSP is locally configured without an explicit path, the ERO object remains empty in the PCRpt generated by PCC. [PR1386935](#)
- Uninitialized EDMEM[0x400094] Read (0x6db6db6d6db6db6d) logs are seen with sampling applied to a subscriber with routing-service applied. [PR1386948](#)
- The rpd might crash when traceoptions are enabled. [PR1387050](#)
- On MX2000 routers, the backup CB's chassis environment status shows up as "Testing" even after the backup CB becomes online by removal or insert operation. [PR1387130](#)
- The bbe-smgd process might crash when two subscribers log in with the same framed-route prefix and preference values. [PR1387690](#)
- Some SFBs might go down when one of the PSMs in the chassis generates a bad output voltage that is out-of-range. [PR1387737](#)
- FPC core file is seen at sensor_export_time_exceed_limit agent_health_monitor_data_reap when **Jinsight** is configured. [PR1388112](#)
- Psec IKE keys are not cleared when delete or clear notification is received from peer on GRES-enabled DUT. [PR1388290](#)
- Fabric drops might be seen if using a newer generation of MPC with SFB2. [PR1388780](#)
- Incorrect value for flow packets or octets fields might be seen in an inline J-Flow scenario, [PR1389145](#)
- IGMP group threshold exceed log message prints an incorrect demux logical interface. [PR1389457](#)
- MX204: Excluding the **speed** CLI option under the interface level. [PR1389918](#)
- The jnxFruState might show incorrect PIC state after replacing an MPC with another MPC with less number of PICs. [PR1390016](#)
- Traffic destined to VRRP VIP gets dropped because the filter is not updated to the related logical interface. [PR1390367](#)
- The **delete chassis redundancy** command with routing-options nonstop-routing is not giving a commit warning. [PR1390575](#)
- Delay in CLI output with second or more **show subscriber <> extensive** queries when the first session is sitting at -(more)- prompt displaying **show subscribers extensive** command output. [PR1390762](#)
- Trailing chars are seen in GNMI get API reply. [PR1390967](#)
- All the BBE and ESSM subscriber sessions might be lost after GRES or unified ISSU. [PR1391409](#)
- The **routing-engine-power-off-button-disable** statement does not work on MX204 and MX10003 routers. [PR1391548](#)
- The bbe-smgd process might crash after committing configuration changes. [PR1391562](#)
- The bbe-smgd process might crash in a corner case if family inet6 is used in the dynamic profile. [PR1391845](#)

- On MX2000, fans start spinning at high speed upon inserting previously offlined FPC. [PR1393256](#)
- Third-generation FPC reboot loop is caused because of having internal intf issues. [PR1393643](#)
- Junos OS enhancement configuration statement added to modify mcontrol watchdog timeout. [PR1393716](#)
- If FPGA on the new master CB has a specific hardware failure, the chassis might keep crashing after GRES switchover. [PR1393884](#)
- Expected entries like "UI_COMMIT_PROGRESS" are not getting populated while checking with Junos OS script session for obtaining the syslog output. [PR1394780](#)
- MPC7, MPC8, and MPC9 might not boot in the MX Series Virtual Chassis. [PR1396268](#)
- Adding IRB to bridge-domain with PS interface causes a kernel crash. [PR1396772](#)
- The MS-MPC might generate a core file when mspmand receives a non-synchronized packet of TCP. [PR1396785](#)
- A smid process memory leak occurs, and it does not come down from 100 percent. [PR1397643](#)
- PFT MX10008: Inline-services enabling the **Flex-Flow-Sizing** take more than 12 minutes to move to steady state. [PR1397767](#)
- The **show system errors active** command is not showing the error for MPC3E next-generation HQoS. [PR1398084](#)
- Kernel core file is generated on vMX. [PR1398320](#)
- MPLSoUDP tunnels do not come up on interface route - dyn_tunnel_fwd_route_eligible - because next-hop type is configured as interface. [PR1398362](#)
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. [PR1398398](#)
- IPsec tunnel cannot be established, because the tunnel SA and rule are not installed in the PIC. [PR1398849](#)
- The bbe-smgd process might crash when executing the **show pppoe lockout** command. [PR1398873](#)
- Wrong timestamp is displayed in the jvision collector log file. [PR1399829](#)
- ARP refresh functionality might fail in an EVPN scenario. [PR1399873](#)
- JET/PRPD incompatibility for the rib_service.proto field RouteGateway.weight occurs from Junos OS Release 18.4R1 to Release 18.4R2 and onward. [PR1400563](#)
- The mgd-api crashes due to memory leak. [PR1400597](#)
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- The authd might crash when issuing the **show network-access requests pending** command during the authd restarting. [PR1401249](#)
- The **show | compare** output on global group changes loses the diff context after a rollback or "load update" is performed. [PR1401505](#)
- The subscriber route installation fails because some interface states are not properly installed. [PR1401506](#)

- FPC core files are generated due to a corner case scenario (race condition between RPF, IP flow). [PR1401808](#)
- JET authentication does not work for usernames and passwords of certain lengths. [PR1401854](#)
- Traffic loss is seen for IGMP subscribers after GRES. [PR1402342](#)
- The MPC might crash due to the CPU hogging by dfw thread. [PR1402345](#)
- Some error logs might be seen on FPC when reading attempt from uninitialized memory location. [PR1402484](#)
- FPC might crash after you offline or online MIC-3D-16CHE1-T1-CE-H. [PR1402563](#)
- DHCP subscriber cannot reconnect over dynamic VLAN demux interfaces because of RPF check failure. [PR1402674](#)
- Host outbound traffic might be dropped on MPC7, MPC8, and MPC9 [PR1402834](#)
- Observed rpd core files when a few colored LSPs are changed to uncolored LSPs. The cores are at <<< #0 tag_cmp_tag (tag1=0x0, tag_label1=0x0, tag2=0x98b6628, tag_label2=0x98b6644) at ../../../../src/junos/usr/sbin/rpd/lib/mpls/label_mgr/core/mpls_label.c:473 473 if (tag1->tagt_mtu != tag2->tagt_mtu) >>> [PR1403208](#)
- Reported Log Variance might be incorrect if the PTP profile is changed from G.8275.2 to SMPTE or another multicast IP profile. [PR1403219](#)
- Smg-service might become unresponsive. [PR1403480](#)
- The time synchronization through PTPoE might not work when Enhanced Subscriber Management is enabled on MX Series routers. [PR1404002](#)
- The FPC might crash in a CoS scenario. [PR1404325](#)
- The repd continues to generate core files on VC-Bm when there are too many IPv6 addresses on one session. [PR1404358](#)
- The **targeted-broadcast** statement does not work on the IRB interface. [PR1404442](#)
- Configuration load override or load replace resets ANCP neighbors. [PR1405318](#)
- NAT64 translation issues of ICMPv6 packet-too-big message occur with MS-MPC/MS-PIC. [PR1405882](#)
- Fabric performance drop is seen on MPC7, MPC8, MPC9E, and SFB2-based MX2000 routers. [PR1406030](#)
- Traffic impact might be seen if **auto-bandwidth** is configured for RSVP LSPs. [PR1406822](#)
- New CLI option is introduced to display DF and MLR in split format. [PR1406884](#)
- Layer 2 VPN will flap repeatedly after link up between PE and CE devices under "asynchronous-notification" and "some types of MICs" conditions. [PR1407345](#)
- NPC core files are generated after daemon restart in #0 jnh_get_oif_nh (ifid=0x51a51a80, ifl=0x6aeb52e0, family_mtu=0, max_mpls_labels=0 '\000', pad_ge_frame_check=< optimized out>, ret_jnh=0x483a54a8) at ../../../../src/pfe/common/pfe-arch/trinity/toolkits/jnh/jnh_if.c:15248. [PR1407765](#)

- Ephemeral database might get stuck during commit. [PR1407924](#)
- Traffic forwarding fails when crossing VCF members. [PR1408058](#)
- Alarm mismatch in total memory is detected after running the **reboot vmhost both** command. [PR1408480](#)
- TFTP of MPC line cards images fails when performing unified ISSU. [PR1408558](#)
- Python script might stop working due to **Too many open files** error. [PR1408936](#)
- interface-set meta-data needs to include the CoS TCP names in order to aid collector reconciliation with queue-stats data. [PR1409625](#)
- FPC might generate core files during next-hop change due to FPC reboot or interface flap when using MPLS inline J-Flow. [PR1409807](#)
- When using SFP+, the interface optic output might be non-zero even when the interface has been disabled. [PR1410465](#)
- Traffic loss might be seen on MPC8E and MPC9E after requesting one of the SFB2s to go offline/online. [PR1410813](#)
- Kernel replication failure and vmcore are seen because **add IPv6 route prefix** operation is not supported with the next-hop to be ATM interface. [PR1411376](#)
- MX10003: The rpd crash with switchover-on-routing-crash does not trigger a Routing Engine switchover and the rpd on the master Routing Engine goes into STOP state. [PR1412322](#)
- During unified ISSU from Junos OS Release 16.1R4-S11.1 to Junos OS Release 18.2R2-S1.2, CoS GENCFG write failures are observed: [**COS(cos_rewrite_do_pre_bind_add_action:676): Binding of table 44226 to ifl 1073744636 failed, table already bound to ifl**]. [PR1413297](#)
- MPC10E line card will not power up in old MidPlane MX chassis when using Junos OS Release 19.1R1. [PR1413373](#)
- Broken support of [family inet6 filter] on the ATM interface. [PR1413663](#)
- The bbe-smgd process might have memory leak while running the **show system subscriber-management route route-type <> routing-instance <>** command. [PR1415922](#)
- In the scenario where the MX and the peer both try to bring an IPsec tunnel up but the peer side does not answer the MX requests, we can bring the peer initiated tunnel down. [PR1420293](#)
- The bbe-smgd process might crash and might not recover in a rare scenario. [PR1420376](#)

Infrastructure

- The error of **jlaunchd: disk-monitoring is thrashing, not restarted** might be seen. [PR1380032](#)

Interfaces and Chassis

- In case of MPLS, DMR packets are sent with different MPLS expiration bits if the MX Series router receives CFM DMM packets with varying expiration values on the MPLS header. [PR1365709](#)
- Constant dcpfe process crash might be seen if using an unsupported GRE interface configuration. [PR1369757](#)
- The jpppd process might crash if the EPD value contains a format specifier. [PR1384137](#)
- DCD core file can be seen after FPC restart if **channelized interfaces** are configured. [PR1387962](#)
- All DPCs might crash while adding or deleting a logical interface from the aggregated Ethernet bundle. [PR1389206](#)
- The interface-control process thrashes and dcd does not restart after adding an invalid demux interface to the configuration. [PR1389461](#)
- Decoupling of Layer 2 logical interface configuration from bridge-domain or EVPN configuration. [PR1390823](#)
- Interim accounting updates might not be sent for subscribers after Junos OS selective update. [PR1391011](#)
- A dcd memory leak might be seen when committing configuration change on static route tag. [PR1391323](#)
- Error message might be seen if GR interface is configured. [PR1393676](#)
- DCD crashes on deleting the sub interface from VPLS routing-instance when the same sub interface is also part of mesh-group. [PR1395620](#)
- The **MIC Error code: 0x1b0002** alarm might not be cleared for MIC on MPC6 when the voltage has returned to normal. [PR1398301](#)
- The backup Routing Engine might get stuck in amnesiac mode after reboot. [PR1398445](#)
- All dcd operations might be blocked if profile-db is corrupt. [PR1399184](#)
- Certain otn-options cause interface flapping during commit. [PR1402122](#)
- Subscriber might not be able to access the device due to the conflicted assigned address. [PR1405055](#)
- The cfmd might fail to start after it is restarted. [PR1406165](#)
- The aaa-options configuration statement for PPPoE subscribers does not work on the MX80 and MX104 routers. [PR1410079](#)

J-Web

- Junos OS: Persistent XSS vulnerability in J-Web (CVE-2019-0047). [PR1410400](#)
- Junos OS: Session fixation vulnerability in J-Web (CVE-2019-0062). [PR1410401](#)

Layer 2 Features

- The unicast traffic from IRB interface toward LSI might be dropped due to Packet Forwarding Engine mismatch at egress processing. [PR1381580](#)
- Flow label is still used by ingress PE though the egress PE is not configured/supporting for Flow label in a VPLS multihomed Scenario. [PR1393447](#)
- In a Layer 2 domain (for example, bridge-domain, VPLS), there is unexpected flooding of unicast traffic at approximately every 40 seconds toward all local CE-facing interfaces. [PR1406807](#)

Layer 2 Ethernet Services

- The subscriber's authentication might fail when the link-layer address that is encoded in the DHCPv6 DUID is different from the actual link-layer hardware address. [PR1390422](#)
- The SNMP query on LACP interface might lead to lacpd crash. [PR1391545](#)
- Log messages `dot1xd[]: task_connect: task ESP CLIENT:....: Connection refused` might be reported in Junos OS Release 17.4 or later. [PR1407775](#)

MPLS

- The rpd might crash on the backup Routing Engine after switchover. [PR1382249](#)
- MPLS LSP will remain in the down state due to routing loop detection after flapping link between PE router and egress PE router. [PR1384929](#)
- Configured bandwidth 0 does not get applied on the RSVP interface. [PR1387277](#)
- The bypass LSP might pass through an unexpected path that includes the same SRLG as the protected TE link that is down. [PR1387497](#)
- The rpd process might keep crashing repeatedly if the LSP destination address is set to be 0.0.0.0. [PR1397018](#)
- The rpd might crash when the LDP route with indirect next-hop is deleted. [PR1398876](#)
- A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based statistics are used. [PR1401152](#)
- Resources might be reserved for stale RSVP LSP when RSVP is disabled on the interface. [PR1410972](#)
- LDP crash is caused by an `ldp_label_bind_route` assert condition. [PR1413231](#)
- LDP native IPv6 loopback remains in inet6.3 after removing the IPv6 address from the core interface. [PR1414965](#)

Platform and Infrastructure

- MQCHIP CPQ block should report major alarm. [PR1276132](#)
- Some line cards might crash in a subscriber scenario enabled with distributed IGMP. [PR1355334](#)
- The FPC might crash continuously when the filters in the same filter list refer to a same nested filter [PR1357531](#)

- The kernel and ksyncd core files are generated after dual cb flap at rt_nhfind_params: rt_nhfind() found a next-hop different from that on the master 30326. [PR1372875](#)
- The traffic traversing an IRB interface might not be tagged with a VLAN if the packets go through an additional routing instance. [PR1377526](#)
- IPv6 ping might fail for spine node in a EVPN scenario. [PR1380590](#)
- Packet drops on interface occur if the **gigether-options loopback** statement is configured. [PR1380746](#)
- The jcrypto syslog help package and events are not packaged even when the following error message is compiled: **Below KMD help syslogs are missing help syslog | match KMD**
KMD_IKEV2_CONFIG_PAYLOAD_FAILED IKE Phase-1 Config Payload failed
KMD_INSTALL_JUNOS_IKE Inform user to install junos-ike package **KMD_PEER_CERT_VERIFY_FAILED**
IKE Phase-1 peer certificate verification failed **KMD_VPN_DFBIT_STATUS_MSG** VPN DF bit status set **KMD_VPN_DOWN_ALARM_USER** VPN monitor detected IPsec SA is down
KMD_VPN_PV_LIFETIME_CHANGED Config key lifetime test changed **KMD_VPN_PV_PHASE1** ike-phase1 failures over configured auditable threshold **KMD_VPN_PV_PHASE2** ike-phase2 failures over configured auditable threshold **KMD_VPN_PV_PSK_CHANGED** Pre-shared Key changed
KMD_VPN_PV_REAUTH_FREQUENCY_CHANGED Config key reauth frequency changed
KMD_VPN_PV_XAUTHC_AUTH_FAIL ike-xauth-client authentication failure **KMD_VPN_TS_MISMATCH** IPsec traffic-selector mismatch **KMD_VPN_UP_ALARM_USER** VPN monitor detected that an IPsec SA is up. **KMD_XAUTH_AUTHENTICATE_FAILED** IKE Phase-1 XAUTH authentication failed [PR1290089](#)
- The dfwd might crash with **DFWD_TRASHED_RED_ZONE** log messages. [PR1380798](#)
- Traffic loss is seen in Layer 2 VPN with GRE tunnel. [PR1381740](#)
- MAC learning might get stuck on MX Series routers with DPC and MPC. [PR1383233](#)
- jlock hog is reported at restart routing. [PR1389809](#)
- Individual command authorization might cause mgd crash. [PR1389944](#)
- Traffic is being dropped when passing through MS-DPC to MPC. [PR1390541](#)
- The RADIUS authentication does not work through management-instance for IPv6 family. [PR1391160](#)
- The lockout-period might not work for the user being locked out. [PR1393839](#)
- RVT interface might start flapping. [PR1399102](#)
- In a scaled scenario (500 TWAMP control sessions and 500 TWAMP test sessions), a few TWAMP connections might fail to establish. [PR1399547](#)
- Syslog error messages: **[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_pfe:3211): pfe_id 0 cchip 0**
[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_pfe:3272): No PFE found for pfe_id_start 0.
[PR1402377](#)
- MAP-E some ICMP types cannot be encapsulated/decapsulated on SI interface. [PR1404239](#)
- When a non-root user tries to archive the **var/log**, some files are missing if a **cscript.log** file exists. [PR1405903](#)

- Abnormal queue-depth counters appear in **show interface queue** command output on interfaces associated to XM2 and 3. [PR1406848](#)
- Ipv6 drops due to **output trunk vlan lookup failed**. [PR1407200](#)

Routing Policy and Firewall Filters

- The **set metric multiplier offset** might overflow or underflow. [PR1349462](#)
- The rpd process might crash if **then next-hop** is configured for the LDP export policy. [PR1388156](#)
- The rpd process might crash when **routing-options flow** configuration is removed. [PR1409672](#)

Routing Protocols

- BGP might not advertise routes on the existing BGP peer after adding a Layer 3 VPN instance. [PR1237006](#)
- Migrate from syslog API to Errmsg API: `/src/junos/usr.sbin/ppmd`. [PR1284621](#)
- The BGP session might be stuck with high BGP OutQ value after GRES on both sides. [PR1323306](#)
- The VRF static route might not be exported when **route-distinguisher-id** is used on RR in a BGP Layer 3 VPN scenario. [PR1341720](#)
- The dynamic next-hop template cache does not shrink when the application frees the next-hop template and there are surplus templates in cache. [PR1346984](#)
- vFPC might continuously crash on vMX platform. [PR1364624](#)
- Ukern memory leak and core crash are seen in BGP environment. [PR1366823](#)
- The static route might persist even after its BFD session goes down. [PR1385380](#)
- The rpd might crash after issuing the **show route detail** operational command for RIP route. [PR1386873](#)
- Penultimate-hop router does not install BGP LU label, causing traffic to be discarded without notification. [PR1387746](#)
- IGMPv3/MLD membership requests might not work normally. [PR1389119](#)
- Unexpected packet loss might be seen for some multicast groups during failure recovery with both MoFRR and PIM automatic MBB join load-balancing features enabled. [PR1389120](#)
- FPC might crash when BGP multipath is configured with protection. [PR1389379](#)
- BGP IPv6 routes with IPv4 next-hop causes the rpd to crash. [PR1389557](#)
- All the BGP sessions will flap after switchover. [PR1391084](#)
- The ppmd on the Routing Engine might run with high CPU utilization after a Routing Engine switchover. [PR1392704](#)
- The rpd core files are generated on the backup Routing Engine during neighborship flap when using an authentication-key with more than 20 characters. [PR1394082](#)
- The rpd process might crash when **rp-register-policy** is configured with more than 511 terms. [PR1394259](#)

- The best and the second-best routes might have the same weight value if BGP PIC is enabled. [PR1395098](#)
- DMZ LINK BANDWIDTH - not able to aggregate bandwidth, when applying the policy. [PR1398000](#)
- The rpd soft core might be seen when Layer 2 VPN is used. [PR1398685](#)
- The rpd might crash in BGP setup with NSR enabled. [PR1398700](#)
- The rpd might crash when BGP **add-path send** is configured and NSR is enabled. [PR1401948](#)
- BGP router on the same broadcast subnet with its neighbors might cause IPv6 routing issue on the neighbor from other vendors. [PR1402255](#)
- EVPN multihoming MAC might not be installed by the remote PE device. [PR1403881](#)
- Memory leaks when labeled-isis transit routes are created as chain composite next-hop. [PR1404134](#)
- Extended traffic loss might be seen after link recovery when source-packet-routing is used on OSPF P2P links. [PR1406440](#)
- SBFD failure occurs with a special IP address like 127.0.0.1 under interface lo0. [PR1406631](#)
- The rpd crashes with BGP functions `bgp_peer_tcpwriteerror_gracefully`. [PR1410553](#)

Services Applications

- L2TP subscribers might be stuck in init state in a corner case. [PR1391847](#)
- The spd might crash when **any-ip** is configured in the 'from' clause of the NAT rule with the static translation type. [PR1391928](#)
- IP ToS bits are not copied to the outer IPsec header. [PR1398242](#)
- Invalid Layer 4 checksum might be observed on IPv4 packets generated by NAT64 with MS-DPC after translating fragmented IPv6 UDP/TCP packets. [PR1398542](#)
- The ICMPv6 packet with embedded IPv6 fragment might not be translated correctly to an IPv4 ICMP packet in a NAT64 with MS-DPC deployment. [PR1402450](#)
- Inconsistent content might be observed to the access line information between ICRQ and PPPoE message. [PR1404259](#)
- The stale si- logical interface might be seen when L2TP subscribers with duplicated prefixes or framed-route log in. [PR1406179](#)
- The kmd process might crash on MX/ACX platforms when IKEv2 is used. [PR1408974](#)

Subscriber Access Management

- The subscribers might be stuck in terminating state if RADIUS redirect is used. [PR1376265](#)
- Multiple IPv6 IANA addresses are assigned for one session in a IPv6 PD binding failure scenarios. [PR1384889](#)
- Dual-stacked DHCPv6-PD client connection terminated after commit when RADIUS address assignment is not defined within the range of a local pool. [PR1401839](#)

- The authd crash might be seen due to a memory corruption issue. [PR1402012](#)
- Adding a firewall filter service through the **test aaa** command causes a crash in dfwd. [PR1402051](#)
- JSRC used RADIUS service accounting protocol instead of JSRC for SRC installed service. [PR1403835](#)
- Continuous log message **authd[18454]: %DAEMON-3-LI: liPollTimerExpired returned 0** is seen. [PR1407923](#)

User Interface and Configuration

- The **show configuration** and **rollback compare** commands are causing high CPU usage. [PR1407848](#)

VPNs

- The receivers belonging to a routing instance might not receive multicast traffic in an Extranet next-generation MVPN scenario. [PR1372613](#)
- The **accept-remote-source** statement configured on the core interface might cause traffic outage. [PR1375716](#)
- High rpd CPU utilization on the backup Routing Engine might be observed in an MVPN with NSR scenario. [PR1392792](#)
- The rpd process crashes when the LSP template for a provider tunnel is changed. [PR1395353](#)
- Downstream interface is not removed from multicast route after getting PIM prune. [PR1398458](#)
- Routes with multiple communities might be rejected in an inter-AS next-generation MVPN scenario. [PR1405182](#)
- With rosen MVPN configuration with data-mdt, the **show pim mdt data-mdt-limit instance <interface name>** with family option causes high CPU usage of the rpd. [PR1405887](#)

SEE ALSO

[What's New | 88](#)

[What's Changed | 111](#)

[Known Limitations | 121](#)

[Open Issues | 125](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 195](#)

Documentation Updates

IN THIS SECTION

- [Spanning Tree Protocol User Guide | 195](#)

This section lists the errata and changes in Junos OS Release 19.1R3 documentation for MX Series.

Spanning Tree Protocol User Guide

- Documentation on configuring Layer 2 Protocol Tunneling (L2PT) for spanning-tree protocols on MX Series and ACX Series routers has been removed from the [Spanning-Tree Protocols User Guide](#) and merged into [Layer 2 Protocol Tunneling](#) in the [Ethernet Switching User Guide](#).

SEE ALSO

[What's New | 88](#)

[What's Changed | 111](#)

[Known Limitations | 121](#)

[Open Issues | 125](#)

[Resolved Issues | 142](#)

[Migration, Upgrade, and Downgrade Instructions | 195](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.1 | 196](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 197](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 199](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 201](#)

- Upgrading a Router with Redundant Routing Engines | 201
- Downgrading from Release 19.1 | 202

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 18.3R1 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 19.1

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 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 Instal 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-x86-32-19.1R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.1R3.9-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-x86-32-19.1R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.1R3.9-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](#).

NOTE: After you install a Junos OS Release 19.1 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-19.1R3.9-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-19.1R3.9-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 19.1 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 19.1

To downgrade from Release 19.1 to another supported release, follow the procedure for upgrading, but replace the 19.1 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		88
What's Changed		111
Known Limitations		121
Open Issues		125
Resolved Issues		142
Documentation Updates		195

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New](#) | [203](#)
- [What's Changed](#) | [206](#)
- [Known Limitations](#) | [207](#)
- [Open Issues](#) | [209](#)
- [Resolved Issues](#) | [212](#)
- [Documentation Updates](#) | [216](#)
- [Migration, Upgrade, and Downgrade Instructions](#) | [216](#)

These release notes accompany Junos OS Release 19.1R3 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

- [What's New in Release 19.1R3 | 203](#)
- [What's New in Release 19.1R2 | 203](#)
- [What's New in Release 19.1R1 | 204](#)

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

NOTE: For information about NFX product compatibility, see [NFX Product Compatibility](#).

What's New in Release 19.1R3

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 19.1R3.

What's New in Release 19.1R2

There are no new features or enhancements to existing features for NFX Series in Junos OS Release 19.1R2.

What's New in Release 19.1R1

Hardware

- **xDSL SFP modules (NFX250 NextGen)**—Starting in Junos OS Release 19.1R1, NFX250 NextGen devices support xDSL SFPs. The xDSL SFPs are supported on the SFP and SFP+ ports on the devices. Note that the xDSL SFPs are not supported on the extension modules. The xDSL SFP supports ADSL2/2+ and VDSL2.

[See [ADSL2 and ADSL2+ Interfaces on NFX250 \(NextGen\) Devices](#).]

[See [ADSL2 and ADSL2+ Interfaces on NFX250 Devices](#).]

[See [VDSL2 Interfaces on NFX250 Devices](#).]

Application Security

- **Application Quality of Experience (AppQoE) on NFX150 dual CPE deployments**—Starting in Junos OS Release 19.1R1, you can configure Application Quality of Experience (AppQoE) on NFX150 dual CPE deployments. AppQoE effectively prioritizes, segregates, and routes business-critical applications traffic without compromising performance or availability.

[See [Application Quality of Experience on NFX Devices](#).]

- **AppQoE scaling support (NFX250)**—Starting in Junos OS Release 19.1R1, Application Quality of Experience (AppQoE) enforces the configuration limit for overlay paths, metric profiles, probe parameters, and SLA rules per profile when you configure application-specific SLA rules and associate the SLA rules to an APBR profile. If you configure more parameters than the allowed limit, an error message is displayed after you commit the configuration.

[See [Application Quality of Experience on NFX Devices](#).]

Architecture

- **Reoptimized architecture support (NFX250 NextGen)**—Starting in Junos OS Release 19.1R1, NFX250 devices support a reoptimized architecture, which enables you to use the Junos Control Plane (JCP) as the single point of management to manage all the components.

NOTE: For documentation purposes, the NFX250 devices that use this architecture are referred to as NFX250 NextGen.

Key components in the software include the JCP, Juniper Device Manager (JDM), Layer 2 dataplane, Layer 3 dataplane, and virtualized network functions (VNFs). The JDM functions in the background. Users cannot access the JDM directly.

[See [How to Configure NFX250 \(NextGen\)](#).]

Firewall User Authentication

- **Firewall user authentication (NFX150)**—Starting in Junos OS Release 19.1R1, pass-through firewall user authentication is supported on NFX150 devices. Pass-through authentication restricts users who attempt to access a resource in another zone using FTP, Telnet, HTTP, or HTTPS. If the traffic matches a security policy that specifies pass-through authentication, the user is required to provide login information. The device validates the username and password against the information stored in the local database or on an external authentication server. The device supports the external authentication servers RADIUS, LDAP, and SecurID.

[See [Integrated User Firewall](#).]

High Availability

- **High availability (NFX150)**—Starting in Junos OS Release 19.1R1, NFX150 devices support the high availability feature. You can configure a cluster of two NFX150 devices to act as primary and secondary devices for protection against device failures. The high availability feature supports Layer 2 and Layer 3 features in dual CPE deployments.

By default, the heth-0-0 interface functions as the control interface. One of the remaining front panel interfaces can be configured as the fabric interface. On the LAN, the active/backup mechanism is used. If the primary device fails, the secondary device takes over the operation. On the WAN, both active/active and active/backup mechanisms are supported.

[See [Chassis Cluster on NFX150](#).]

Performance modes

- **Performance modes (NFX150 and NFX250 NextGen)**—Starting in Junos OS Release 19.1R1, NFX150 and NFX250 NextGen devices provide the following three performance modes:
 - **Throughput mode**—Provides maximum resources (CPU and memory) for Junos software and remaining resources, if any, for third-party VNFs. 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 [NFX150 Feature Overview.](#)]

[See [NFX250 NextGen Overview.](#)]

Wireless WAN

- **LTE support in dual CPE deployments (NFX150)**—Starting in Junos OS Release 19.1R1, you can provide a backup WAN connection by configuring LTE modules on a pair of NFX150 devices operating in cluster mode.

[See [Configuring the LTE Module on NFX Devices.](#)]

SEE ALSO

[What's Changed | 206](#)

[Known Limitations | 207](#)

[Open Issues | 209](#)

[Resolved Issues | 212](#)

[Documentation Updates | 216](#)

[Migration, Upgrade, and Downgrade Instructions | 216](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.1R3 | 207](#)
- [What's Changed in Release 19.1R2 | 207](#)
- [What's Changed in Release 19.1R1 | 207](#)

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

What's Changed in Release 19.1R3

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 19.1R3 for NFX Series devices.

What's Changed in Release 19.1R2

Factory-Default Configuration

- **Plug-and-play configuration (NFX150 and NFX250 NextGen devices)**—Starting in Junos OS Release 19.1R2, the factory-default configuration is modified to include the secure router plug-and-play configuration. [PR1401704](#)

What's Changed in Release 19.1R1

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 19.1R1 for NFX Series devices.

SEE ALSO

What's New 203
Known Limitations 207
Open Issues 209
Resolved Issues 212
Documentation Updates 216
Migration, Upgrade, and Downgrade Instructions 216

Known Limitations

IN THIS SECTION

- [Security | 208](#)
- [Platform and Infrastructure | 208](#)
- [Virtual Network Functions \(VNFs\) | 208](#)

Learn about known limitations in Junos OS Release 19.1R3 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.

Security

- Starting in Junos OS Release 19.1R1, the TCP and ICMP RPM probes take the best-effort queue of the outgoing interface instead of the network control queue on NFX150 and NFX250 NextGen devices. As a workaround, configure a DSCP value, such as nc1, to make the RPM probes take the network control queue. [PR1329643](#)

Platform and Infrastructure

- The Routing Engine boots from the secondary disk when you:
 - Press the reset button, on the RCB front panel, while Routing Engine is booting up but before Junos OS is up.
 - Upgrade software, by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network.
 - Upgrade BIOS and the upgrade fails.
 - Reboot and the system hangs before Junos OS is up.

[PR1344342](#)

- Starting in Junos OS Release 18.4, NFX150 devices support two versions of disk layout. In the older version of the disk layout, you could upgrade or downgrade from Junos OS Release 18.4. With the new disk layout, a downgrade to releases later than Junos OS Release 18.4 is not possible. As a workaround, avoid operations that reformat the disk layout. [PR1379983](#)

Virtual Network Functions (VNFs)

- After you create or delete a VNF on NFX150 and NFX250 NextGen devices, the **request virtual-network-functions console vnf-name** command gives an error that the VNF domain is not found. VNFs are reachable through SSH in this state. [PR1433204](#)

SEE ALSO

[What's New | 203](#)

[What's Changed | 206](#)

[Open Issues | 209](#)[Resolved Issues | 212](#)[Documentation Updates | 216](#)[Migration, Upgrade, and Downgrade Instructions | 216](#)

Open Issues

IN THIS SECTION

- [Interfaces | 209](#)
- [High Availability | 210](#)
- [Platform and Infrastructure | 210](#)
- [Virtualized Network Functions \(VNFs\) | 211](#)

Learn about open issues in Junos OS Release 19.1R3 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 NFX Series devices, if the IRB interface configuration and DHCP service configuration on JDM are removed and rolled back while retaining the VLAN mapping to the IRB interface, the DHCP service fails to assign IP addresses to the corresponding VNF interfaces and the service chaining fails. As a workaround, remove the VLAN mapping to the IRB interface along with IRB and DHCP service configuration on JDM. [PR1234055](#)
- When you issue a **show interface** command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. [PR1306191](#)
- On NFX150 devices, when you reboot the fpc0 interface, a few error messages are seen in the VTY console. [PR1326487](#)
- Starting in Junos OS Release 18.3R1, the reboot time has increased for fpc0 and fpc1 interfaces on NFX150 devices. [PR1355527](#)
- On NFX Series devices, while configuring **vmhost vlans** using **vlan-id-list**, the system allows duplicate VLAN IDs in the VLAN ID list. [PR1438907](#)

- When a DHCP server assigns a conflicting IP address to the NFX device interfaces, the NFX device does not send a DHCP DECLINE message in response. [PR1398935](#)
- On NFX150 devices, only the CFM cells that are configured for MEP levels are exchanged across xDSL MEPs. Other MEP-level CFM packets are dropped, whereas for Ethernet all MD levels along with above level are exchanged. [PR1409576](#)
- On NFX150 devices, when the interface configuration has the encapsulation **flexible-ethernet-services** enabled on a 10-Gigabit Ethernet interface, traffic gets dropped. [PR1425927](#)

High Availability

- On an NFX150 high availability (HA) chassis cluster, when one of the nodes is down or rebooting, any configuration commit on the peer active node triggers a reboot of the active node. This leads to loss of network connectivity until any one of the nodes in the HA topology becomes active. [PR1427550](#)

Platform and Infrastructure

- Jumbo frames are not supported through OVS on NFX250 NextGen devices. [PR1420630](#)
- Starting in Junos OS Release 19.1R1, Linux bridge mode is not supported on NFX250 devices. [PR1410598](#)
- Starting in Junos OS Release 18.1R1, the file transfer rate from external media over the network to an NFX150 device is around 40–50 Mbps. [PR1290263](#)
- During FTP on NFX150 devices, the following error message appears: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. [PR1315605](#)
- On NFX Series devices, the IPSec-NM does not start when the device comes up after a reboot. This issue appears after the system reboots during an upgrade process. It is caused by an inconsistency state of the Docker engine storage during the upgrade process. As a workaround, reboot the device. [PR1439577](#)
- When the NFX250 devices are operating on Linux bridge mode, the memory might be insufficient to launch a CLI session from JDM. This results in generating multiple JDM core files while spinning up a vSRX VNF. As a workaround:
 1. Check whether the `/var/third-party/jdm-config/last_1048576kB_nr_hugepages_value` or `/var/third-party/jdm-config/last_2048kB_nr_hugepages_value` file is present on the hypervisor. If it is, then delete it.
 2. Reboot the device.
 3. Upgrade to the release where this issue is fixed, if not already upgraded.

[PR1440427](#)

- NFX Series devices running Junos OS Release 19.2R1 do not take the statically assigned MAC address from the CLI configuration. [PR1458554](#)
- On NFX250 devices, vector packet processing (VPP) is not loaded in dual CPE, and at times in single CPE. [PR1461238](#)
- On NFX Series devices, after a power outage, JDMD might become unresponsive due to the `/etc/hosts` file getting corrupted. [PR1477151](#)

Virtualized Network Functions (VNFs)

- When you issue the **show virtual-network-functions *vnf-name*** command, the system creates a defunct process due to the presence of `popen()` calls and `pclose()` calls that do not match. This issue is fixed in Junos OS Release 15.1X53-D497 onward by ensuring that `pclose()` calls match the `popen()` calls. [PR1415210](#)
- While instantiating a vSRX VNF, multiple JDM core files are generated. As a workaround, verify that the `/var/third-party/jdm-config/last_1048576kB_nr_hugepages_value` and `/var/third-party/jdm-config/last_2048kB_nr_hugepages_value` files exist on the hypervisor. If the files exist, then delete the files and reboot the device. [PR1440427](#)
- On NFX250 devices, JDM does not allow to spin up the VM if the VNF name contains the word JDM. JDM can be used as a substring for VNF names and it is not case sensitive. However, the VNF name must not be equal to JDM. For example, `jdm123`, `abcJDM`, and `abcJDM123` are valid VNF names, whereas `jdm`, `JDM`, `Jdm`, or `JDm` are not valid VNF names. As a workaround, do not use JDM (case insensitive) as part of the VNF name. [PR1463963](#)

SEE ALSO

[What's New | 203](#)

[What's Changed | 206](#)

[Known Limitations | 207](#)

[Resolved Issues | 212](#)

[Documentation Updates | 216](#)

[Migration, Upgrade, and Downgrade Instructions | 216](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3 | 212](#)
- [Resolved Issues: 19.1R2 | 213](#)
- [Resolved Issues: 19.1R1 | 215](#)

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

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

Resolved Issues: 19.1R3

Class-of-Service (CoS)

- On NFX150 devices, when you configure a CoS rewrite rule for the st0 interface, the CoS value is not applied on the corresponding forwarding class. Hence, CoS does not work as expected and network traffic is impacted. [PR1439401](#)

High Availability

- On an NFX150 high availability chassis cluster, the host logs updated in the system log messages might not show the correct timestamp. As a workaround, convert the UTC timestamp to the local time zone. [PR1394778](#)

Platform and Infrastructure

- On NFX250 devices, when you issue the **request support information** command, the configuration and counter data are missing for JDM. [PR1413674](#)
- If you are using **init-descriptor filename vsrx.xml** to upgrade the NFX Series devices, the upgrade process reverts the file to default and the JDM subsystem becomes unavailable. [PR1456900](#)
- Starting in Junos OS Release 19.2R1, NFX Series devices do not take the statically assigned MAC address from the CLI configuration. This issue was caused by an issue with the intel drivers. [PR1458554](#)
- On NFX250 devices, Virtual Port Peer (VPP) is not running on dual CPE and occasionally on single CPE. [PR1461238](#)
- On NFX Series devices, if there are any conditional groups, the l2cpd process might crash and generate a core file when interfaces are flapping and the LLDP neighbors are available. This issue might cause the

dot1x process to fail and all the ports have a short interruption at the time of the process crash. As a workaround, delete the conditional group in the device. [PR1431355](#)

- On an NFX250 device with a secondary disk, the device reboots to the secondary disk during a password recovery process. [PR1467569](#)
- When you upgrade NFX250 devices from Junos OS Release 15.1X53-D496 to Junos OS Release 18.4R2, the upgrade fails. [PR1468586](#)
- When you upgrade NFX250 devices from Junos OS Release 15.1X53-D470 to Junos OS Release 18.4R2, the IPSec-NM does not start when the device comes up after reboot. It is caused by inconsistency state of docker engine storage during upgrade. [PR1439577](#)
- After a power outage, JDMD is not responsive because the `/etc/hosts` file is corrupted. [PR1477151](#)

Interfaces

- On NFX150 devices, the WAN ports (heth-0-4 and heth-0-5) do not function properly if you remove a cable connected to these ports or flap the link. As a workaround, use one of the following options:
 - Flap the link again.
 - Enable or disable the interface from the CLI.

[PR1449278](#)

- When traffic goes through vSRX 3.0 instances, core files are generated and traffic is dropped. This issue might cause all interfaces to go down and the Packet Forwarding Engine might not come up. [PR1465132](#)

Virtualized network Functions (VNFs)

- NFX250 devices do not allow *jdm* (case-insensitive) as a VNF name. You can use *jdm* as part of the name. For example, *jdm123*, *abcJDM*, *abcJDM123* are valid VNF names, but *jdm*, *JDM*, *Jdm*, *JDm* are not valid VNF names. [PR1463963](#)

Resolved Issues: 19.1R2

Class-of-Service (CoS)

- In the NFX Series device configuration, traffic is being sent to the incorrect queue when configuring CoS with **forwarding-classes** class versus queue. The **forwarding-classes** class is not supported and is hidden. As a workaround, use **forwarding-classes** queue when you configure CoS. [PR1436408](#)

Interfaces

- On NFX250 devices using xDSL SFP transceivers on the fiber ports, the status of the transceiver is displayed under the **Adsl Status** field in the output of the **show interfaces int-name** command. If you hot-swap an xDSL SFP with another xDSL SFP on the same port, then the **Adsl Status** field is not displayed in the output of the **show interfaces** command. [PR1408597](#)
- 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](#)
- On NFX150 devices, FPC0 might not be online after an upgrade and a device reboot is required. [PR1430803](#)
- When you run the **show chassis fpc** or **show chassis fpc details** command, the **Temperature** field in the command output message is displayed as **Testing**. [PR1433221](#)
- 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](#)

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 might need more time to get an IP address assigned. The subscriber might remain in this state until its lease expires if it had previously bound with the address in the option 50. [PR1435039](#)

Platform and Infrastructure

- On an NFX250 device, the console is not accessible and JDM stops working. These issues occur because the libvirtd process stops responding. [PR1341772](#)
- On an NFX250 device, if the **idle-time out** parameter for a user login class on JDM is configured in minutes, the system considers the configured idle timeout value in seconds. The user is logged out based on the idle timeout value in seconds. [PR1435310](#)
- On NFX150 devices, the **show security dynamic-address** command does not work for port 3. [PR1448594](#)
- Version compare in phc might fail causing the phc to download the same image. [PR1453535](#)
- When applying firewall filters on lo0.0 on an NFX250 NextGen device, FPC0 disappears. [PR1448246](#)
- When the REST API receives several continuous HTTP requests, the REST service might become unresponsive. [PR1449987](#)

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 jnxlpSecTunMonInDecryptedPkts`
- `show snmp mib walk jnxlpSecTunMonLocalGwAddr`
- `show snmp mib walk jnxlpSecTunMonLocalGwAddrType`

[PR1386894](#)

- On NFX250 devices, the **request-load-configuration** command output from the device does not match with YANG modules for Junos OS Release 18.4. [PR1416106](#)

Virtual Network Functions (VNFs)

- When you downgrade from Junos OS Release 19.2 to Junos OS Release 18.4, the **show virtual-network-functions vnf-name** command does not display the VNF information. [PR1437547](#)

Resolved Issues: 19.1R1

NFX250

- Junos Device Manager (JDM) depends on the libvirtd daemon to manage the guest VMs through CLI. On NFX250 devices running Junos OS Release 19.1R1, the libvirtd daemon is inactive and the vjunos VM start up fails. This results in inband connectivity failure, guest VMs fails to start, and the console hangs. [PR1314945](#)

NFX150

- On NFX150 devices running Junos OS Release 19.1R1, software upgrade does not delete all images from the previous installation. This occupies 1 GB of storage per upgrade and leads to depletion of storage after several upgrades. [PR1408061](#)

SEE ALSO

[What's New | 203](#)

[What's Changed | 206](#)

[Known Limitations | 207](#)

[Open Issues | 209](#)

[Documentation Updates | 216](#)

[Migration, Upgrade, and Downgrade Instructions | 216](#)

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for NFX Series.

SEE ALSO

- [What's New | 203](#)
- [What's Changed | 206](#)
- [Known Limitations | 207](#)
- [Open Issues | 209](#)
- [Resolved Issues | 212](#)
- [Migration, Upgrade, and Downgrade Instructions | 216](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 216](#)
- [Basic Procedure for Upgrading to Release 19.1 | 217](#)

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.

NOTE: For information about NFX product compatibility, see [NFX Product Compatibility](#).

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 19.1

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 **bundle** package, only when so instructed by a Juniper Networks support representative.

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 19.1R3:

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 | 203](#)

[What's Changed | 206](#)

[Known Limitations | 207](#)

[Open Issues | 209](#)

[Resolved Issues | 212](#)

[Documentation Updates | 216](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- [What's New | 219](#)
- [What's Changed | 231](#)
- [Known Limitations | 236](#)
- [Open Issues | 238](#)
- [Resolved Issues | 240](#)

- Documentation Updates | 247
- Migration, Upgrade, and Downgrade Instructions | 248

These release notes accompany Junos OS Release 19.1R3 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 19.1R3 | 219
- What's New in 19.1R2 | 219
- What's New in 19.1R1 | 220

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

What's New in 19.1R3

There are no new features or enhancements to existing features for PTX Series Packet Transport Routers in Junos OS Release 19.1R3.

What's New in 19.1R2

There are no new features or enhancements to existing features for PTX Series Packet Transport Routers in Junos OS Release 19.1R2.

What's New in 19.1R1

Hardware

- **QFX10000-60S-6Q line card (PTX10008 and PTX10016 routers)**—Starting with Junos OS Release 19.1R1, the QFX10000-60S-6Q line card provides 60 SFP+ ports that can be configured at either 10-Gbps or 1-Gbps, and six flexible configuration ports for 100-Gbps and 40-Gbps. By default, all the ports will be in the 10-Gbps mode.

Of the six flexible configuration ports, two ports have QSFP28 sockets that support either 100-Gbps, 40-Gbps, or 10-Gbps speeds. The remaining four ports have QSFP+ sockets that can be configured as either a native 40-Gbps port or four 10-Gbps ports using a breakout cable. With breakout cables, the line card supports a maximum of 84 logical 10-Gigabit Ethernet ports.

- **Support for 40-Gbps ports to operate at 10-Gbps or 1-Gbps speed (PTX1000, PTX10008, and PTX10016)**—Starting in Junos OS Release 19.1R1, you can use the Mellanox 10-Gbps pluggable adapter (QSFP+ to SFP+ adapter or QSA; 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-Gigabit Ethernet port to a 10-Gbps port or a 1-Gbps port. You can then plugin a 10-Gbps SFP+ transceiver or a 1-Gbps SFP transceiver into the QSA adapter, which is inserted into the QSFP or QSFP+ ports of the PTX1000 router or the PTX10K-LC1101 and PTX10K-LC1102 line cards of the PTX10008 and PTX10016 routers.

[See [PTX1000 Transceivers](#), [PTX10008 Transceivers](#), and [PTX10016 Transceivers](#).]

Authentication, Authorization, and Accounting (AAA)

- **Support for SFTP global disablement (PTX Series)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the statement **sftp-server** at the **[edit system services ssh]** hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections were globally enabled by default.

[See [Configuring sftp-server](#).]

Class of Service

- **Support for class of service (CoS) (PTX10001-20C)**—Starting in Junos OS Release 19.1R1, PTX10001-20C routers support class of service (CoS) functionality for IPv6 traffic. Only default and custom INET, DSPC, and DSPC IPv6 classifiers are supported. Rewrite rules are not supported.

CoS is the assignment of traffic flows to different service levels. Service providers can use router-based CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows.

[See [CoS Features and Limitations on PTX Series Routers](#).]

Forwarding and Sampling

- **Support for tracking static RPM routes across multiple next hops (PTX Series)**—Starting in Junos OS Release 19.1R1, you can use **rpm-tracking** to track up to 16 next hops for RPM-controlled static routes. This feature supports both IPv4 and IPv6 static rpm-tracked routes, and extends the single-hop [rpm-tracking](#) introduced in Junos OS Release 18.4.

[See [show route rpm-tracking](#).]

- **Support for using IP addresses in a SR-TE LSP segment list (PTX series)**—Starting in Junos OS Release 19.1R1, you can use IP addresses (IPv4 or IPv6) for next hops in a segment routing traffic engineering (SR-TE) list of label-switched paths (LSPs). This work extends the support for traffic steering based on a segment routing policy that was introduced in Junos OS Release 17.4R1, wherein the controller can specify a segment routing policy consisting of multiple paths to steer labeled or IP traffic.

[See **auto-translate** in [segment-list](#) and **retry-timer** in [source-packet-routing](#).]

Junos Telemetry Interface

- **Support for the Junos telemetry interface (JTI) (PTX10002)**—Starting with Junos OS Release 19.1R1, you can provision sensors through the Junos telemetry interface to export telemetry data for several network elements without involving polling. You can stream data through UDP or gRPC.

Only the following sensors are supported on PTX10002 routers:

- Physical interfaces statistics
- Label-switched-path (LSP) statistics
- Network processing unit (NPU) memory

- NPU memory utilization
- CPU memory

To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level.

To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models.

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

- **Transit SPRING sensor support on Junos telemetry interface (JTI) (PTX3000 and PTX5000 with FPC2)**—Starting in Junos OS Release 19.1R1, JTI sensor support is available for Source Packet Routing in Networking (SPRING), also known as segment routing. Segment routing is a control-plane architecture that enables an ingress router to steer a packet through a specific set of nodes and links in the network.

Segment routing statistics must first be enabled before the sensor can be configured and statistics streamed to an outside collector by means of JTI.

To enable collection of statistics, configure **set protocols isis source-packet-routing sensor-based-stats per-sid ingress** through the Junos CLI.

To configure the sensor for statistics to be issued to an outside collector, include the following path for either UDP (native) or gRPC streaming:

- **/junos/services/segment-routing/sid/usage/**

For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

To provision the 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 [sensor \(Junos Telemetry Interface\)](#), [Configuring a Junos Telemetry Interface Sensor \(CLI Procedure\)](#), and [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **RSVP interface OpenConfig model support and self-ping logs on Junos telemetry interface (JTI) (PTX10003)**—Starting in Junos OS Release 19.1R1, JTI sensor support is enhanced for RSVP interfaces to include delivery of more statistics. The level of support is equivalent to the output delivered when using the **show rsvp interface detail** operational mode command.

To configure the sensor for statistics to be issued to an outside collector, include the following path for gRPC streaming:

- `/network-instances/network-instance/mpls/signaling-protocols/rsvp-te/interface-attributes/interfaces/interface/*`

To provision the 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 [gRPC Services for Junos Telemetry Interface](#) and [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **Support for LSP statistics on Junos telemetry interface (JTI) (PTX10001-20C)**—Starting with Junos OS Release 19.1R1, you can provision the LSP statistics sensor `/junos/services/label-switched-path/usage/` to monitor per-MPLS LSP statistics on the PTX10001-20C router and export telemetry data through JTI to external collectors. You can stream data at configurable intervals through gRPC without involving polling.

JTI support is only for RSVP LSPs.

Statistics that are streamed are similar to the output displayed by the operational mode command **show mpls lsp bypass statistics**.

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).

To enable statistics for export from the Junos OS, include the **sensor-based-stats** statement at the `[edit protocols mpls]` hierarchy level.

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

- **Export of data associated with the Junos kernel through Junos Telemetry Interface (JTI) (PTX Series)**—Starting in Junos OS Release 19.1R1, you can export data associated with the Junos kernel through remote procedure calls (gRPC) and JTI. Kernel telemetry data includes information on Veriexec state, graceful Routing Engine switchover (GRES), in-service software upgrade (ISSU), and Routing Engine ifstate. Junos kernel sensors can be used by device monitoring and network analytics applications to provide insight into the health status of the Junos kernel.

Junos kernel sensors introduced in Junos OS Release 19.1R1 support both periodical and ON_CHANGE streaming. The following Junos kernel resource paths support periodical streaming only:

- `/junos/kernel-ifstate/dead-ifstates-cnt`
- `/junos/kernel-ifstate/alive-ifstates-cnt`
- `/junos/kernel-ifstate/delayed-unrefs-cnt`
- `/junos/kernel-ifstate/delayed-unrefs-max`

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

Layer 3 Features

- **Support for Layer 3 unicast features (PTX10001)**—Starting in Junos OS Release 19.1R1, PTX10001 routers support the following Layer 3 forwarding features for unicast IPv4 and IPv6 traffic:
 - Basic IPv6 Forwarding
 - Virtual router (VRF-lite) for both IPv4 and IPv6
 - Layer 3 subinterfaces support for both IPv4 and IPv6
 - VRF-lite, subinterfaces and IPv6 forwarding support on link aggregation group (LAG)
 - Statistics support for Layer 3 subinterfaces
 - 32-way equal-cost multipath (ECMP)
 - Centralized Bidirectional Forwarding Detection (BFD)
 - IPv4 Layer 3 protocols such as
 - OSPF
 - IS-IS
 - BGP
 - IPv6 Layer 3 protocols such as
 - OSPFv3
 - IS-ISv6
 - BGPv6

MPLS

- **Flexible MPLS label stack depth (PTX Series routers with third-generation FPCs)**—Currently, Junos OS supports push of up to a maximum of five labels per component of the next hop chain, even though the underlying device capability can be higher. Starting in Junos OS Release 19.1R1, the device capability of pushing more than 5 labels can be leveraged for features, such as, segment routing traffic engineering (TE) LSPs and RSVP-TE pop-and-forward LSPs.

The number of labels that can be pushed for an MPLS next hop is the number of labels the device is capable of pushing, or the maximum-labels configured under **family mpls** of the outgoing interface, whichever is smaller.

[See [Configuring the Maximum Number of MPLS Labels, maximum-labels](#).]

- **Support for MPLS ping and traceroute for segment routing (PTX Series)**—Starting in Junos OS Release 19.1R1, MPLS ping and traceroute are supported for segment routing for protocols IS-IS and OSPF over IPv4. This feature also supports ECMP traceroute for protocols IS-IS and OSPF.

In Junos OS Release 19.1R1, MPLS ping and traceroute for segment routing supports IPv4 IGP-Prefix segment FEC validation. FEC validation for IGP-Adjacency Segment ID is not supported.

[See [ping mpls segment routing isis](#), [ping mpls segment routing ospf](#), [traceroute mpls segment-routing ospf](#), [traceroute mpls segment-routing isis](#).]

- **Enhancements to MPLS for LSP path selection (PTX Series)**—Starting in Junos OS Release 19.1R1, the following enhancements to MPLS have been added for LSP path selection and optimization:

- Earlier when LSP active paths were modified, the LSP path gets cleared and gets resignaled immediately. From Junos OS Release 19.1R1 onward, if a secondary path is available, then Junos OS selects the secondary path as active, clears and resignals the primary path after the expiry of the **optimize-hold-dead-delay** timer. When the primary LSP path is established, the **revert-timer** gets started. After the **revert-timer** expires, the primary LSP path becomes active.

If the primary LSP path is not active with **revert-timer** on and when there is a change to the primary LSP path, then the LSP path gets cleared and resignaled immediately. When the primary LSP path is established, the revert-timer gets restarted.

- Earlier, if there was any Constrained Shortest Path First (CSPF) failure then the current LSP path becomes invalid because it did not match with the configured constraints. In this case, the current LSP path gets cleared immediately. From Junos OS Release 19.1R1 onward, if a secondary LSP path is available, then Junos OS selects the secondary LSP path as active and clears the primary path after the expiry of the **optimize-hold-dead-delay** timer.
- The CLI command **no-bypass-statistics-polling** added under the [**edit protocols mpls statistics**] hierarchy now provides information on bypass LSP statistics.
- A new CLI command **delay** has been introduced under the [**edit protocols mpls optimize-adaptive-teardown**] hierarchy and the value for delay is in the range of (3..65535 seconds). When the **adaptive-teardown** configuration is triggered, the **delay** CLI command further delays the tearing down of old optimized LSP paths based on the configured value.

[See [statistics \(Protocols MPLS\)](#), [optimize-adaptive-teardown](#).]

- **Control transport address used for targeted LDP session (PTX Series)**—Currently, only the router ID or interface address is used as the LDP transport address. Starting in Junos OS Release 19.1R1, you can configure any other IP address as the transport address of targeted LDP sessions, session groups, and interfaces. This new configuration is applicable only for configured LDP neighbors that have Layer 2 circuit, MPLS, and VPLS adjacencies.

This feature is beneficial when you have multiple loopback interface addresses, and different IGPs associated with LDP interfaces, and you can control the session established between targeted LDP neighbors with the configured transport address.

[See [Control Transport Address Used for Targeted-LDP Session](#).]

- **MPLS egress traffic statistics for label IS-IS routes at ingress device (PTX Series)**—Currently, sensors are available for collecting segment routing statistics for MPLS transit traffic, which is MPLS-to-MPLS in nature. Starting in Junos OS Release 19.1R1, additional sensors are introduced to collect segment

routing statistics for MPLS egress traffic at the ingress provider edge (PE) device, which is IP-to-MPLS in nature.

With this feature, you can enable sensors for label IS-IS segment routing egress traffic only, and stream the statistics to a gRPC client.

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

- **Policy-based multipath routes (PTX Series)**—In segment routing networks with multiple protocols in the core, you can combine segment routing traffic-engineered (SR-TE) LDP routes and SR-TE IP routes to create a multipath route that is installed in the routing information base (also known as routing table). You can resolve BGP service routes over the multipath route through policy configuration and steer traffic differently for different prefixes.

[See [Policy-Based Multipath Routes Overview](#).]

- **Use of SID labels as first hop for resolving non-colored static segment routing LSPs (PTX Series)**—Currently, for a static non-colored segment routing traffic-engineered LSP to be usable, the first hop of the segment list must be an IP address. Only the second to *n*th hop could be segment identifier (SID) labels. Starting in Junos OS Release 19.1R1, this requirement does not apply. You can now configure SID labels as the first hop in the segment list.

With this configuration, static non-colored segment routing LSPs are resolved using MPLS fast reroute (FRR) and weighted equal-cost multipath. Without this configuration, by default, the LSPs are resolved using IP address.

[See [Static Segment Routing Label Switched Path](#).]

- **Support of install statement for segment routing LSPs (PTX Series)**—The *install destination-prefix* statement which is currently supported at the `[edit protocols mpls label-switched-path lsp-name]` and `[edit protocols mpls static-label-switched-path lsp-name ingress]` hierarchy levels is now also supported at the `[edit protocols source-packet-routing source-routing-path lsp-name]` hierarchy level for both colored and non-colored static segment routing label-switched paths (LSPs).

You can associate one or more prefixes with a segment routing LSP using the **install** statement. When the LSP is up, all the prefixes are installed as entries into the **inet.3** or **inet6.3** routing table.

[See [install \(Protocols MPLS\)](#).]

Network Management and Monitoring

- **sFlow performance improvements (PTX Series)**—Starting in Junos OS Release 19.1R1, the following improvements have been added to the sFlow technology feature:
 - For MX Series, PTX Series, and QFX Series, you can configure forwarding class and DSCP values per collector.
 - For PTX Series and QFX Series, you can configure IPv6 addresses for the **source-ip** and **agent-id**.
 - Enhancements are made to the following CLI commands: **show sflow collector**, **show sflow collector address *ip-address***, and **show sflow interface**.

[See [Understanding How to Use sFlow Technology for Network Monitoring](#), `collector`, `agent-id`, `source-ip`, `show flow collector`, and `show flow interface`.]

Routing Policy and Firewall Filters

- **Support for IPv6 firewall filters (PTX100020C)**— Starting with Junos OS Release 19.1R1, you can configure a firewall filter with match conditions for IPv6 traffic (ingress direction only). You configure firewall filters under the `[edit firewall]` hierarchy level.

This feature was previously supported in an "X" release of Junos OS.

[See [IPv6 Firewall Filter Match Conditions and Actions \(PTX10001-20C\)](#).]

Routing Protocols

- **Support for BGP graceful shutdown (PTX Series)**— Starting in Junos OS Release 19.1R1, graceful traffic migration from one BGP next hop to another is supported, without traffic interruption. Also, BGP administrative shutdown communication can be sent to the BGP peer.

You can configure both `graceful-shutdown` and `shutdown` statements at the `[edit protocols bgp]`, `[edit protocols bgp group group-name]`, and `[edit protocols bgp group group-name neighbor address]` hierarchy levels.

NOTE: Graceful shutdown is disabled by default.

[See: [graceful-shutdown \(Protocols BGP\)](#), [shutdown \(Protocols BGP\)](#).]

- **Support for export of BGP Adjacency-RIB-Out through BGP Monitoring Protocol (BMP) (PTX Series)**— Starting in Junos OS Release 19.1R1, BMP is enhanced to support route monitoring of pre and post `rib-out` policy.

You can configure `post-policy` and `pre-policy` under `rib-out` statement at `[edit protocols bgp bmp]`, `[edit protocols bgp group group-name bmp]`, and `[edit protocols bgp group group-name neighbor address bmp]` hierarchies.

NOTE: The default monitoring mode of `rib-out` is `pre-policy`.

[See: [Understanding the BGP Monitoring Protocol](#).]

- **Support for configurable SRGB used by SPRING in OSPF protocols (PTX Series)**— Starting in Junos OS Release 19.1R1, 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 19.1R1, 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

[See [source-packet-routing \(Protocols IS-IS and OSPF\)](#).]

- **Support for BGP advertising aggregate bandwidth across external BGP links for load balancing (PTX Series)**—Starting in Junos OS Release 19.1R1, BGP uses a new link bandwidth extended community, **aggregate-bandwidth**, to advertise an aggregated bandwidth of multipath routes across external links. BGP calculates the aggregate of multipaths that have unequal bandwidth allocation and advertises the aggregated bandwidth to external BGP peers. A threshold to the aggregate bandwidth can be configured to restrict the bandwidth usage of a BGP group. In earlier Junos OS releases, a BGP speaker receiving multipaths from its internal peers advertised the link bandwidth associated with the active route.

To advertise the aggregated bandwidth of multipath routes and to set a maximum threshold, configure a policy with the **aggregate-bandwidth** and **limit bandwidth** actions at the **[edit policy-options policy-statement name then]** hierarchy level.

[See [Advertising Aggregate Bandwidth Across External BGP Links for Load Balancing Overview](#).]

- **Support for policy-based allocation for IPv4 BGP-labeled unicast (PTX Series)**—Starting in Junos OS Release 19.1R1, this feature supports:
 - Allocating policy-based labels for IPv4 BGP-LU prefixes in per-prefix label allocation mode
 - 1:1 mapping between prefixes and labels
 - Map policy for labels
 - Fallback actions of dynamic and reject for handling error conditions

[See [policy-options, route-filter-list](#).]

- **Support for BGP link-state distribution with SPRING extensions (PTX Series)**—Starting in Junos OS Release 19.1R1, BGP link-state extensions export segment routing topology information to software-defined networking controllers. Controllers can get the topology information by either being a part of an interior gateway protocol (IGP) domain or through BGP link-state distribution.

BGP link-state distribution is supported on inter-domain networks and provides a scalable mechanism to export the topology information. This feature benefits networks that are moving to source packet routing in networking (SPRING) but also have RSVP deployed, and continue to use both SPRING and RSVP in their networks.

In this release, OSPF link-state protocol is supported, which pushes SPRING information to the BGP link-state address family.

[See [Link-State Distribution Using BGP Overview](#).]

- **Scalability for LDP-over-RSVP and BGP labeled unicast services (PTX Series)**—Starting in Junos OS Release 19.1R1, this feature enhances RPD to produce the chain next hop for various services. The RPD is enhanced to build a translation layer between RIB and FIB to segment multi protocol transport tunnels.

Segmentation happens as follows:

- Only LDP, RSVP, LDP-over-RSVP, and LDP-over-RSVP-over-BYPASS ingress tunnels are considered for segmentation.
- Segmentation does not happen if there is only one label in the stack.
- Segmentation happens at the application boundary. A next hop with two LDP labels in its stack or two RSVP labels will not be split into two next hops with one label each.

Any route resolution over LDP or LDP-over-RSVP is changed from INH->FNH to CNH->INH->CNH->FNH in kernel and Packet Forwarding Engine, and for LDP routes in INET.3, where the indirect next hop (INH) is an application installed in direction toward the final next hop (FNH). Any segmented stack introduces the composite chain next hop (CNH), where the segmented portion of the label stack precedes an INH or an FNH. The chain is collapsed and the resulting label stack is encoded in the packet header by the hardware before forwarding the packet.

By chaining labels instead of stacking them, PTX Series memory is made available for FNH label operations, as well as CNH by grouping CNHs within the same unicast next hop (for ECMP) based on the label space identifier.

The following applications are supported:

- Transit
 - LBGP stitching with LDP over RSVP
- Ingress
 - 6PE BGP-V6-Route->LBGP(Explicit V6 NULL label) over LDP over RSVP
 - 4PE BGP-V4-Route->LBGP(Explicit V4 NULL label) over LDP over RSVP
 - BGP-L3VPN over LDP over RSVP
 - BGP-V6-VPN over LDP over RSVP
- BGP route with indirection resolving over LDP over RSVP
 - IBGP-V4-ROUTE over LDP over RSVP
 - IBGP-V6-ROUTE over LDP over RSVP

[See [Tunneling LDP LSPs in RSVP LSPs Overview](#), [BGP Route Resolution Overview](#).]

Services Applications

- **Support for IPv4 and IPv6 inline active flow monitoring on IRB interfaces (PTX1000)**—Starting in Junos OS Release 19.1R1, you can perform inline active flow monitoring for IPv4 and IPv6 traffic on integrated routing and bridging (IRB) interfaces. Both IPFIX and version 9 templates are supported.

IRB interfaces enable a switch to identify packets that are being sent to local addresses to be bridged whenever possible and to be routed only when required. Switching or bridging uses fewer layers of processing than routing, thus reducing the number of address lookups.

[See [Inline Active Flow Monitoring on IRB interfaces](#).]

- **Support for automatic restart of Two-Way Active Measurement Protocol (TWAMP) Client (PTX Series)**—Starting in Junos OS Release 19.1R1, the TWAMP client restarts automatically after a network failure, a configuration change, or an IP connectivity issue. However, for the client to reconnect to the TWAMP server automatically, you must use 0 as the *test-count* value in the **set rpm twamp client control-connection test-count** command. Also, at the TWAMP server side, the default value of *max-connection-duration* in the **set rpm twamp server max-connection-duration** must also be 0. You can display the test results after the network recovers, or after the server is reachable, by using the **set services rpm twamp client control-connection c1 persistent-results** command.

[See [Understanding TWAMP Auto-Restart](#).]

- **Port mirroring support for the IPv6 address family (PTX10001)**—Starting in Release 19.1R1, Junos OS supports port mirroring on the PTX10001 router for the IPv6 address family. Port mirroring copies packets entering or exiting a port and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns as correlating events. The PTX Series routers support the IPv6 (inet6) address family only.

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

SEE ALSO

[What's Changed | 231](#)

[Known Limitations | 236](#)

[Open Issues | 238](#)

[Resolved Issues | 240](#)

[Documentation Updates | 247](#)

[Migration, Upgrade, and Downgrade Instructions | 248](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.1R3-S4 | 232](#)
- [What's Changed in 19.1R3 | 232](#)
- [What's Changed in 19.1R2 | 232](#)
- [What's Changed in 19.1R1 | 234](#)

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

What's Changed in Release 19.1R3-S4

Infrastructure

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

What's Changed in 19.1R3

General Routing

- **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.

Interfaces and Chassis

- **Updates to the show interfaces and show policer commands (PTX Series)**—Starting in Junos OS Release 19.1R3, 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_`) although PTX Series routers do not support the default ARP policer.

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.

What's Changed in 19.1R2

General Routing

- **User confirmation prompt for configuring the sub options of request vmhost commands (MX Series and PTX series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm a **[yes,no]** for the sub options also.
 - **request vmhost reboot**
 - **request vmhost poweroff**
 - **request vmhost halt**

In previous releases, the confirmation prompt was available for only the main options.

Network Management and Monitoring

- **The `show system schema` command and `<get-yang-schema>` RPC require specifying an output directory (PTX Series)**—Starting in Junos OS Release 19.1R2, when you issue the `show system schema` operational mode command in the CLI or execute the `<get-yang-schema>` RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the `<output-directory>` element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.

Routing Protocols

- **Change in the default behavior of `advertise-from-main-vpn-tables` configuration statement**—BGP now advertises EVPN routes from the main `bgp.evpn.0` table. You can no longer configure BGP to advertise the EVPN routes from the routing instance table. In earlier Junos OS Releases, BGP advertised EVPN routes from the routing instance table by default.

[See [advertise-from-main-vpn-tables](#).]

Software-Defined Networking

- **Increase in the maximum value of `delegation-cleanup-timeout` (PTX Series)**—You can now 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](#).]

What's Changed in 19.1R1

EVPNs

Interfaces and Chassis

- **Support to get optics loopback status for QSFP-100GE-DWDM2 transceivers (PTX Series)**—Starting in Junos OS Release 19.1R1, you can get the optics loopback status of QSFP-100GE-DWDM2 transceivers along with the regular Ethernet loopback status by issuing the `show interfaces interface-name` or `show interfaces interface-name brief` command. The new output field **Optics Loopback** is added under **Link-level type** when the `show interfaces interface-name` CLI command is executed.
- **New XML tag element `<lacp-hold-up-state>` added in `show lacp interfaces` XML display (PTX Series)**—Starting in Junos OS Release 19.1R1, the `show lacp interfaces | display xml` command displays a new XML tag element `<lacp-hold-up-state>`. The `<lacp-hold-up-state>` displays the time interval an

interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold-up information for all interfaces was in a single **<lacp-hold-up-information>** XML tag. Now, for each interface it is displayed in a separate **<lacp-hold-up-information>** XML tag.

- **Support for creating Layer 2 logical interface independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Releases 18.4R1, 18.4R2, 19.1R1, and later, PTX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

MPLS

- **New debug statistics counter (PTX Series)**—The **show system statistics mpls** command has a new output field, called **Packets dropped, over p2mp composite nexthop**, to record the packet drops over composite point-to-multipoint next hops.

Network Management and Monitoring

- **NETCONF <kill-session> operation returns different values in <rpc-error> when the session identifier is equal to the current session ID (PTX Series)**—Starting in Junos OS Release 19.1R1, when you execute the **<kill-session>** NETCONF operation and the session identifier is equal to the current session ID, the values of the **<error-type>** and **<error-tag>** elements in the resulting **<rpc-error>** are **application** and **invalid-value**, respectively. In earlier releases, the **<error-type>** and **<error-tag>** values are **protocol** and **operation-failed**.

[See [<kill-session>](#).]

- **sysName.0 MIB object displays the fully qualified domain name (PTX Series)**—Starting in Junos OS Release 19.1R1, the sysName.0 MIB object displays the fully qualified domain name. That is, if the hostname and domain name are configured on the system, both show up for the sysName.0 MIB object: **host-name.domain-name**. Previously, only the hostname showed up.

[see [show snmp mib](#).]

- **Change in error severity (PTX10016)**—Starting in Junos OS Release 19.1R1, on PTX10016 routers, the severity of the FPC error, shown in the syslog as **PE Chip::FATAL ERROR!! from PE2[2]: RT: Clear Fatal if it is detected LLMEM Error MEM:llmem, MEMTYPE: 1**, is changed from fatal to non fatal (or minor). In case of this error, only a message is displayed for information purposes. To view the error details, you can use the commands **show chassis fpc errors** and **show chassis errors active**.

[See [show chassis fpc errors](#).]

Services Applications

- **Support for enabling hardware timestamping of RPM probe messages (PTX Series)**—Starting in Junos OS Release 19.1R1, PTX Series routers support timestamping of RPM probe messages on the Packet Forwarding Engine. The following configuration statements at the `[edit services rpm probe owner test test-name]` hierarchy level are supported:
 - **hardware-timestamp**—To enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor.
 - **one-way-hardware-timestamp**—To enable timestamping of RPM probe messages for one-way delay and jitter measurements.

These features are supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types.

User Interface and Configuration

- **Options for monitor traffic interfaces statement added (PTX Series)**—Starting in Junos OS Release 19.1R1, the options **write-file** and **read-file** under the **monitor traffic** command are included in the visible CLI.
[See [monitor traffic](#).]

SEE ALSO

What's New 219
What's Changed 236
Open Issues 238
Resolved Issues 240
Documentation Updates 247
Migration, Upgrade, and Downgrade Instructions 248

Known Limitations

IN THIS SECTION

- [General Routing | 237](#)

Learn about known limitations in this release 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 PTX switches with FPC3, if the filter action of policer and count are configured in one same filter term, then it might cause the policer counter to be higher than the actual count. [PR1089330](#)
- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. [PR1268678](#)
- The Routing Engine boots from the secondary disk when you:
 - Press the reset button, on the RCB front panel, while the Routing Engine is booting up but before Junos OS is up.
 - Upgrade software, by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network.
 - Upgrade BIOS and the upgrade fails.
 - Reboot and the system hangs before Junos OS is up. [PR1344342](#)
- PTX1000 sFlow sampling output has different VLAN priority in extended switch data fields with the same dual-tag configuration when egress sampling is configured. This issue is dependent on the sequence in which sampling and MAC rewrite happen. In MX Series MAC rewrite occurs after sampling and in the case of PTX Series sampling happens after MAC rewrite. [PR1387468](#)
- The **request vmhost power-off** command does not actually power off the system in the latest releases. It only does a reboot and the system comes back up. [PR1393061](#)
- Frames that cannot be fragmented and are larger than the outgoing interface MTU size are dropped; however, the **show interface statistics extensive** output might not show these dropped frames against output errors and MTU errors. [PR1408576](#)
- In case of stacked VLAN Series tagging, VLAN tagged frames counter is not supported for LC1101, LC1102, and LC1103 series card in PTX Series. [PR1412987](#)

SEE ALSO

[What's New | 219](#)

[What's Changed | 231](#)

[Open Issues | 238](#)

[Resolved Issues | 240](#)

[Documentation Updates | 247](#)

Open Issues

IN THIS SECTION

- [General Routing | 238](#)
- [Interfaces and Chassis | 240](#)
- [Routing Protocols | 240](#)

Learn about open issues in this release 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 a PTX Series PIC with the CFP2-DCO-T-WDM transceiver installed, after repeated configuration rollbacks, the link sometimes takes a long time to come up. [PR1301462](#)
- When CFP2-DCO-T-WDM-1 is plugged in a PTX Series PIC, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)
- In a specific case, when semigraceful RCB reboot is initiated by the internal shell command **vhclient init 0**, GRES takes longer to complete. That is, 3 minutes as opposed to 21 seconds. We don't see this delay when we use the regular CLI command **request vmhost reboot** (graceful reboot) or disconnect the Routing Engine by plugging out the cable and reconnect it by plugging the cable back in (ungraceful reboot). [PR1312065](#)
- 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](#)
- The user might not be able to stop the zero-touch provisioning (ZTP) bootstrap when a PTX10016 or PTX10008 router with more number of line cards is powered on with factory-default configuration. [PR1369959](#)
- Control plane switch management (CPSM) daemon memory leak occurs in the VM host. The issue might also cause logrotate not to work, and result in a large CPSM log size. [PR1387903](#)

- The **show ephemeral-configuration** command syntax was changed between Junos OS Release 17.2X75 and Release 18.2X75. In Junos OS Release 18.2X75, the correct command is **show ephemeral-configuration instance**. [PR1391488](#)
- Because of the small counter size present in the ASIC, the normal discard counter reported in the CLI is less than the actual packet drop rate. [PR1394979](#)
- The rx_power value streamed to the telemetry server is the raw value (mW) returned directly from the transceiver driver. The Junos CLI value has been transformed in the transportd daemon into different units: Rx input total power (0.01dBm). [PR1411023](#)
- Sometimes the SFP+ read does not work on one or more ports of the LC1103-2C/6Q/60X. If this happens, the corresponding SFP+ module does not get detected and is not displayed at the Routing Engine CLI in the output of **show chassis hardware**. As a workaround, re-seat the SFP+ module. [PR1412897](#)
- The telemetry statistics might not account correctly for the traffic on SRTE policies (both byte count and packet count) on PTX Series platforms. This is a sensor-related issue. [PR1413680](#)
- VTY command **show filter index number counter** shows values as zero at 28-02-HOSTBOUND_NDP_DISCARD_TERM on a 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. There are no issues with NDP functionality, and DDoS for NDP is also working. [PR1420057](#)
- The firewall counter for the lo0 interface does not increment for FPC3-PTX-U2 and FPC3-PTX-U3. [PR1420560](#)
- An em2 interface configuration causes an FPC to crash during initialization and the FPC does not come online. After you delete the em2 configuration and restart the router, the FPC comes online. [PR1429212](#)
- More packet loss is seen after unified ISSU with inter-AS Layer 3 VPN Option B configuration. [PR1435578](#)
- Memory leaks are observed while running the longevity check. [PR1438358](#)

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after the upgrade. This is because of the presence of the old version of `/var/db/cfm.db`. [PR1281073](#)

Routing Protocols

- With bidirectional forwarding detection (BFD) configured on an aggregated Ethernet interface, if you disable or enable the aggregated Ethernet interface, then that interface and the BFD session might not come up. [PR1354409](#)

SEE ALSO

[What's New | 219](#)

[What's Changed | 231](#)

[Known Limitations | 236](#)

[Resolved Issues | 240](#)

[Documentation Updates | 247](#)

[Migration, Upgrade, and Downgrade Instructions | 248](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3 | 241](#)
- [Resolved Issues: 19.1R2 | 242](#)
- [Resolved Issues: 19.1R1 | 245](#)

Learn which issues were resolved in the Junos OS main and maintenance releases 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.

Resolved Issues: 19.1R3

General Routing

- Line-card interface stays down after maintenance. [PR1412126](#)
- On PTX3000 routers, the **core-olympus-fpc3-sevfpc.elf.0.tgz** core file is generated after switchover with GRES. [PR1415145](#)
- LACP packet does not pass through Layer 2 circuit. [PR1424553](#)
- The host-bound traffic might be dropped after you perform a change operation related to the prefix-list configuration. [PR1426539](#)
- On PTX Series platforms, reclassification policy applied on the route prefixes might not work. [PR1430028](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- Changing the hostname triggers LSP on-change notification, and currently IS-IS is sending the hostname instead of the system ID in OC paths. [PR1449837](#)
- JNP10K-LC2101 FPC generates Voltage Tolerance Exceeded major alarm for EA_chip_2v5 sensors. [PR1451011](#)
- On PTX5000 router, the 100-Gigabit Ethernet interface might not come up after flapping. [PR1453217](#)
- Traffic gets dropped silently upon interface flap after DRD autorecovery. [PR1459698](#)
- The forwarding option is missed in routing-instance type. [PR1460181](#)
- Hardware failure in CB2-PTX causes traffic interruption. [PR1460992](#)
- On PTX10008 routers, the **sample/syslog/log** action in output firewall filter with packet size less than 128 bytes might cause an ASIC wedge (all packet loss). [PR1462634](#)
- On PTX3000 and PTX5000 routers, PICs might restart if the QSFP optics is overheated. [PR1462987](#)
- On PTX10000 platforms, the FPC might restart during run time. [PR1464119](#)
- On PTX Series platforms, EBUF parity interrupt is not seen. [PR1466532](#)
- IPv6 traffic between two CE devices might get dropped in a Layer 3 VPN. [PR1466659](#)
- On PTX5000 platforms, Packet Forwarding Engine error logs **prds_packet_classify_notification: Failed to find fwd nh for flabel 48** might be reported when IGMP packets are sampled. [PR1466995](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- Incorrect counter values are observed for the arrival rate and peak rate for DDoS commands. [PR1470385](#)
- Traffic loops for pure Layer 2 packets coming over an EVPN tunnel with destination MAC matching IRB MAC. [PR1470990](#)
- On PTX10000 and PTX1000 platforms, MACsec traffic over a Layer 2 circuit might not work after upgrading from Junos OS Release 15.1 to later versions. [PR1475089](#)

Forwarding and Sampling

- On PTX Series or TVP platforms, the pfd might crash and might not be able to come up. [PR1452363](#)

Infrastructure

- On all Junos OS VM-based platforms, FPCs might reboot if jlock hog occurs. [PR1439906](#)

Interfaces and Chassis

- After a member interface flaps, the aggregated Ethernet interface remains down on 5-port 100-Gigabit Ethernet DWDM CFP2-ACO PIC. [PR1429279](#)

Multiprotocol Label Switching (MPLS)

- Kernel might crash and device might restart. [PR1478806](#)

Routing Protocols

- Route churn might be seen after the **maximum-prefixes** configuration is changed from value A to value B. [PR1423647](#)
- The rpd might crash with SRTE configuration change. [PR1442952](#)
- SSH login might fail if a user account exists in both the local database and RADIUS or TACACS+. [PR1454177](#)
- The rpd process might crash with BGP multipath enabled and the active route withdrawn occasionally. [PR1481589](#)

Resolved Issues: 19.1R2

General Routing

- The agentd sensor transmits multiple interface telemetry statistics per FPC slot. [PR1392880](#)
- Confirmation message is missing when issuing **request vmhost reboot re***. [PR1397912](#)
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. [PR1399683](#)
- Incorrect memory statistics message is seen in FPC logs of PTX Type 1 FPC. [PR1404088](#)
- ZTP upgrade might fail if there is more than one 10-Gigabit Ethernet interface connected to the DHCP server. [PR1404832](#)
- On PTX3000 and PTX5000, the backup Control Board's chassis environment status keeps "Testing" after the backup Control Board is removed or reinserted. [PR1405181](#)
- Openconfig-network-instance:network-instances support for IS-IS must be hidden unless supported. [PR1408151](#)
- The port at FPC(For example, JNP10K-LC1101) might fail to come up. [PR1409585](#)
- The CPU might be hogged by jsd process in JET scenario. [PR1409639](#)

- Hostname does not update at FPC shell after system configuration change on CLI. [PR1412318](#)
- Junos PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- The PTX1000/PTX10002 might stop forwarding packets after the chassis-control process restarts. [PR1414434](#)
- Traffic loss could be seen for duration of hold-time down timer when flapping an interface with hold-time down timer configured. [PR1418425](#)
- RX alarms are not set as according to the threshold value configured for the DCO Tunable Optics. [PR1419204](#)
- An interface might go to down state on QFX10000/PTX10000 platform. [PR1421075](#)
- Virtual Chassis might become unstable and FXPC generates core files when there are a lot of configured filter entries. [PR1422132](#)
- Packet Forwarding Engine wedge might be observed after performing the command **show forwarding-options load-balance ...** [PR1422464](#)
- The 4x10G interfaces on PTX3000/PTX5000 FPC type 3 might not come up after frequent flap ping occurs for a long time. [PR1422535](#)
- While committing a huge configuration, customer is seeing the error **error: mustd trace init failed**. [PR1423229](#)
- Traffic is dropped after FPC reboot with AE member links deactivated by remote device. [PR1423707](#)
- JDI-I2circuit-REGRESSIONS:AE not coming up with LACP enabled over the ccc circuit between R0 and R3. [PR1424553](#)
- The **per-interface-per-member-link** command is hidden for PTX5000 FPC. [PR1425372](#)
- Specific interface on P3-15-U-QSFP28 PIC card remains down until another interface comes up. [PR1427733](#)
- An interface with port 7, 9, 17, 19, 27, or 29 might go to down state on 30-port 40-Gigabit Ethernet or 100-Gigabit Ethernet line cards. [PR1427883](#)
- When an interface is configured with jumbo frames support (For example, MTU = 9216), the effective MTU size for locally sourced traffic is 24 bytes less than the expected value. [PR1428094](#)
- Inline Jflow might cause PECHIP Major error. [PR1429419](#)
- IPFIX J-Flow timestamp is not matching with NTP synchronized system time. [PR1431498](#)
- **SIB Link Error** detected on a specific Packet Forwarding Engine might cause complete service impact. [PR1431592](#)
- Scaled filter leads to packet drop as `flt.Dispatcher.flt_err` on PTX Series router. [PR1433648](#)
- IPv6 neighbor solicitation packets are getting dropped on PTX Series router. [PR1434567](#)
- On PTX10016 platforms, mastership cannot switch over immediately while facing SSD failure. [PR1437745](#)

- On PTX10002, no chassis alarm is raised when PEM is removed or power is lost to PEM. [PR1439198](#)
- Interfaces on PTX Series router might not come up after FPC restart or port flap. [PR1442159](#)
- BCM FW needs to be upgraded to DE2E. [PR1445473](#)
- Receipt of a malformed packet for J-Flow sampling might create a FPC process core file. [PR1445585](#)
- The jdhcpd process might crash after issuing the command **show access-security router-advertisement-guard**. [PR1446034](#)
- Egress sampling for sFlow might stop working for more than eight interfaces on PTX Series platforms. [PR1448778](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)
- Firewall filter applied at interface level is not working when entropy level is present in certain scenarios. [PR1452716](#)
- The FPC might crash when the severity of error is modified. [PR1453871](#)

Infrastructure

- Command **request system recover oam-volume** might fail on PTX Series router. [PR1425003](#)
- Unsupported package warning is seen after system upgrade. [PR1427344](#)

Interfaces and Chassis

- Syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** occurs upon LFM-related configuration commit on AE interfaces. [PR1423586](#)
- After interface flap, the LAG/AE remains down on 5X100GE DWDM CFP2-ACO MIC. [PR1429279](#)
- Some ports on PTX might remain down after rebooting the FPC/device at remote side. [PR1429315](#)

Layer 2 Ethernet Services

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

MPLS

- Services dependent on LDP might be impacted if committing any configuration changes. [PR1416032](#)
- RSVP Path error received on PSB:2 (new path calculated by CSPF) is not treated as optimization when CSPF is computed and optimization retry is not honoring $2^{\text{retry}} + \text{rsvp-error-hold-time}$. [PR1416948](#)
- The dynamic bypass RSVP LSP tears down when being used to protect LDP LSP. [PR1425824](#)
- The transit packets might be dropped if an LSP is added or changed. [PR1447170](#)

Platform and Infrastructure

- REST API process will get non responsive when a number of requests come in at a high rate. [PR1449987](#)

Routing Protocols

- Routing Engine based micro BFD packets do not go out with configured source IP when the interface is in logical system. [PR1370463](#)
- Syslog message is seen whenever prefix SID coincides with the node SID. [PR1403729](#)
- Dynamic routing protocol flapping with vmhost Routing Engine switchover occurs on next-generation Routing Engine. [PR1415077](#)
- PTX Series device cannot intercept PIM BSR message. [PR1419124](#)
- RPD might crash with OSPF overload configuration. [PR1429765](#)

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 L2 circuit flaps repeatedly. [PR1282875](#)
- Memory leak might happen if PIM messages are received over an MDT (mt- interface) in Draft-Rosen MVPN scenario. [PR1442054](#)

Resolved Issues: 19.1R1

General Routing

- Repeated log messages %PFE-3 fpcX expr_nh_index_tree_ifl_get and expr_nh_index_tree_ipaddr_get are observed when sampling packet is discarded with log (or syslog) statement under firewall filter. [PR1304022](#)
- On PTX Series platform, multicast traffic packet drop seen is more than 50 percent when having FPC1/FPC2 mix with FPC3. [PR1339481](#)
- On PTX10001 platform, the FRR link-protection convergence during FRR and MBB with various MPLS optimize-timers is observed. [PR1355953](#)
- The netproxy service client component fails to start after issuing **request vmhost reboot** command. [PR1365664](#)
- The IPLC card might take a long time to come up after requesting it online from an offline state. [PR1368637](#)
- Some harmless log messages are suppressed on the backup SPMB. [PR1369731](#)
- On PTX10001 platform, 100G-LR4 optics and 100G-ER4 optics are not supported. [PR1371590](#)
- Inline BFD might keep flapping when inline sampling is configured. [PR1376509](#)
- Traffic might be dropped on third-generation FPCs on PTX Series routers. [PR1378392](#)

- BFD sessions flap when restarting one FPC on PTX10000. [PR1383703](#)
- Packet Forwarding Engine based local repair does not happen for IP routes pointing to unilist of composites with indirect next hops. [PR1383965](#)
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent the major alarm. [PR1384435](#)
- Packet drop might be seen in AF3 queue on PTX Series platforms. [PR1385454](#)
- The system might hang after issuing **request system reboot**. [PR1386279](#)
- BFD flaps are seen on PTX Series platforms with inline BFD. [PR1389569](#)
- Forwarding issue on mixed link-speed aggregated Ethernet interface after FPC reloads. [PR1390417](#)
- PTX10002-60C FPC might not be detected after the ukern crashes. [PR1396507](#)
- High jsd or na-grpcd CPU usage might be seen even when JET or JTI is not used. [PR1398398](#)
- CPU hog might be observed on PTX Series platform. [PR1399369](#)
- Log message **JAM HW data base open failed for ptx5kpic_3x400ge-cfp8** occurs during commit. [PR1403071](#)
- On PTX3000, FPCs are not able to come online for tens of minutes after a reboot of the chassis. [PR1404611](#)
- The 100G SR4 optics with part number 740-061405 should be displayed as **QSFP-100G-SR4-T2**. [PR1405399](#)
- Layer 2 VPN flaps repeatedly after linkup between PE and CE devices under asynchronous notification and some types of MIC conditions. [PR1407345](#)
- For PTX10001-20C devices, the DHCP relay functionality and binding of DHCP do not work. [PR1407476](#)
- On PTX3000 router, the rpd crash is observed at `if_addr_link`, `krt_chnh_template_create_restart`, `krt_chnh_create_restart`, `krt_comp_add_comp_nh`, `krt_build_comp_nh`, `krt_build_nexthop`, `krt_rt_add_sock`, `krt_decode_rt`, `krt_sysctl_read_consume`, `krt_rt_read`, `krt_sys_rtread`, `krt_var_init`, `ctx_handle_node`, `ctx_walk_features`, `task_read_config`, `main`. [PR1409051](#)

Interfaces and Chassis

- PE Chip:pe0[0]: IPW: **oversize_drop error** causes major error on FPC. [PR1375030](#)

MPLS

- MPLS LSP will remain in down state because of routing loop detection after flapping link between PE router and egress PE. [PR1384929](#)
- The rpd might crash when LDP route with indirect next hop is deleted. [PR1398876](#)
- A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based-stats are used. [PR1401152](#)

Platform and Infrastructure

- On PTX Series routers, the **RPM hardware-timestamp** and **one-way-hardware-timestamp** statements are not enabled. [PR1399842](#)
- When non-root user tries to archive the **var/log**, some files are missing from the cscript.log file. [PR1405903](#)

Routing Protocols

- The rpd process generates a core file on the backup Routing Engine during neighborship flap when using authentication key with size larger than 20 characters. [PR1394082](#)
- The rpd RT_NEXTHOPS_TEMPLATE memory leaks while using segment routing for IS-IS protocols. [PR1404134](#)

SEE ALSO

What's New 219
What's Changed 231
Known Limitations 236
Open Issues 238
Documentation Updates 247
Migration, Upgrade, and Downgrade Instructions 248

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for the PTX Series.

SEE ALSO

What's New 219
What's Changed 231
Known Limitations 236
Open Issues 238
Resolved Issues 240
Migration, Upgrade, and Downgrade Instructions 248

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.1 | 248](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 251](#)
- [Upgrading a Router with Redundant Routing Engines | 251](#)

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

Basic Procedure for Upgrading to Release 19.1

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 19.1R3:

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. 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 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 jinstal 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 composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot  
source/junos-install-ptx-x86-64-19.1R3.9.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/junos-install-ptx-x86-64-19.1R3.9-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](#).

NOTE: After you install a Junos OS Release 19.1jinstall 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 | 219](#)

[What's Changed | 231](#)

[Known Limitations | 236](#)

Open Issues 238
Resolved Issues 240
Documentation Updates 247

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- What's New | 252
- What's Changed | 267
- Known Limitations | 272
- Open Issues | 276
- Resolved Issues | 282
- Documentation Updates | 303
- Migration, Upgrade, and Downgrade Instructions | 303

These release notes accompany Junos OS Release 19.1R3 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 19.1R3-S2 | 253
- What's New in Release 19.1R3 | 253
- What's New in Release 19.1R2 | 253
- What's New in Release 19.1R1 | 255

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

NOTE: The following QFX Series platforms are supported in Release 19.1R3: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.

What's New in Release 19.1R3-S2

Routing Policy and Firewall Filters

- **Loopback firewall filter scale optimization (QFX5120-48Y and EX4650)**—Starting with Junos OS Release 19.1R3-S2, you can configure up to 768 loopback filter terms for IPv6, and up to 1152 terms for IPv4. To do so, you configure an ingress firewall filter, apply it to the loopback interface, and then enable the **loopback-firewall-optimization** command at the **[edit chassis]** hierarchy level (this triggers the packet forwarding engine (PFE) to restart).

Terms that include a reserved multicast destination (for example 224.0.0.x/24) and terms with a time-to-live (TTL) of 0/1 are not directly supported. Instead, you need to configure a separate filter for these terms. For example, to count OSPF packets on the loopback interface, you would create a separate filter with terms for the protocol (OSPF) to count packets destined to a reserved multicast address (such as 224.0. 0.6).

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

What's New in Release 19.1R3

There are no new features and enhancements to the existing features for the QFX Series switches in Junos OS Release 19.1R3.

What's New in Release 19.1R2

EVPN

- **EVPN-VXLAN support (QFX10002-60C switches)**—Starting in Junos OS Release 19.1R2, the QFX10002-60C switch can function as a Layer 2 or Layer 3 VXLAN gateway in both EVPN-VXLAN centrally-routed and edge-routed bridging overlays (EVPN-VXLAN topologies with two-layer and collapsed IP fabrics). In these roles, the switch supports the following features:
 - Enterprise style of Layer 2 interface configuration
 - Active/active multihoming
 - Default routing instance

- Multiple routing instances of type virtual switch, and VLAN-aware service on the virtual switch routing instance
- Pure type-5 routes
- Proxy ARP use and ARP suppression, and proxy NDP use and NDP suppression on an IRB interface
- ESIs on physical and aggregated Ethernet interfaces
- OSPF, IS-IS, BGP, and static routing on IRB interfaces
- DHCP relay
- IPv6 support for user data traffic
- EVPN-VXLAN with MPLS as transport layer
- MAC mobility

[See [EVPN User Guide](#).]

- **BPDU protection in EVPN-VXLAN (QFX5100, QFX5110, and QFX5200 switches)**—Starting in Junos OS Release 19.1R2, you can enable BPDU protection in an EVPN-VXLAN configuration. With a spanning tree protocol configured on an edge port, you can enable BPDU protection. If a BPDU is received on the edge port, the edge port is disabled and it stops forwarding all traffic. You can also configure BPDU protection on VXLAN interfaces without a spanning tree protocol configured, or enable BPDU protection and have other traffic forwarded. Only the BPDUs are dropped, and all other traffic is forwarded. Additionally, you can unblock an interface either automatically or manually.
 - To enable BPDU protection with RSTP on an edge port on access and leaf devices:
`set protocols rstp interface interface-name edge`
`set protocols rstp bpdv-block-on-edge`
 - To enable BPDU protection with a spanning tree protocol on access and leaf devices:
`set protocols layer2-control bpdv-block interface interface-name`
 - To enable BPDU protection but still forward other traffic on access and leaf devices:
`set protocols layer2-control bpdv-block interface interface-name drop`
 - To automatically unblock an interface using an expiry timer on access and leaf devices:
`set protocols layer2-control bpdv-block disable-timeout time in seconds`
 - To manually unblock an interface on access and leaf devices:
`run clear error bpdv interface all`
- **Support for EVPN-VXLAN features (QFX5120-32C)**—Starting in Junos OS Release 19.1R2, QFX5120-32C switches support the following features in an EVPN-VXLAN environment:
 - Firewall filtering and policing
 - Graceful restart

- Class of service (CoS)
- Virtual machine traffic optimization (VMTO) for ingress traffic
- MAC limiting (firewall filter-based)
- Storm control
- Port mirroring and analyzers
- Core isolation

[See the [EVPN User Guide](#).]

What's New in Release 19.1R1

Hardware

- **QFX5120-32C switches**— Starting with Release 19.1R1, Junos OS supports the fixed-configuration QFX5120-32C switch. This switch provides 100-Gbps spine-and-leaf connectivity in Layer 2 and Layer 3 fabrics for cloud and Web services.

The QFX5120-32C has 2 SFP+ ports that operate at 10-Gbps speed, and 32 ports that can operate at 40-Gbps (with QSFP+ transceivers) and 100-Gbps speeds (with QSFP28 transceivers). You can use breakout cables to channelize the 40-Gbps ports into four 10-Gigabit Ethernet interfaces and the 100-Gbps ports into four 25-Gigabit Ethernet interfaces.

The QFX5120-32C is available with AC power supplies and with front-to-back or back-to-front airflow.

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for SFTP global disablement (QFX Series)**—Starting in Junos OS Release 19.1R1, we have globally disabled incoming SSH File Transfer Protocol (SFTP) connections by default. You can enable incoming SFTP connections globally by configuring the statement **sftp-server** at the **[edit system services ssh]** hierarchy level. Prior to Junos OS Release 19.1R1, the incoming SFTP connections were globally enabled by default.

[See [Configuring sftp-server](#)]

Class of Service (CoS)

- **Support for per-port buffer monitoring (QFX5000 switches)**—Starting with Junos OS Release 19.1R1, to keep track of peak buffer occupancy for each queue or priority group on a port, you can enable per-port buffer monitoring on a QFX5000 Series switch by setting **buffer-monitor-enable** at the **[edit chassis fpc slot-number traffic-manager]** hierarchy level. You can then monitor the buffer occupancy on the designated ports by executing the **show interfaces priority-group interface-name buffer-occupancy** or **show interfaces queue interface-name buffer-occupancy** command.

[See [traffic-manager](#).]

- **Support for class of service (CoS) on QFX5120-32C switches (QFX Series)**—Starting in Junos OS Release 19.1R1, QFX5120-32C switches support most class of service (CoS) features. IP precedence classification is not supported; DSCP classifiers are supported but can't be set at ingress. Also, as with other QFX5200 series switches, CoS flexible hierarchical scheduling (ETS) is not supported.

CoS is the assignment of traffic flows to different service levels. Service providers can use router-based CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows.

[See [CoS Operational Comparison Between QFX5100, QFX5120, QFX5200, and QFX5210 Switches](#).]

EVPNs

- **EVPN proxy ARP and ARP suppression, and proxy NDP and NDP suppression without IRB interfaces (QFX10000 switches)**—Starting in Junos OS Release 19.1R1, QFX10000 switches that function as Layer 2 VXLAN gateways in an EVPN-VXLAN environment support proxy ARP and ARP suppression, and proxy NDP and NDP suppression on non-IRB interfaces. Now, any interface configured on these Layer 2 VXLAN gateways can deliver ARP and NDP requests from both local and remote devices.

In addition, you can now control the following aspects of the MAC-IP address bindings database on a QFX10000 switch:

- The maximum number of MAC-IP address entries in the database
- The amount of time a locally learned MAC-IP address binding remains in the database

[See [EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression](#).]

Forwarding and Sampling

- **Customizing hashing parameters and shared-buffer alpha values for better load balancing (QFX5100, QFX5110, QFX5200, and QFX5210 switches)**—These switches achieve load balancing through use of a hashing algorithm, which determines how to forward traffic over LAG bundles or to next-hop devices when ECMP is enabled. The hashing algorithm makes hashing decisions based on values in various packet fields. Starting with Junos OS Release 19.1R1, you can explicitly configure some hashing parameters to make hashing more efficient. The shared-buffer pool is a global memory space that all ports on the switch share dynamically as they need buffers. The switch uses the shared-buffer pool to absorb traffic bursts after the dedicated-buffer pool is exhausted. The shared-buffer pool threshold is dynamically calculated based on a factor called “alpha”. Also starting with Junos OS Release 19.1R1, you can specify the alpha, or dynamic threshold, value to determine the change threshold of shared buffer pools for both ingress and egress buffer partitions.

To specify hashing parameters:

```
user@switch# set forwarding-options enhanced-hash-key hash-parameters (ecmp | lag)
```

To specify a threshold value for a particular queue:


```
user@switch# set class-of-service shared-buffer (ingress|egress) buffer-partition buffer
dynamic-threshold value
```

[See [hash-parameters](#) and [buffer-partition](#)].

General Routing

- **Supported features on new hardware (QFX5120-32C)**—Starting with Junos OS Release 19.1R1, the following Junos OS features are supported on QFX5120-32C switches:

- **Layer 2 unicast features:**
 - 802.1Q VLAN trunking
 - 802.1p
 - PVLAN
 - Routed VLAN interface (RVI)
 - Layer 3 VLAN-tagged logical interfaces
 - 4096 VLANs
 - MAC address filtering
 - MAC address aging configuration
 - Static MAC address assignment for interface
 - Per-VLAN MAC learning (limit)
 - MAC learning disable
 - Persistent MAC (sticky MAC)
 - Q-in-Q Tag manipulation
 - MAC address limit per port
 - MAC limiting
 - MAC limiting per port, per VLAN
 - MAC move limiting
 - PVLAN on Q-in-Q
 - 802.1D
 - 802.1w (RSTP)
 - 802.1s (MST)
 - BPDU protection
 - Loop protection
 - Root protection

- VSTP
- RSTP and VSTP running concurrently
- Link aggregation (static and dynamic) with LACP (fast and slow LACP)
- LLDP
- Multiple VLAN Registration Protocol (802.1ak)

[See [Ethernet Switching User Guide](#).]

- **Layer 2 multicast features:**

- IGMP snooping for IGMPv1, IGMPv2, and IGMPv3
- IGMP proxy
- IGMP querier
- Virtual router (VRF-lite) IGMP snooping

[See [Multicast Overview](#).]

- **Layer 3 unicast features:**

- Static routing, ping, and traceroute (IPv4, IPv6)
- OSPFv2 (IPv4) and OSPFv3 (IPv6)
- RIPv2
- BGP (IPv4, IPv6), BGP 4-byte ASN support, and BGP multipath
- MBGP (IPv4)
- IS-IS (IPv4, IPv6)
- BFD (for RIP, OSPF, IS-IS, BGP, PIM)
- Filter-based forwarding (FBF)
- Unicast reverse path forwarding (RPF)
- IP directed broadcast traffic forwarding
- VRRP
- VRRPv3 (IPv6)
- Neighbor Discovery Protocol (IPv6)
- Path MTU discovery
- IPv6 CoS—Behavior aggregate (BA) classifiers, multifield (MF) classifiers and rewrite rules, traffic-class scheduling
- IPv6 stateless address autoconfiguration
- ECMP—32-way

- Hierarchical ECMP
- Virtual router (VRF-lite) IS-IS, RIP, OSPF, BGP

[See [BGP User Guide](#), [IPv6 Neighbor Discovery User Guide](#), [IS-IS User Guide](#), [OSPF User Guide](#), [Protocol-Independent Routing Properties User Guide](#), and [RIP User Guide](#).]

- **Layer 3 multicast features:**

- IGMP version 1 (IGMPv1), version 2 (IGMPv2), and version 3 (IGMPv3)
- IGMP filtering
- PIM sparse mode (PIM-SM)
- PIM source-specific multicast (PIM-SSM)
- PIM dense mode (PIM-DM)
- Virtual router (VRF-lite) PIM, IGMP
- Multicast Source Discovery Protocol (MSDP)

[See [Multicast Overview](#).]

- **VXLAN features:**

- EVPN-VXLAN—Layer 2 and Layer 3 VXLAN gateways
 - Pure type-5 routes. [See [EVPN Type-5 Route with VXLAN encapsulation for EVPN-VXLAN](#).]
 - IGMP snooping. [See [Overview of IGMP Snooping in an EVPN-VXLAN Environment](#).]
 - Tunneling of Q-in-Q traffic. [See [Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network](#).]
 - Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces. [See [Supported Protocols on an IRB Interface in EVPN-VXLAN](#).]
 - Support for IPv6 data traffic. [See [Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay](#).]
 - MAC mobility. [See [Overview of MAC Mobility](#).]
 - EVPN proxy ARP and ARP suppression, and NDP and NDP suppression. [See [EVPN Proxy ARP and ARP Suppression, and NDP and NDP Suppression](#).]
- OVSDB-VXLAN—Layer 2 VXLAN gateway. [See [Understanding the OVSDB Protocol Running on Juniper Networks Devices](#).]
- PIM-based Layer 2 VXLAN gateway. [See [Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches](#).]
- MPLS support. [See [MPLS Feature Support on QFX Series and EX4600 Switches](#).]

- Multichassis link aggregation group (MC-LAG). [See [Multichassis Link Aggregation Features, Terms, and Best Practices](#).]
- **Services support:**
 - sFlow. [See [Overview of sFlow Technology](#).]
 - Port mirroring. [See [Understanding Port Mirroring](#).]
 - Storm control. [See [Understanding Storm Control](#).]
- Resilient hashing support for LAGs and ECMP routes. [See [Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups](#).]
- Distributed denial of service (DDoS) protection. [See [Understanding Distributed Denial-of-Service Protection on QFX Series Switches](#).]
- Unified Forwarding Table (UFT). [See [Understanding the Unified Forwarding Table](#).]

Interfaces and Chassis

- **Multichassis link aggregation groups, configuration synchronization, and configuration consistency check (MC-LAG) (QFX5120 switches)**—Starting in Junos OS Release 19.1R1, MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running spanning tree protocols (STP).
[See [Multichassis Link Aggregation Features, Terms, and Best Practices](#).]
- **Increasing the number of ARP and neighbor discovery entries to 256,000 (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 19.1R1, the number of ARP and neighbor discovery entries has increased to 256,000 when enabling the **enhanced-convergence** statement. Enhanced convergence improves Layer 2 and Layer 3 convergence time during enhanced MC-LAG and VXLAN L3 gateway restoration scenarios.

To increase the number of ARP and neighbor discovery entries, enable the **arp-enhanced-scale** statement at the **[edit system]** hierarchy.

[See [Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies](#).]

- **Channelizing enhancement on QFX5210-64C switches**—Starting in Junos OS Release 19.1R1, the behavior of Flexi-pic mode on QFX5210-64C switches has improved. Channelizing ports in this mode no longer disables a corresponding port. The new behavior allows you to use any port within four designated blocks for channelization as long as the total number of channels does not exceed 128 or 32 in any one of the four blocks. Channelization helps to maximize port utilization.
[See [Channelizing Interfaces on Switches](#).]
- **Channelizing interfaces on QFX5120-32C switches**—The 32 ports on the QFX5120-32C switch support native 40- or 100-Gigabit Ethernet configuration and channelized 10-, 25-, or 40-Gigabit Ethernet configuration. Starting in Junos OS Release 19.1R1, you can channelize the default 100-Gbps ports into

four 25-Gigabit Ethernet or two 50-Gigabit Ethernet interfaces, and the 40-Gbps ports into four 10-Gigabit Ethernet interfaces (using breakout cables).

If you have disabled auto-channelization, then to channelize the ports, manually configure the port speed using the **set chassis fpc slot-number port port-number channel-speed speed** command, where the speed can be set to **10G, 25G, 50G**.

NOTE:

- The last 100-Gbps port (port 31) does not support four 10-Gigabit Ethernet port or four 25-Gigabit Ethernet port channelization. Only 40-Gigabit Ethernet, 100-Gigabit Ethernet and 2x50-Gigabit Ethernet interfaces are supported on port 31.
- You cannot configure channelized interfaces to operate as Virtual Chassis ports.

[See [Channelizing Interfaces on Switches](#).]

Junos Telemetry Interface

- **Support for the Junos telemetry interface (JTI) (QFX10002 and PTX10002)**—Starting with Junos OS Release 19.1R1, you can provision sensors through the Junos telemetry interface to export telemetry data for several network elements without involving polling. You can stream data through UDP or gRPC.

Only the following sensors are supported on QFX10002 switches and PTX10002 routers:

- Physical interfaces statistics
- Label-switched-path (LSP) statistics
- Network processing unit (NPU) memory
- NPU memory utilization
- CPU memory

To provision sensors to stream data through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level.

To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters.

Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module and YANG models.

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

Layer 2 Features

- **L2PT support (QFX5200 switches and QFX5200 Virtual Chassis)**—Starting with Junos OS Release 19.1R1, you can configure Layer 2 protocol tunneling (L2PT) for the following protocols on QFX5200 switches and QFX5200 Virtual Chassis: CDP, E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, STP (including RSTP and MSTP), UDLD, VSTP, and VTP.

[See [Layer 2 Protocol Tunneling](#).]

Licensing

- **QFX5120-32C switch license** —Starting in Junos OS Release 19.1R1, Juniper Networks introduces the QFX5120-32C switch.

The QFX5120-32C switch supports the following licenses models:

- Base features for the QFX5120-32C switch include OSPF, OSPFv3, and RIPng.
- Advanced Feature License (AFL) for QFX5120-32C switch includes BGP, IS-IS, MPLS, VXLAN, and Open vSwitch Database (OVSDb).
- PFL for QFX5120-32C switch includes Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS), Virtual Extensible Local Area Network (VXLAN), and Open vSwitch Database (OVSDb).

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

Management

- **Tracing support for individual JET application files (QFX Series)**—Previously you could configure traceoptions for all applications. Starting in Junos OS Release 19.1R1, you can also configure traceoptions for an individual application file. If you configure trace options both globally (all applications) and locally (by application file), the local configuration has the higher priority. You must commit global traceoptions and the daemonized application configurations at the same time for the global traceoptions for the daemonized application to take effect.

[See [application](#).]

MPLS

- **MPLS scaling enhancements (QFX5100, QFX5110, QFX5200, QFX5210)**—Starting in Junos OS Release 19.1R1, MPLS scaling is enhanced on the switches. For instance, you can increase the scale from its default 1024 to 8192 on the QFX5100. 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](#).]

- **Control transport address used for targeted-LDP session (QFX Series)**—Currently, only the router ID or interface address is used as the LDP transport address. Starting in Junos OS Release 19.1R1, you can

configure any other IP address as the transport address of targeted LDP sessions, session groups, and interfaces. This new configuration is applicable only for configured LDP neighbors that have Layer 2 circuit, MPLS, and VPLS adjacencies.

This feature is beneficial when you have multiple loopback interface addresses, and different IGPs associated with LDP interfaces, and you can control the session established between targeted LDP neighbors with the configured transport address.

[See [Control Transport Address Used for Targeted-LDP Session](#).]

- **Policy-based multipath routes (QFX Series)**—In segment routing networks with multiple protocols in the core, you can combine segment routing traffic engineered (SR-TE) LDP routes and SR-TE IP routes to create a multipath route that is installed in the routing information base (also known as routing table). You can resolve BGP service routes over the multipath route through policy configuration and steer traffic differently for different prefixes.

[See [Policy-Based Multipath Routes Overview](#).]

- **Use of SID labels as first hop for resolving non-colored static segment routing LSPs (QFX Series)**—Currently, for a static non-colored segment routing traffic-engineered LSP to be usable, the first hop of the segment list must be an IP address. Only the second to *n*th hop could be segment identifier (SID) labels. Starting in Junos OS Release 19.1R1, this requirement does not apply. You can now configure SID labels as the first hop in the segment list.

With this configuration, static non-colored segment routing LSPs are resolved using MPLS fast reroute (FRR) and weighted equal-cost multipath. Without this configuration, by default, the LSPs are resolved using IP address.

[See [Static Segment Routing Label Switched Path](#).]

- **Support of install statement for segment routing LSPs (QFX Series)**—The *install destination-prefix* statement which is currently supported at the `[edit protocols mpls label-switched-path lsp-name]` and `[edit protocols mpls static-label-switched-path lsp-name ingress]` hierarchy levels is now also supported at the `[edit protocols source-packet-routing source-routing-path lsp-name]` hierarchy level for both colored and non-colored static segment routing label-switched paths (LSPs).

You can associate one or more prefixes with a segment routing LSP using the **install** statement. When the LSP is up, all the prefixes are installed as entries into the **inet.3** or **inet6.3** routing table.

[See [install \(Protocols MPLS\)](#).]

Network Management and Monitoring

- **Local port mirroring support (QFX10002-60C switch)**—Starting in Junos OS Release 19.1R1, QFX10002-60C switches support local port mirroring. Port mirroring copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

[See [Examples: Configuring Port Mirroring for Local Analysis](#).]

- **sFlow performance improvements (QFX Series)**—Starting in Junos OS Release 19.1R1, the following improvements have been added to the sFlow technology feature:
 - For MX Series, PTX Series, and QFX Series, you can configure forwarding class and DSCP values per collector.
 - For PTX Series and QFX Series, you can configure IPv6 addresses for the **source-ip** and **agent-id**.
 - Enhancements are made to the following CLI commands: **show sflow collector**, **show sflow collector address ip-address**, and **show sflow interface**.

[See [Understanding How to Use sFlow Technology for Network Monitoring](#), [collector](#), [agent-id](#), [source-ip](#), [show flow collector](#), and [show flow interface](#).]

Routing Policy and Firewall Filters

- **Support for IPv6 filter-based forwarding (QFX5100, QFX5110, and QFX5200 switches)**— Starting with Junos OS Release 19.1R1, you can use stateless firewall filters in conjunction with filters and routing instances to control how IPv6 traffic travels in a network. This is called IPv6 filter-based forwarding. To set up this feature, you define a filtering term that matches incoming packets based on the source or destination address and then specify the routing instance to send packets to. You can use filter-based forwarding to route specific types of traffic through a firewall or security device before the traffic continues on its path. You can also use it to give certain types of traffic preferential treatment or to improve load balancing of switch traffic.

This feature was previously supported in an "X" release of Junos OS.

[See [Firewall Filter Match Conditions](#) and [Understanding Filter-Based Forwarding](#).]

- **Support for 2000 Egress Firewall Filters (QFX5110 switches)**—Starting in Junos OS Release 19.1R1, you can configure up to 2000 VLAN firewall filters on the switch. This feature is only supported in the egress direction (traffic exiting the VLAN). To configure, include the **egress-to-ingress** option under the **from** statement at the **[edit firewall]** hierarchy level.

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

- **Support for packet load balancing based on GTP-TEID hashing (QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 19.1R1, you can configure load balancing of IPv4 or IPv6 packets by using GPRS Tunneling Protocol-tunnel endpoint identifier (GTP-TEID) field hash calculations. The GTP-TEID hashing is added to the Layer 2 and Layer 3 field hashing that you have already configured. To enable this feature, configure the **gtp-tunnel-endpoint-identifier** statement at the **[edit forwarding-options enhanced-hash-key family inet]** or the **[edit forwarding-options enhanced-hash-key family inet6]** hierarchy Level. GTP versions 1 and 2 are supported; they support only user data. You must use UDP port number 2152 for both GTP versions.

[See [gtp-tunnel-endpoint-identifier](#).]

- **Support for matching IPv6 source addresses from an inet6 egress interface (QFX5100)**—Starting in Junos OS Release 19.1R1, you can configure a firewall filter on a IPv6 egress interface to match specified IPv6 source or destination addresses, for example, to protect a third-party device connected to the switch.

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

Routing Protocols

- **Support for 128 equal-cost paths for BGP multipath (QFX10000)**—Starting with Junos OS Release 19.1R1, you can configure a maximum of 128 equal-cost paths for external BGP peers. Previously, the maximum number supported was 64. For MPLS routes, the maximum number of equal-cost paths you can configure remains unchanged at 64. To specify 128 equal-cost paths for external BGP peers, include the **maximum-ecmp 128** statement at the **[edit chassis]** hierarchy level. You must also configure a routing policy that exports routes from the routing table into BGP. Define a routing policy by including the **policy-statement *policy-name*** set of statements at the **[edit policy-options]** hierarchy level. Apply the policy to routes exported to the forwarding table by including the **export *policy-name*** statement at the **[edit routing-options forwarding-table]** hierarchy level.

[See [maximum-ecmp](#).]

- **Support for BGP graceful shutdown (QFX Series)**— Starting in Junos OS Release 19.1R1, graceful traffic migration from one BGP next hop to another is supported, without traffic interruption. Also, BGP administrative shutdown communication can be sent to the BGP peer.

You can configure both **graceful-shutdown** and **shutdown** statements at the **[edit protocols bgp]**, **[edit protocols bgp group *group-name*]**, and **[edit protocols bgp group *group-name* neighbor *address*]** hierarchy levels.

NOTE: Graceful shutdown is disabled by default.

[See: [graceful-shutdown \(Protocols BGP\)](#), [shutdown \(Protocols BGP\)](#).]

- **Support for export of BGP Adjacency-RIB-Out through BGP Monitoring Protocol (BMP) (QFX Series)**— Starting in Junos OS Release 19.1R1, BMP is enhanced to support route monitoring of pre and post **rib-out** policy.

You can configure **post-policy** and **pre-policy** under **rib-out** statement at **[edit protocols bgp bmp]**, **[edit protocols bgp group *group-name* bmp]**, and **[edit protocols bgp group *group-name* neighbor *address* bmp]** hierarchies.

NOTE: The default monitoring mode of rib-out is **pre-policy**.

[See: [Understanding the BGP Monitoring Protocol.](#)]

- **Support for policy-based allocation for IPv4 BGP-labeled unicast (QFX Series)**—Starting in Junos OS Release 19.1R1, this feature supports:
 - Allocating policy-based label for IPv4 BGP-LU prefixes in per-prefix label allocation mode.
 - 1:1 mapping between prefixes and labels.
 - Map policy for labels.
 - Fallback actions of dynamic and reject for handling error conditions.

[See [policy-options](#), [route-filter-list](#).]

System Management

- **Support for aggregated Ethernet and loopback interfaces on primary and secondary interfaces using PTP (QFX5110 switches)** —Starting with Junos OS Release 19.1R1, you can configure both primary and secondary interfaces as aggregated Ethernet and loopback interfaces using PTP over IPv4 and IPv6 unicast transport on the IEEE 1588v2 default profile and the G.8275.2 enhanced profile. Although, the loopback interface (lo0.0) is the same for both the primary and secondary aggregated Ethernet interfaces, the IP addresses must be unique.

[See [Understanding the PTP G.8275.2 Enhanced Profile \(Telecom Profile\)Multicast Overview.](#)]

SEE ALSO

Changes in Behavior and Syntax 267
Known Behavior 272
Known Issues 276
Resolved Issues 282
Documentation Updates 303
Migration, Upgrade, and Downgrade Instructions 303

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.1R3 | 268](#)
- [What's Changed in Release 19.1R2 | 269](#)
- [What's Changed in Release 19.1R1 | 270](#)

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

What's Changed in Release 19.1R3

General Routing

- **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.

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 19.1R3, 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](#).]

Network Management and Monitoring

- **entPhysicalTable fetched on QFX10002**—In Junos OS Release 19.1R3, the MIB data for `entPhysicalTable` is fetched on a QFX10002-72Q or QFX10002-36Q switch.

[See [SNMP Explorer](#).]

Platform and Infrastructure

- **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.
- **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 Release 19.1R2

EVPN

- **Support for disabling automatic ESI generation (MX Series and QFX Series)**—Starting with Junos OS Release 19.1R2, Junos OS supports disabling the automatic ESI generation for virtual gateway addresses. We recommend that you disable the automatic ESI generation for EVPN networks with edge-routed bridging to improve performance. To disable automatic ESI generation, include the **no-auto-virtual-gateway-esi** statement at the **[edit interfaces name irb unit logical-unit-number]** hierarchy level.
- **Logical Interface is created along with physical Interface by default (QFX Series switches)**—In Junos OS Release 19.1R2 and later, logical interface is created on **ge**, **et**, **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, earlier when you view the **show interfaces** command, by default, only the physical interface (**ge-0/0/0**), is displayed. Now, the logical interface (**ge-0/0/0.16386**) is also displayed.

Interfaces and Chassis

- **The resilient-hash statement is no longer available under aggregated-ether-options (QFX5200 and QFX5210 switches)**—Starting in Junos OS Release 19.1R2, the **resilient-hash** statement is no longer available in the **[edit interfaces aex aggregated-ether-options]** hierarchy level. Resilient hashing is not supported on LAGs on QFX5200 and QFX5210.

[See [aggregated-ether-options](#).]

- **Logical interfaces created along with physical interfaces by default (QFX10000 and QFX5000 switches)**—On the QFX10000 line of switches, logical interfaces are created along with the physical **et-**, **sxe-**, **xe-**, and channelized **xe-** interfaces. In earlier releases, only physical interfaces are created.

On the QFX5000 line of switches, by default, logical interfaces are created on channelized **xe-** interfaces. In earlier releases, logical interfaces are not created by default on channelized **xe-** interfaces (**xe-0/0/0:1**, **xe-0/0/0:2**, and so on), but they are created on **et-**, **sxe-**, and nonchannelized **xe-** interfaces.

Layer 2 Features

- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—In Junos OS Release 19.1R2, we have introduced the configuration statement **input-native-vlan-push** at the **[edit interfaces interface-name]** hierarchy level. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the **input-vlan-map** configuration statement with a **push** operation is configured.

[See [input-native-vlan-push](#).]

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (QFX Series)**—Starting in Junos OS Release 19.1R2, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.

Services and Applications

- **Commit check for incomplete tunnel encapsulation configuration on flexible tunnel interface (FTI)**—Tunnel encapsulation configuration is mandatory for FTI interfaces. In Junos OS Release 19.1R2, when you try to commit any incomplete tunnel encapsulation configuration on an FTI, the CLI displays a commit error message.

Software-Defined Networking

- **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 might disrupt the PCEP session with the main active stateful PCE.

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

What's Changed in Release 19.1R1

EVPN

- Starting with Junos OS Release 19.1R1, the **no-arp-suppression** configuration statement is no longer supported on EX Series and QFX Series devices.

- **New options in show evpn instance command (QFX series)**—Starting in Junos OS Release 19.1R1, you can use the **show evpn instance esi-info** command to only display the ESI information for a routing instance and **show evpn instance neighbor-info** to only display the IP address of the EVPN neighbor for a routing instance. Information associated with the ESI, such as the route distinguisher, bridge domain, and IRB are filtered out.

Interfaces and Chassis

- **Commit error thrown when GRE interface and tunnel source interface are configured in different routing instances (QFX Series)**—In Junos OS Release 19.1R1, QFX Series switches do not support configuring GRE interface and the underlying tunnel source interface in two different routing instances. If you try this configuration, it will result in a commit error with the following error message:

error: GRE interface (gr-0/0/0.0) and its underlying tunnel source interface are in different routing-instances

error: configuration check-out failed

[See [Understanding Generic Routing Encapsulation](#).]

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (QFX Series)**—Starting in Junos OS Release 19.1R1, the **show lacp interfaces | display xml** command displays a new XML tag element <lacp-hold-up-state>. The <lacp-hold-up-state> displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces was in a single <lacp-hold-up-information> XML tag. Now, for each interface it is displayed in a separate <lacp-hold-up-information> XML tag.
- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, and later, QFX Series switches support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

In earlier Junos OS releases, when you use a Layer 2 logical interface configuration (units with **encapsulation vlan-bridge** configuration), then you must add the logical interface as part of a bridge domain or EVPN routing instance for the commit to succeed.

Network Management and Monitoring

- **sysName.0 MIB object displays the fully qualified domain name (QFX Series)**—Starting in Junos OS Release 19.1R1, the sysName.0 MIB object displays the fully qualified domain name. That is, if the hostname and domain name are configured on the system, both will show up for the sysName.0 MIB object: **host-name.domain-name**. Previously, only the hostname showed up.
[see [show snmp mib](#).]
- **NETCONF <kill-session> operation returns different values in <rpc-error> when the session identifier is equal to the current session ID (QFX Series)**—Starting in Junos OS Release 19.1R1, when you execute

the **<kill-session>** NETCONF operation and the session identifier is equal to the current session ID, the values of the **<error-type>** and **<error-tag>** elements in the resulting **<rpc-error>** are **application** and **invalid-value**, respectively. In earlier releases, the **<error-type>** and **<error-tag>** values are **protocol** and **operation-failed**.

[See [<kill-session>](#).]

Security

- **Syslog or log action on firewall drops packets (QFX5000 switches)**—Starting in 19.1R1, if you configure a syslog or log action on an ingress firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.
- **Firewall warning message (QFX5000 switches)**—Starting in 19.1R1, a warning message is displayed whenever a firewall term includes log or syslog with the accept filter action.

User Interface and Configuration

- **Options for monitor traffic interfaces statement added (QFX Series)**—Starting in Junos OS Release 19.1R1, the options **write-file** and **read-file** under the **monitor traffic** command are included in the visible CLI.

[See [monitor traffic](#).]

SEE ALSO

New and Changed Features 252
Known Behavior 272
Known Issues 276
Resolved Issues 282
Documentation Updates 303
Migration, Upgrade, and Downgrade Instructions 303
Product Compatibility

Known Limitations

IN THIS SECTION

- [EVPN | 273](#)
- [General Routing | 273](#)

- [Infrastructure | 274](#)
- [Layer 2 Features | 274](#)
- [MPLS | 274](#)
- [Platform and Infrastructure | 274](#)
- [Routing Protocols | 276](#)
- [Security | 276](#)
- [Virtual Chassis | 276](#)

Learn about known limitations in this release for QFX Series.

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

EVPN

- On QFX10000 switches configured as type-5 route peers, when only peer 1 advertises routes, that peer might not install the de-encapsulated next-hop route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any type-5 route. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. [PR1191092](#)
- When a VLAN uses an IRB interface as the routing interface, the vlan-id parameter must be set to **none** to ensure proper traffic routing. This issue is platform independent. [PR1287557](#)

General Routing

- On the QFX5100, if a scaled configuration involving a LAG interface, more than 3000 VLANs, and corresponding next hops is removed and a new configuration involving a LAG interface is applied at the same time, the new configuration might not take effect until the previous configuration has been deleted. During this time, FXPC might consume high CPU resources. No other system impact is observed. [PR1363896](#)
- The statement `pm4x25_line_side_phymod_interfa` might throw the error **ERROR: u=0 p=81 interface type 16 not supported by internal SERDES for this speed 50000**. This error message is seen when a channelization is detected in Junos OS Release 18.1R3. [PR1366137](#)
- If out-of-band management link is operated at a speed other than 1000 Mbps (for example, link peer is kept 10/100 Mbps) on QFX Series products within the Junos VM, the corresponding interface always reflects a speed of 1000 Mbps in all aspects. For example, in the output of the **show interfaces em0**

command. The actual speed in use will reflect only on the corresponding interface on the Linux host. [PR1401382](#)

- When doing an RFC MAC learning rate, we achieve a learning rate of only 13,000 MAC entries. For higher learning rate, we see some MACs are not learned, but sometimes the issue is not seen even at higher rates. [PR1403603](#)
- The maximum number of Layer 3 interfaces that can be configured on QFX5100 is 8000, QFX5200 is 8000, and QFX5110 is 12,000. [PR1406107](#)
- On a QFX5120, ARP might not get resolved for an untagged packet coming on an interface with **encapsulation ethernet-bridge** when this interface is in a VXLAN with the **encapsulate-inner-vlan** statement. [PR1454804](#)

Infrastructure

- If Junos OS panics with a file-system-related panic, such as 'dup alloc', recovery through the OAM shell might be needed. From the OAM shell, run 'fsck' on the root volume until it is marked clean. Only at this point it is safe to reboot to the normal volume. [PR1444941](#)

Layer 2 Features

- The **Targeted-broadcast forward-only** command does not broadcast the traffic. [PR1359031](#)
- xSTP configuration is not supported on **flexible-vlan-tagging** interfaces for any of the QFX5000 line of devices (QFX5100, QFX5110, QFX5200, QFX5210, QFX5120). [PR1414659](#)
- Hierarchical ECMP with VXLAN OVERLAY routing is not supported on QFX5000 devices. [PR1456594](#)

MPLS

- There is no warning message about Packet Forwarding Engine restart when the MPLS tunnel extend configuration is deleted. [PR1394722](#)

Platform and Infrastructure

- When the sFlow collector can be reached only through the Routing Engine, large samples due to heavy traffic can cause the Routing Engine CPU to become busy. [PR1332337](#)
- On QFX10002, QFX10008, and QFX10016, ND is incorrectly working on an IRB/Layer 3 interface with a discard filter. [PR1338067](#)
- Hardware watchdog does not work on QFX10008 and QFX10002-60C. [PR1343131](#)

- The 100-Gigabit Ethernet interface goes down after you configure and delete the Ethernet loopback configuration. [PR1353734](#)
- When vlan is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)
- On a QFX5120 switch with 288,000 MAC scale, the Routing Engine **show ethernet-switching table summary** command output shows the learned scale entries after a delay of around 60 seconds. [PR1367538](#)
- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and the QEMU hypervisor. You can recover the device by power cycling it. [PR1385970](#)
- These error logs are expected when routes pointing to target next hops are in turn pointing to hold next hops. These error logs appear for a short time. Later, when the next hop changes from hold next hop to a valid next hop, unicast next hops are walked again and updated with the appropriate weight and reroute counters, and no more error logs are seen. [PR1387559](#)
- Re-ARP request sent without a VLAN ID (so Routing Engine ARP fails). [PR1390794](#)
- The QFX5100 (Junos OS Release 19.1R1) uses SDK version 6.3.7. Unified ISSU with BST configuration is not supported and is a product limitation with regard to BCM chipset running on SDK 6.3.7. Even configuring BST after the unified ISSU might not work. As a workaround, restarting of Packet Forwarding Engine is required after the unified ISSU. For QFX5110, unified ISSU is not supported on Junos OS Release 19.1R1. [PR1395587](#)
- On QFX5120 system, the hardware link scan thread interrupt processing takes significant time due to firmware limitation. This results in greater than 50 ms convergence delay during MPLS FRR. [PR1403082](#)

Routing Protocols

- When an interface is configured with family **mpls**, one label is reserved for the explicit-null case. Only one label is used across the different MPLS interfaces for the explicit-null case. This label will only be deleted when all the interfaces with family **mpls** are deleted. So the maximum number of tunnels you can have is 1. [PR1418733](#)

Security

- —On QFX5000 platforms, if a syslog or log action is configured on a firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.

Virtual Chassis

- A Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop (>2 seconds) might occur. [PR1347902](#)

SEE ALSO

[New and Changed Features | 252](#)

[Changes in Behavior and Syntax | 267](#)

[Known Issues | 276](#)

[Resolved Issues | 282](#)

[Documentation Updates | 303](#)

[Migration, Upgrade, and Downgrade Instructions | 303](#)

Open Issues

IN THIS SECTION

- [EVPN | 277](#)
- [General Routing | 277](#)
- [Infrastructure | 280](#)
- [Interfaces and Chassis | 280](#)

- Layer 2 Features | 280
- MPLS | 281
- Platform and Infrastructure | 281
- Routing Protocols | 281

Learn about open issues in this release for QFX Series.

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

EVPN

- At times, when l2ald is restarted, a race condition occurs where VTEP notification comes in from the kernel before lo0. As a result, l2ald is unable to process the VTEP add request and gets stuck in an infinite loop. [PR1384022](#)

General Routing

- Layer 3 multicast traffic does not converge to 100 percent and continuous drops are observed after the downstream interface goes down or comes up or while an FPC comes online after FPC restart. This happens with multicast replication for 1000 VLANs or IRB interfaces. [PR1161485](#)
- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED stays unlit. [PR1317750](#)
- On the QFX10002-60C, the **filter** operation with **log** action is not supported for protocols other than Layer 2, IPv4, and IPv6. The following message is seen in the firewall logs: **Protocol 0 not recognized**. [PR1325437](#)
- Backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- QFX10000 platform drops the access point (AP) heartbeat packets, as result the WAP cannot work. [PR1352805](#)
- Interface flapping is observed only on a peer port with 100GBASE-LR4 optics in the warm boot stage of VMs during a unified ISSU process. As a workaround, do not use 100GBASE-LR4 during a unified ISSU. [PR1353415](#)
- A mib2d core file is generated in mib2d_write_snmpidx at snmpidx_sync.c on both active directories while bringing up a base traffic profile. [PR1354452](#)

- When a rpd reads next hops from the kernel on restart, in the INH -> FWD NH{List NH} -> {Chain NH} scenario, the rpd should not create old-style list next hop for the forwarding next hop. [PR1360354](#)
- On the QFX5100, if a scaled configuration involving a LAG interface, more than 3000 VLANs, and corresponding next hops is removed and a new configuration involving a LAG interface is applied at the same time, the new configuration might not take effect until the previous configuration has been deleted. During this time, FXPC might consume high CPU resources. No other system impact is observed. [PR1363896](#)
- The statement `pm4x25_line_side_phymod_interfa` might throw the error **ERROR: u=0 p=81 interface type 16 not supported by internal SERDES for this speed 50000**. This error message is seen when channelization is detected in Junos OS Release 18.1R3. [PR1366137](#)
- On the QFX10000 line of switches, with EVPN-VXLAN, the following error is seen: `expr_nh_fwd_get_egress_install_mask:nh type Indirect of nh_id: # is invalid`. [PR1367121](#)
- The user might not be able to stop the ZTP bootstrap, when an QFX10016 or an QFX10008 switch with more number of line cards is powered on with factory-default configuration. [PR1369959](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- On QFX10008 and QFX10016 platforms, traffic loss might occur because of switch modular failure on the Control Board (CB). This failure further causes all SIBs to be marked as faulty and causes FPCs to restart until Routing Engine switchover occurs. [PR1384870](#)
- 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 is received only on the 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](#)
- Control plane switch management (CPSM) daemon memory leak occurs in the VM host. It might also result in logrotate not working, and cause large CPSM log size. [PR1387903](#)
- In the Junos OS Release 18.4R1 branch, intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface. [PR1388082](#)
- DCPFE didn't come up in some instances of abruptly powering off and powering on the QFX5120-EX4650. To recover the device, power cycle it or reboot the host. [PR1393554](#)
- Layer 2 multicast and broadcast convergence is high while deleting and adding back the scale configurations of VLANs and VXLAN. [PR1399002](#)
- A QFX10000's FPC may restart if an operator configures VXLAN's VNI 0 identifier. [PR1401215](#)
- On the QFX5120, OVSDB-managed VXLAN experiences traffic loss. [PR1401943](#)
- If the USB storage device is not removed from the device after an upgrade, the system might come up and might reboot repeatedly. As a workaround, you need to manually change the boot sequence from the BIOS menu to select boot from SSD. [PR1404717](#)
- On the QFX10002, traffic drop is observed with MSTP configuration (65 instances and 64 interfaces with 3840 VLANs) [PR1408943](#)

- You might see multiple reconnect logs, JTASK_IO_CONNECT_FAILED, during the device initialization. There is no functionality impact due to these messages. These messages can be ignored. [PR1408995](#)
- Intermittently chassis alarms not raised after power-cycle of the device. Chassis alarms can be recovered by restarting lcmd from CLI - request app-engine service restart chassis-manager or, restart chassis-control [PR1413981](#)
- On QFX5110 and QFX5120 platforms, uRPF check in strict mode does not work properly. [PR1417546](#)
- When a bad optics is connected to the device that could inhibit EEPROM failure conditions or I2C read failure conditions, the device could end up in this condition. Please check the description of the issue mentioned in the PR. [PR1420874](#)
- When NSSU is done from Junos OS Release 18.1R3 to any later image on QFX5100 Virtual Chassis with LACP link protection configured, there might be around 5 minutes of traffic loss. Traffic loss is not seen during NSSU if the link protection configuration is not present. [PR1435519](#)
- When routing process is restarted, if the system is configured with EVPN service, memory of the Layer 2 address learning daemon (l2ald) increases by 4000 when you use **show system processes extensive | match l2ald**. [PR1435561](#)
- Unified ISSU might fail from Junos OS Release 17.2X75-D43.2 to some target versions on QFX5200 platforms. And dcpfe crash might be seen. [PR1438690](#)
- Unified ISSU will is not supported for QFX5200 from Junos OS Release 17.2X75-D4x to Junos OS Release 19.2R1. [PR1440288](#)
- On QFX10000 switches and EVPN-VXLAN (spine-leaf) scenario, QFX10000 spine switches are configured with VXLAN Layer 3 gateway (utilizing the virtual gateway) on an IRB interface. If you enable and then subsequently remove the VXLAN Layer 3 gateway on this IRB interface on one or some of these spine switches, traffic drop might be observed. As a workaround, if all virtual gateways are configured with an unique IPv4 or IPv6 MAC address, this issue does not occur. [PR1446291](#)
- Whenever any member in RSPAN VLAN is removed from that VLAN, you must reconfigure the analyzer session for that RSPAN VLAN. [PR1452459](#)
- In an EVPN-VXLAN with service-provider style configuration, if the VLAN name associated with access ports is changed, then the virtual bridge domain might not be created. This is because the bridge domain add notification for the new VLAN comes before the bridge domain delete notification for the old VLAN. Because of this, virtual bridge domain might not be created and MAC's might not be learned. [PR1454095](#)
- After changing the VLAN name on the trunk interface, local host MAC learning will be hold for more than 30 seconds. [PR1454274](#)
- Enabling maintenance configuration on MH device without disabling the ESI link might lead to a traffic loop. We recommend that you disable the ESI link instead of the maintenance configuration on the MH device. [PR1456349](#)
- On QFX5110, FEC errors might be seen on the other side. [PR1457266](#)

- Change of VTEP source address by changing the loopback address will trigger reduction in Vport and VNI. [PR1467158](#)
- On QFX5100 Virtual Chassis, 10-Gigabit VCP ports will not be active. [PR1494980](#)

Infrastructure

- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock (No such file or directory)** messages are seen during FTP. [PR1315605](#)

Interfaces and Chassis

- Flooding of ARP reply unicast packets is seen as a result of an ARP request sent for the device's VRRP MAC address. The ARP reply that is flooded in the VLAN by the device has the correct DMAC of the originator of the ARP request. That is, the ARP reply is flooded but with the correct unicast DMAC. The ARP reply is not broadcasted. [PR1454764](#)

Layer 2 Features

- If the access-side interfaces are used as SP-style interfaces, when a new logical interface is added and if there is already a logical interface on the physical interface, there is a 20-50 ms traffic drop on the existing logical interface. [PR1367488](#)
- On the QFX5120, during new tenant addition, there might be a few transient packet drops (2-15 packets) for a couple of random intra-VNI traffic streams in an EVPN-VXLAN topology for the existing tenants. The drop is almost negligible and is autorecovered. [PR1455654](#)
- With QFX5110 and QFX5200 platforms, if storm control IS enabled on the interfaces along vxlan configuration, storm control will not get effected with ARP REQ packets coming more than storm control threshold. [PR1469837](#)
- On QFX5110 and QFX5120 platforms, changing the lo0 IP address might sometimes either result in stale entry of IP in mpls_entry table or missing IP entry, which results in traffic drop for VXLAN traffic. [PR1472333](#)

MPLS

- A lingering RSVP state might keep some labeled routes programmed in the Packet Forwarding Engine longer than they should be. This RSVP state eventually expires and then delete the RSVP MPLS routes from FIB. However, traffic loss is not anticipated because of this lingering state or the corresponding label routes in the FIB. In the worst case, in a network, where there is persistent link flapping going on, this lingering state might interfere with the LSP scale being achieved. [PR1331976](#)

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced I2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)

Routing Protocols

- Higher convergence time for LFA with BFD in Junos OS Release 18.1. [PR1337412](#)
- The **pimd_rtrequest_v4(1133), IS_MASTER_RE: 1, Process: rpd, RTM_ID: 5, error: 17, errmsg: rt exists; ifindex = 340** error messages are cosmetic and expected logs. These logs are not harmful and have no functional impact; they just show the state of PIM register messages. These logs are already LOG_DEBUG for external builds, and you do not need to make any change in any of the components. [PR1371431](#)
- When a MOLEX QSFP+ DAC cable is connected to the QFX5210, the link might not come up. A DCPFE might generate a core file and the fxpc process might not come up. [PR1397158](#)
- There is no functionality impact because of the following error message:
BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(),128:I3 nh 6594 unintsall failed in h/w with Mini-PDT base configurations. [PR1407175](#)
- In Junos OS Release 19.1R3, the MUX state of LACP interface will not change sometimes when **force-up** is onfigured. [PR1484523](#)

SEE ALSO

[New and Changed Features | 252](#)

[Changes in Behavior and Syntax | 267](#)

[Known Behavior | 272](#)

[Resolved Issues | 282](#)

[Documentation Updates | 303](#)

[Migration, Upgrade, and Downgrade Instructions | 303](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.1R3 | 282](#)
- [Resolved Issues: 19.1R2 | 288](#)
- [Resolved Issues: 19.1R1 | 298](#)

Learn which issues were resolved in the 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: 19.1R3

Class of Service (CoS)

- Shaping does not work after the reboot if the **shaping-rate** is configured. [PR1432078](#)
- The traffic is placed in network-control queue on extended port even if it comes in with different dscp marking. [PR1433252](#)
- On QFX5120-32C, while moving unicast traffic to the multicast queue through the MF classifier, the **show interface queue <>** command does not display any statistics. [PR1459281](#)

EVPN

- ARP request/NS might be sent back to the local segment by the DF router. [PR1459830](#)
- The rpd might crash after the EVPN-related configuration is changed. [PR1467309](#)

Forwarding and Sampling

- The l2ald process might experience memory leak on platforms running Junos OS. [PR1455034](#)
- Type 1 ESI/AD route might not be generated locally on EVPN PE in the all-active mode. [PR1464778](#)

General Routing

- On QFX5100 switches, the LR4 QSFP transceiver might take up to 15 minutes to come up after a Virtual Chassis reboot. [PR1337340](#)
- The 10-Gigabit fiber interfaces might flap frequently when they are connected to other vendor's switch [PR1409448](#)

- The optic module 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](#)
- Part of routes could not be provided into the Packet Forwarding Engine when both IPv4 and IPv6 are used. [PR1412873](#)
- The **show interface** command indicates **Media type: Fiber** on QFX5100-48T running the QFX 5e Series image. [PR1419732](#)
- Ports might get incorrectly channelized if they are already 10-Gbps ports and they are channelized to 10-Gbps again. [PR1423496](#)
- The host-bound traffic might be dropped after performing change configuration related to prefix-list. [PR1426539](#)
- The l2cpd process might crash and generate a core file when the interfaces are flapping. [PR1431355](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- The line card might crash because the installed SFP-T module is not supported. [PR1432809](#)
- BGP neighbourship might not come up if the MACsec feature is configured. [PR1438143](#)
- The EX4600 or QFX5100 Virtual Chassis might not come up after the Virtual Chassis port fiber connection is replaced with DAC cable. [PR1440062](#)
- MAC addresses learned on the RTG might not be aged out after a Virtual Chassis member is rebooted. [PR1440574](#)
- CRC error might be seen on the VCPs of the QFX5100 Virtual Chassis. [PR1449406](#)
- The FPC does not restart immediately after the system is rebooted, which might cause packet loss. [PR1449977](#)
- CoS classification does not work on QFX10000. [PR1450265](#)
- FPC core files might be seen after changing the configuration of PTP or Synchronous Ethernet. [PR1451950](#)
- The l2ald and eventd hogs 100 percent after **clear ethernet-switching table** is issued. [PR1452738](#)
- The classifier configuration does not get applied to the interface in an EVPN-VXLAN environment. [PR1453512](#)
- The **show chassis led** command shows a wrong status. [PR1453821](#)
- On QFX5100-VC, the vgd process hogs the CPU without the **switch-options vtep-source-interface lo0.0** configuration. [PR1454014](#)
- After a reboot, the master FPC might come up in master state again instead of backup state. [PR1454343](#)
- Dcpfe should crash because usage of data is not NULL terminated on QFX5000. [PR1454527](#)
- On QFX10002-60c, in an EVPN-VXLAN environment, the MAC+IP count is shown as zero. [PR1454603](#)
- The untagged hosts ARP/NS requests might not be resolved when it is connected on encapsulation ethernet-bridge\ interface. [PR1454804](#)

- A firewall filter might not be able to be applied in a particular VC/VCF member as TCAM space running out. [PR1455177](#)
- In a 16+ member QFX5100 VCF, the **FROM** column under the **show system users** command output reports **feb0**, **feb1**, **feb2**, and **feb3** for **fpc16**, **fpc17**, **fpc18**, and **fpc19** respectively. [PR1455201](#)
- The PFC feature does not work on the QFX10000 line of switches. [PR1455309](#)
- The cosd crash might be observed if the **forwarding-class-set** is directly applied on the child interface of an aggregated Ethernet interface. [PR1455357](#)
- Link-up delay and traffic drop might be seen on mixed SP L2, SP L3, and EP L2 type configurations. [PR1456336](#)
- The Packet Forwarding Engine process might crash after a Routing Engine switchover on QFX10000 platforms. [PR1457414](#)
- Overtemperature SNMP trap messages are displayed after update even though the temperatures are within the system thresholds. [PR1457456](#)
- QFX5110 switches port 51 has one LED blinking amber Junos OS Release 19.1R1.6. [PR1457516](#)
- Dual tag Q-in-Q not working with EVPN-VXLAN [PR1458206](#)
- On a QFX5210, the LED does not light on port 64 and 65 after upgrade to Junos OS Release 19.2R1. [PR1458514](#)
- The BPDU packet might be looped between leaf DF switch and non-DF switch and causes traffic blocking. [PR1458929](#)
- JDI-QFX5200-REGRESSION-SWITCHING-QFX5200: DHCPv6 LDRA relay bounded count is not as expected after DHCP is configured. [PR1459499](#)
- The fxpc process might crash due to several BGP IPV6 session flaps. [PR1459759](#)
- The **forwarding** option is missing from the **routing-instance type** configuration. [PR1460181](#)
- The **accept-source-mac** feature with VXLAN is not working on QFX5000 platforms. [PR1460885](#)
- The entPhysicalTable MIB is not fetching expected data on QFX10002-72Q / 36Q platforms. [PR1462582](#)
- The fxpc process might generate core files when changing the MTU in a VXLAN scenario with firewall filters applied on QFX5000 platforms. [PR1462594](#)
- JDI-RCT : QFX 5100 VC/VCF : Observing the error **BRCM-VIRTUAL,brcm_vxlan_walk_svp(),6916:Failed to find L2-iff for ifl:** while cleaning up Evpan-VxLAN configurations with Mini-PDT base. configurations. [PR1463939](#)
- The FPC might restart during run time on PTX10000 or QFX10000 platforms. [PR1464119](#)
- The dcpfe might crash when changing the firewall filter on QFX5000 platforms. [PR1464352](#)
- The interface might not come up on FPC restart on QFX10000 platforms. [PR1464650](#)
- PEM is not present spontaneously on QFX5210. [PR1465183](#)

- A 10G interface might not come up on QFX5100-48T switches or negotiate at speed 1G when connected with Broadcom 57800-T daughter card. [PR1465196](#)
- The QSFP-100G-PSM4 could not be correctly identified on QFX5200 or QFX5110 platforms. [PR1465214](#)
- The physical interface of aggregated Ethernet might take time to come up after disable/enable. [PR1465302](#)
- The broadcast and multicast traffic might be dropped over IRB or LAG interface in QFX/EX VC scenario. [PR1466423](#)
- BGP Open messages with specific types of BGP Optional Capabilities causing BMP messages not been encoded correctly when sent to the BMP Collector. [PR1466477](#)
- EBUF parity interrupt is not seen on QFX10K/PTX platforms. [PR1466532](#)
- IPv6 traffic might get dropped in Layer 3 VPN network. [PR1466659](#)
- Slow packet drop might be seen on QFX5000 switches. [PR1466770](#)
- DHCPvX ACK messages do not receive response to broadcast INFORM packets with Junos OS Release 19.1R2.3. [PR1467182](#)
- Ingress drops are to be included at the CLI from interface statistics and added to 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 Data Center environment. [PR1468732](#)
- If continuous interface flaps at ingress/egress of PE devices, IP routed packets might be looped on the MPLS PHP node. [PR1469998](#)
- Incorrect counter value for "Arrival rate" and "Peak rate" for ddos commands. [PR1470385](#)
- The speed 10m might not be configured on the GE interface. [PR1471216](#)
- Traffic loss might occur when a VTEP source interface is configured in multiple routing instances. [PR1471465](#)
- QFX5K: Egress PAACL size will be half in Junos OS Release 19.4R1. [PR1472206](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- Detached interface in LAG processing xSTP BPDUs. [PR1473313](#)
- An l2ald crash might be seen when around 16000 VLAN IDs sharing the same VXLAN tunnel and when the Packet Forwarding Engine is rebooted. [PR1473521](#)
- On QFX5000 switches, RIPv2 routes that are being forwarded across an Layer 2 circuit connection are dropped. [PR1473685](#)
- Continuous error log messages might be raised on QFX5000 switches in an EVPN-VXLAN scenario. [PR1474545](#)
- Layer 2 circuit might fail to communicate through VLAN 2 on QFX5000 switches. [PR1474935](#)

- MACsec traffic over Layer 2 circuit might not work on QFX10000, PTX10000, and PTX1000 platforms after upgrading from Junos OS Release 15.1 to later versions. [PR1475089](#)
- DAC cables are not being properly detected in Packet Forwarding Engine in QFX5200 on Junos OS Release 18.4R2-S2.4. [PR1475249](#)
- There might be traffic drop on QFX5110/5120 switches acting as leaf switch in a multicast environment with VxLAN. [PR1475430](#)
- QFX Platforms are exhibiting invalid PFE PG counter pairs to copy, src 0xfffff80, dst 0. [PR1476829](#)
- Continuous Error logs on the device: `prds_ptc_wait_adoption_status: PECHIP[1] PTC[1]: timeout on getting adoption valid bit[8] asserted`. [PR1477192](#)
- ARP request packets for unknown host might get dropped in remote PE in EVPN-VXLAN scenario. [PR1480776](#)

Interfaces and Chassis

- Traffic might be forwarded to wrong interfaces in an MC-LAG scenario. [PR1465077](#)
- Executing commit might hang because of the struck dcd process. [PR1470622](#)
- Commit error is not thrown when member link is added to multiple aggregation group with different interface specific options. [PR1475634](#)

Layer 2 Features

- Storm control configuration may be disabled for the interface. [PR1354889](#)
- Physical layer and MAC/ARP learning might not work for copper base SFP-T on QFX5100/QFX5110/EX4600. [PR1437577](#)
- The LLDP function might fail when a Juniper device connects to a non-Juniper one. [PR1462171](#)
- QFX5110-32Q: Some of the MAC addresses are missing from the MAC table in software after restarting the Packet Forwarding Engine. [PR1467466](#)
- The fxpc core might be seen when committing the configuration all together, for example, after a reboot. [PR1467763](#)
- Ingress traffic might be blackholed if underlying interfaces flap in EVPN/VXLAN scenario. [PR1469596](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)

MPLS

- On QFX10002 switches, the `show mpls static-lsp | display xml` command produces invalid XML. [PR1469378](#)
- Traffic might be lost over QFX5100 switches acting as a transit PHP node in the MPLS network. [PR1477301](#)

Platform and Infrastructure

- The SLAX script might be lost after upgrading software. [PR1479803](#)

Routing Protocols

- OSPF VRF sessions might take a long time to come up when the host table is full and host routes are in the LPM table. [PR1358289](#)
- Invalid VRRP mastership election is seen on QFX5110-VC peers. [PR1367439](#)
- The traffic with destination UDP port 520 (RIP) or 521 (RIPng) gets dropped on QFX5000 and EX4600 switches. [PR1429543](#)
- Host destined packets with filter log action might not reach to the routing engine if log/syslog is enabled. [PR1379718](#)
- QFX5100 : BGP v4/v6 convergence & RIB install/delete time degraded in Junos OS Releases 19.1R1/19.2R1/19.3R1/19.4R1. [PR1414121](#)
- On the QFX5000 Series platforms acting as Layer 2 circuit PE (tunnel terminating node), if VLAN 2 is used for Layer 2 circuit communication with CE node, the VLAN 2 packets might be dropped on PE. [PR1474935](#)
- CRC errors might be seen on QFX5100 Virtual Chassis. [PR1444845](#)
- Core files might be generated when the EVPN Type-5 routing instance is being added or removed. [PR1455547](#)
- The egress interface in the Packet Forwarding Engine for some end hosts might not be correct on the Layer 3 gateway switch after it is rebooted. [PR1460688](#)
- On QFX 5100 Virtual Chassis or Virtual Chassis Fabric, the **brmc_ipmc_route_counter_delete:3900Multicast stat destroy failed (-10:Operation still running)** error is seen after performing unified ISSU with Mini-PDT base configurations. [PR1460791](#)
- When deleting an IRB I terface on the Layer 3 gateway, IRB does not get removed from the Paacket Forwarding Engine and might cause traffic to be discarded silently to IRB MAC address. [PR1463092](#)
- The mcsnoopd crash might be seen if one BD/VLAN is configured as part of EVPN and it has any multicast router interfaces (static/dynamic). [PR1468737](#)
- Traffic might not be forwarded over an ECMP link in an 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 get forwarded in a VRRP scenario. [PR1477073](#)

Resolved Issues: 19.1R2

Class of Service (CoS)

- On QFX10008, FPC0 crashes and generates a core file after running the Packet Forwarding Engine command **show cos sched-usage**. [PR1449645](#)
- **show cos scheds-per-pfe** and **show cos pfe-scheduler-ifds** Packet Forwarding Engine commands will restart forwarding planes on QFX10008 switches [PR1452013](#)

EVPN

- The rpd process crashes with EVPN type-3 route churn. [PR1394803](#)
- The **show evpn instance extensive esi** command does not filter the output of desired ESI or neighbor information of an EVPN instance. [PR1402175](#)
- ARP entry is still pointing to failed VTEP after PE-CE link fails for multihomed remote ESI. [PR1420294](#)
- Multicast MAC address might be learned in the Ethernet switching table on QFX5000 and QFX10000 platforms with EVPN-VXLAN configured. [PR1420764](#)
- The device may proxy the ARP probe packets in an EVPN environment. [PR1427109](#)
- Unexpected next-hop operation error from Kernel to L2ald in a Layer 2 gateway during the MAC movement operation. [PR1430764](#)
- Asynchronous between ARP table and Ethernet switching table happens if EVPN ESI link flap multiple times. [PR1435306](#)
- The multihomed mac-ip table entry might not be cleaned when host MAC is deleted from MAC table. [PR1436712](#)
- Configuring ESI on a single-homed 25G port might not work. [PR1438227](#)
- 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 multi-home ESI. [PR1446957](#)
- VLAN configuration change with l2ald restart might cause Kernel synchronization issues due and impact forwarding. [PR1450832](#)
- When there is a VXLAN with VLAN ID of 2 on a QFX5100, ARP will not get resolved. [PR1453865](#)
- ARP request/NS might be sent back to the local segment by DF router. [PR1459830](#)

Forwarding and Sampling

- Commit error and dfwd core file might be observed when applying a firewall filter with action "then traffic-class" or "then dscp". [PR1452435](#)

General Routing

- Certain QFX Series devices are vulnerable to 'Etherleak' memory disclosure in Ethernet padding data. [PR1063645](#)
- The 1G copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. [PR1286709](#)
- On QFX10002-60C, commit might be denied when L2 and L3/L4 mix-match conditions are configured on a L2 filter. [PR1326715](#)
- On QFX5100 platforms, LR4 QSFP might 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. [PR1344395](#)
- On QFX5210, when filter with routing instance is applied to a family inet logical interface, traffic gets discarded on unrelated interfaces. [PR1364020](#)
- On QFX5120 and EX4650 line of switches, the convergence delay between PE1 and P router link is more than expected delay value. [PR1364244](#)
- Traffic spikes generated by IPFIX might be seen on QFX10002. [PR1365864](#)
- The backup member switch might fail to become the master switch after switchover on QFX5100, QFX5200, and EX4600 Virtual Chassis platform. [PR1372521](#)
- RIPv2 update packets might not send with IGMP snooping enabled. [PR1375332](#)
- New configuration statement 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 updated for QFX5110. [PR1382803](#)
- Static default route with next-table inet.0 does not work. [PR1383419](#)
- The rpd end up with krt queue stuck might be seen in vrf scenario. [PR1386475](#)
- Error message **portmod_port_core_access_get: Invalid parameter** seen in log messages. [PR1388591](#)
- ARP received on SP-Style interface not sent to all RVTEPs in case of QFX5100 VC only, normal BUM traffic works fine. [PR1388811](#)
- When **show** command is taking a long time to display results, the STP might change states because BPDUs are no longer processed and cause lots of outages. [PR1390330](#)
- On QFX5110 fan LED turns Amber randomly. [PR1398349](#)
- The interrupt process consumes high CPU because of the `intr{swi4: clock (0)}` on QFX5100-48t-6Q running a QFX5100 Series image and Junos OS Release 18.x code. [PR1398632](#)

- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. [PR1399683](#)
- On QFX5100, traffic initiated from a server connected to an interface will be dropped at the interface on the switch if the interface was configured with family ethernet-switching with VXLAN and the configuration is changed to family inet. [PR1399733](#)
- On QFX5110 platforms, from Junos OS Release 17.3 and later, the interfaces with SFP-LX10 transceivers and auto-negotiation enabled(default configuration) might be down. [PR1399878](#)
- On QFX5120-32C Error logs for flex counter seen with GRE configuration. [PR1400515](#)
- QSFP-100GBASE-SR4/LR4 might take a long time to come up after disabling interface or reboot [PR1402127](#)
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface [PR1403528](#)
- Executing command "request system configuration rescue save" may fail with error messages [PR1405189](#)
- DHCP Not working for some clients in dual AD fusion setup on EP ports. [PR1405495](#)
- Ping over loopback might not work over TYPE 5 tunnel on QFX10000 platforms [PR1405786](#)
- QFX5120 : In VxLAN-EVPN configuration, transition from collapsed to non-collapsed L2/L3 GW and vice versa needs switch reload [PR1405956](#)
- QFX5200/5100 might not be able to send out control plane traffic to the peering device [PR1406242](#)
- QFX10002 showing error fpc0 prds_ptc_clear_all_pulse_and_samples:
prds_ptc_clear_all_pulse_and_samples PE 4 PTC 2: after clearing sample, sample still valid 1 [PR1407095](#)
- No inner VLAN tag is added even with **input-vlan-map push** configured on QFX10000 platforms. [PR1407347](#)
- MAC address movement might not happen in Flexible Ethernet Services mode when family inet/inet6 and vlan-bridge are configured on the same physical interface. [PR1408230](#)
- Fan failure alarms might be seen on QFX5100-96S after upgrade to Junos OS Release 17.3R1. [PR1408380](#)
- Restarting line card on QFX10008 and QFX10016 with MC-LAG enhanced-convergence, the intra-vlan traffic might silently be dropped or discarded. [PR1409631](#)
- LLDP memory leak when ieee dcbx packet is received in auto-neg mode followed by another dcbx packet with none of ieee_dcbx tlvs present. [PR1410239](#)
- On QFX5120 platform with QSFP-100G-PSM4 transceiver, because of the timing fault on FPGA (Field Programmable Gate Array) hardware, the link might go down as TX laser being disabled. [PR1410687](#)
- On EX2300-24P or QFX5100, error message **dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family** might be seen. [PR1410717](#)
- 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](#)
- Storm control not shutting down mc-ae interface. [PR1411338](#)

- The spfe on satellite device in Junos fusion setup might crash and it could cause the satellite device to get offline. [PR1412279](#)
- PEM alarm for backup FPC will be remained on master FPC though backup FPC is detached from Virtual Chassis. [PR1412429](#)
- Junos PCC might reject PCUpdate/PCCreate message if there is a metric type other than type 2. [PR1412659](#)
- On QFX5120 devices route table is full for IPv6 routes in some scenario. [PR1412873](#)
- QFX5K: EVPN / VxLAN: Multicast NH limit is 4K. [PR1414213](#)
- The QFX10002 might stop forwarding packets after the "chassis-control" process restarts. [PR1414434](#)
- VC Ports using DAC may not establish link on QFX5200 [PR1414492](#)
- DC output information is missing in the "show chassis environment pem" output for whitebox [PR1414703](#)
- VXLAN Encapsulation nexthop (VENH) doesn't get installed during BGP flap or restart routing. [PR1415450](#)
- Changing FEC parameter for 100GE interfaces with QSFP-100GBASE-SR4 optics is not taking effect [PR1416376](#)
- Two instances of Junos are running after Junos upgrade to 18.1R3-S3.7 [PR1416585](#)
- The dcpfe crash might be seen in EVPN-VXLAN scenario [PR1416925](#)
- MAC learning might not happen on trunk mode interface in EVPN/MPLS scenario. [PR1416987](#)
- Extended remote port mirroring traffic is not tagged when the output interface is a trunk port. [PR1418162](#)
- Traffic loss might be seen on the ae interface on QFX10000 platforms [PR1418396](#)
- Traffic loss might be seen after NSSU operation. [PR1418889](#)
- Rebooting QFX5200-48Y using "request system reboot" doesn't take physical links offline immediately [PR1419465](#)
- The 100G PSM4 optics connected ports go down randomly during a repeated power cycle [PR1419826](#)
- Traffic drop might be observed when transit static LSP is configured on EX4650 and QFX5120 platforms [PR1420370](#)
- Ping fails over Type-5 tunnel on IRB interfaces under EVPN-VXLAN scenario [PR1420785](#)
- An interface may go to downstate on QFX10000/PTX10000 platform [PR1421075](#)
- QFX5120-32C: DHCP binding on client might fail when QFX5120-32C acting as DHCP server, this is seen only for channelized port [PR1421110](#)
- fusion: ETS config not applied on non-cascade ports when AD is rebooted [PR1421429](#)
- BFD might stuck in slow mode on QFX10002/QFX10008/QFX100016 platform [PR1422789](#)
- QFX5100-48T 10G interface might be auto-negotiated at 1G speed instead of 10G [PR1422958](#)

- The interface can not get up when the remote-connected interface only supports 100M in QFX5100 VC setup [PR1423171](#)
- IPv6 multicast traffic received on one VC member might be dropped when egressing on other VC member if MLD snooping is enabled [PR1423310](#)
- ON QFX5120-32C , BUM traffic coming over irb underlay interface gets dropped on destination vtep in PIM based VxLAN [PR1423705](#)
- Traffic is dropped after FPC reboot with AE member links deactivated by remote device [PR1423707](#)
- The Jflow export might fail when channelization is configured on FPC QFX10000-30C [PR1423761](#)
- Ping over EVPN type-5 route to QFX10000 does not work. [PR1423928](#)
- All interfaces will be down and the dcpfe will get crash if SFP-T is inserted on QFX5210. [PR1424090](#)
- IPv6 communication issue might be seen after passing through QFX10002-60C platforms. [PR1424244](#)
- QFX5120 QSFP-100G-PSM4 become undetected and come back up as channelized interfaces. [PR1424647](#)
- All interfaces creation failed after NSSU [PR1425716](#)
- The dcpfe or PFE might not start on AS7816-64X and QFX5K TVP platform devices. [PR1426737](#)
- QFX5210: Received LLDP frames on em0 not displaying in LLDP neighbor output [PR1426753](#)
- Heap memory leak might be seen on QFX10000 platforms [PR1427090](#)
- CRC errors can be seen when other manufacturer device is connected to QFX10000 with QSFP-100GBASE-LR4-T2 optics. [PR1427093](#)
- Rebooting or halting VC member might cause 30 seconds down on RTG link. [PR1427500](#)
- QFX5100-VCF 'rollback' for uncommitted configuration takes 1 hour. [PR1427632](#)
- On QFX10000 platforms certain interfaces might go to down state. [PR1427883](#)
- The dcpfe process might crash and restart in MC-LAG scenario when the ARP/NDP next hop is changed. [PR1427994](#)
- QFX5120-48Y interface with optic "QSFP-100GBASE-ER4L" is not coming up in Junos OS Release 18.3R1-S2.1. [PR1428113](#)
- Licenses used flag for ovsdb on **show system license** might not be flagged even though ovsdb is configured and working. [PR1428207](#)
- In correct display of MAC/MAC+IP and count values, after setting **global-mac-limit** and **global-mac-ip-limit**. [PR1428572](#)
- EVPN-VXLAN I2ald process might generate a core file when number of VXLAN HW IFBDS exceeds the maximum limit of 16382. [PR1428936](#)
- On QFX10008 after Routing Engine switchover, the LED status is not set for missing fan tray. [PR1429309](#)
- DHCP-relay may not work in an EVPN-VxLAN scenario [PR1429506](#)

- Extra incorrect MAC move might be seen when the host moves continuously between the different ESI. [PR1429821](#)
- Interface on QFX5120 switches does not come up after the transceiver is replaced with one having different speed. [PR1430115](#)
- In collapsed VGA4 script ping on shared ESI R6 to R7 IRB address is failing. [PR1430327](#)
- Traffic impact might be seen on QFX10000 platforms with interface **hold-down timer** configured. [PR1430722](#)
- On QFX Series switches, the **Validation of meta data files failed** message is seen on hypervisor. [PR1431111](#)
- **SIB Link Error** error message is detected on a specific Packet Forwarding Engine might cause complete service impact. [PR1431592](#)
- The dcpfe might crash on all line cards on QFX10000 in scaled setup. [PR1431735](#)
- The et- interfaces might not come up on QFX10000-60S-6Q. [PR1431743](#)
- All ingress traffic might be dropped on 100m fixed speed port with **no-auto-negotiation** enabled. [PR1431885](#)
- The optical power of interface might gradually reduce the optical power for almost 3 minutes after issuing **request system reboot at now** on QFX5110 and QFX5120. [PR1431900](#)
- L2 traffic drop on QFX10000 with interface MTU lower than 270 bytes. [PR1431902](#)
- Outer VLAN tag may not be pushed in the egress VXLAN traffic towards the host for QinQ scenario [PR1432703](#)
- Traffic loss might be seen on QFX10000 platforms using LC1105. [PR1433300](#)
- L3 filters applied to PVLAN IRB interface might not work after ISSU. [PR1434941](#)
- SIB/FPC link rrror alarms might be observed on QFX10000 due to a single CRC. [PR1435705](#)
- 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 trap (1.3.6.1.4.1.2636.3.69.1.0.0.1) sent on QFX5110 restart. [PR1436968](#)
- The FPC might crash if both the aggregated Ethernet bundle flapping on local device and the configuration change on peer device occur at the same time. [PR1437295](#)
- QFX5110, QFX5200, QFX5210 There is no jnxFruOK SNMP trap message when only the Power cable is disconnected and connected back. [PR1437709](#)
- 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 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 struck 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 come 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 rebooted. [PR1440574](#)
- Layer 2 and Layer 3 traffic drop is seen on disabling and then re-enabling mclag. [PR1440732](#)
- On QFX5110 switches, Layer 2 and Layer 3 logical interfaces on physical interfaces flexible-ethernet-services VXLAN passing over Layer 2 physical breaks, Layer 3 P2P communication. [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](#)
- On QFX5200, the error **DCBCM[bcore_init]: ioctl call failed ret:0** failure message is observed 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](#)
- Traffic is 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** does not 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](#)
- Except one aggregated Ethernet member link, the other links do not send out sFlow sample packets for ingress traffic. [PR1449568](#)

- On QFX5120, the incoming L3-encapsulated packets are dropped on L3VPN MPLS PE-CE interface. [PR1451032](#)
- vgd core files might be generated on any platforms supporting OVSD. [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 does not get applied to the interface in an EVPN/VXLAN environment. [PR1453512](#)
- **show chassis led** shows incorrect status. [PR1453821](#)
- QFX10002-60c EVPN-VXLAN, MAC+IP count is shown as zero. [PR1454603](#)
- 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 up after update even though the temperature are within the system thresholds. [PR1457456](#)
- The BPDU packet might be looped between leaf DF switch and non-DF switch and blocks traffic. [PR1458929](#)
- The forwarding option is missed in routing instance type. [PR1460181](#)
- In EVPN scenario memory leak might be observed when **proxy-macip-advertisement** is configured. [PR1461677](#)

Interfaces and Chassis

- Changing the value of mac-table-size to default might reboot all the FPCs. [PR1386768](#)
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. [PR1402606](#)
- The logical interfaces in EVPN routing instances might flap after committing configurations. [PR1425339](#)
- VRRP-V6 state is flapping with init and idle states after configuring vlan-tagging. [PR1445370](#)
- An ARP entry is not learned at one of mc-lag device at QFX10000. [PR1449806](#)
- Flooding of ARP reply unicast packets for switch VRRP MAC address through every port in VLAN. [PR1454764](#)
- The traffic might be forwarded to incorrect interfaces in MC-LAG scenario. [PR1465077](#)

Layer 2 Ethernet Services

- LACP PDU might be looped towards peer MC-LAG nodes. [PR1379022](#)
- BFD might flap when some of underlay ECMP interfaces are disabled in the leaf nodes. [PR1416941](#)

- The malfunction of core isolation feature in EVPN-VxLAN scenarios causes traffic drop. [PR1417729](#)
- The DHCP decline packets are not forwarded to DHCP server when forward-only is set within dhcp-reply. [PR1429456](#)

Layer 2 Features

- Storm control configuration may be disabled for the interface. [PR1354889](#)
- VxLAN next hop entry leak issue on QFX5000 platforms. [PR1387757](#)
- With IGMP snooping enabled on the LEAF switches, multicast traffic is forwarded to VLAN/VNI which does not have active receiver. [PR1388888](#)
- On QFX Series line of switches, the following error message **Failed with error (-7) while deleting the trunk 1 on the device 0** is observed when adding or removing local-bias setting on SP style LAG interface. [PR1393276](#)
- QinQ might be malfunctioning if vlan-id-lists are configured. [PR1395312](#)
- On QFX5000 line of switches, symmetric hashing can be configured with the hashing options, though it cannot be enabled and stored in the Junos OS configuration. [PR1397229](#)
- On QFX Series EVPN-VXLAN, unicast IPv6 NS message floods on L3 gateway. Therefore, both IPv4 and IPv6 traffic drops on L2SW. [PR1405814](#)
- IGMP-snooping on EVPN-VXLAN might impact OSPF hello packets flooding after VTEP leaf reboot. [PR1406502](#)
- QFX5110 Virtual Chassis generates DDoS messages of different protocols on inserting a 1G/10G SFP or forming VCP connection. [PR1410649](#)
- Packet loss might be seen when one of the Spine switch fails or reboots. [PR1421672](#)
- Stale entries might be observed in a layer 3 VXLAN gateway scenario. [PR1423368](#)
- The fxpc might continually crash when firewall filter is applied on a logical unit of a dsc interface. [PR1428350](#)
- ERPS nodes do not converge to IDLE state after failure recovery or reboot. [PR1431262](#)
- EVPN-VXLAN non-collapsed JTASK and multimove depth failed errors seen after HALT. [PR1434687](#)
- Transit DHCPv6 packets might be dropped on QFX5100 and QFX5200 platforms. [PR1436415](#)
- The MAC/ARP learning might not work for copper base SFP-T on QFX5100 and QFX5110. [PR1437577](#)
- QFX5000 switches are not properly hashing MPLS transit traffic from VXLAN to L2 LAG. [PR1448488](#)
- Unequal LAG hashing is seen on QFX5100 running Junos OS Release 14.1X53-D28.17. [PR1455161](#)

MPLS

- Traffic loss might be observed after changing configuration under **protocols mpls** in ldp-tunneling scenario. [PR1428081](#)

- The I2circuit traffic might silently get dropped or discarded at **EVPN SPINE/MPLS LSP TRANSIT** device if VXLAN access interface flaps on remote PE node. [PR1435504](#)
- Packet loss might occur when ECMP resilient-hash is enabled on QFX5200 switch. [PR1442033](#)

Platform and Infrastructure

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

Routing Protocols

- Some storm control error logs might be seen on QFX Series platforms. [PR1355607](#)
- Value added in hexa after unknown ext-community is getting reset to 0. [PR1371448](#)
- Host destined packets with filter log action might not reach to the Routing Engine if log or syslog is enabled. [PR1379718](#)
- The IRB transit traffic might not be counted for EVPN/VXLAN traffic. [PR1383680](#)
- EVPN VXLAN non-collapsed: AUTONEG errors and flush operation failed error are seen after the power cycle of the device. [PR1394866](#)
- On QFX5110, the firewall filter applied on VxLAN mapped VLAN is not supported in EVPN-VxLAN scenario. [PR1398237](#)
- ERACL firewall group will operate in double wide mode for QFX5110 in Junos OS Release 19.1R1. [PR1408670](#)
- ICMPv6 RA packets generated by Routing Engine might be dropped on the backup member of Virtual Chassis if **igmp-snooping** is configured. [PR1413543](#)
- The dcpfe might crash when any interface flap. [PR1415297](#)
- The QFX and EX switch may not install all IRB MAC addresses in the initialization [PR1416025](#)
- The same traffic flow might be forwarded to different ECMP next hops on QFX5K. platforms [PR1422324](#)
- The traffic loss might start after deleting IRB logical interface. [PR1424284](#)
- The rpd might generate a core file because of the improper handling of graceful restart stale routes. [PR1427987](#)
- BGP statement **multipath multiple-as** does not work in specific scenario. [PR1430899](#)
- BGP session might go into down status once the traffic flow starts. [PR1431259](#)
- fxpc core file is generated once during reboot due to Bad Chip ID. [PR1432023](#)
- Ping fails over type-5 tunnel on IRB interfaces under EVPN-VXLAN scenario. [PR1433918](#)
- The IPv4 fragmented packets might be broken if PTP transparent clock is configured. [PR1437943](#)
- The bandwidth value of the DDoS-protection might cause the packets loss after the device reboot. [PR1440847](#)

- Traffic might be dropped after the QinQ enabled interface is flapped or a change is made to the **vlan-id-list**. [PR1441402](#)
- On QFX5210, firewall filter DSCP action modifier does not work when firewall filter is mapped to IRB. [PR1441444](#)
- The rpd process might crash in inter-AS option B L3VPN scenario if CNHs is used. [PR1442291](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- PIM (S,G) joins might 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 instance might not work on 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 or 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](#)

User Interface and Configuration

- QFX5100 devices are unable to commit baseline configuration after zeroize. [PR1426341](#)

Resolved Issues: 19.1R1

EVPN

- A few minutes of traffic loss might be observed during recovery from link failure. [PR1396597](#)
- VNI is not updated on default route 0.0.0.0/0 advertised by EVPN type 5 prefix when the local is configuration changed. [PR1396915](#)
- EVPN routes might show **Route Label: 0** in addition to the real label. [PR1405695](#)
- The rpd might crash after NSR switchover. [PR1408749](#)

Interfaces and Chassis

- Constant dcpfe process crash might be seen if using an unsupported GRE interface configuration. [PR1369757](#)

Layer 2 Ethernet Services

- After GRES switchover, LACP will be down on the peer device and never recover automatically. [PR1395943](#)

Layer 2 Features

- The IPv6 NS/NA packets coming from the remote VTEP are not getting forwarded to the local host. [PR1387519](#)
- The dcpfe process might crash after VXLAN overlay ping. [PR1388103](#)
- With IGMP snooping enabled on the leaf switches, multicast traffic is forwarded to VLAN/VNI, which doesn't have an active receiver. [PR1388888](#)
- RTG MAC refresh packets will be sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. [PR1389695](#)
- Packets destined to 01:00:0c:cc:cc:cc are not forwarded on QFX10000. [PR1389829](#)
- EVPN-VXLAN: Dcpfe is restarted at the _bcm_field_td_counter_last_hw_val_update routine after upgrading spine with latest image. [PR1398251](#)
- On QFX5000, dcpfe process crash might be observed during restart of Packet Forwarding Engine or system with scaled EVPN/VXLAN configuration. [PR1403305](#)
- The IPv6 NS/NA packets received over VTEP from an ESI host are incorrectly flooded back to the host. [PR1405820](#)
- With Junos OS releases before 19.1R1, on devices with cut-through configuration enabled, after reboot of the device, cut-through mode will be disabled on the channelized interfaces. [PR1407706](#)
- With arp-suppression/proxy-arp feature, QFX5100 or QFX5110 might not forward IPv6 Router Solicitations or Advertisements. [PR1414496](#)

MPLS

- LSP "statistics" and "auto-bandwidth" functionality might not take effect with single-hop LSPs. [PR1390445](#)

Network Management and Monitoring

- Log files might not get compressed during the upgrade. [PR1414303](#)

Platform and Infrastructure

- The 1-Gigabit Ethernet copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms. [PR1286709](#)
- Optics BiDi: FEC incorrectly displayed on QFX5110 and QFX10002-36Q. [PR1360948](#)

- SFP-T might not work on QFX5100/QFX5110 devices. [PR1366218](#)
- The first 2 characters out of 14 of AS7816-64 serial number is truncated. [PR1371126](#)
- For the Junos OS 18.1R1 or later, USB image installation on QFX5210-64C, an AMI BIOS upgrade is required. [PR1371199](#)
- Packet Forwarding Engine is in a bad state after performing optics insertion or removal on a port. [PR1372041](#)
- The IPv6 routed packet might be transmitted through an interface whose VRRP state is in non-master. [PR1372163](#)
- QFX5110 ethernet-switching flood group shows incorrect information. [PR1374436](#)
- Packet Forwarding Engine wedge might be observed if there are interfaces going to the down state. [PR1376366](#)
- EM policy update is needed on QFX5210-64C. [PR1380077](#)
- The overlay ECMP might not work as expected on QFX5110 in an EVPN-VXLAN environment. [PR1380084](#)
- IPv6 ping might fail for spine node in EVPN scenario. [PR1380590](#)
- Traffic black hole is caused by FPC offline in MC-LAG scenario. [PR1381446](#)
- The QFX-QSFP-40G-SR4 transceiver might not be recognized after upgrading Junos OS on QFX5100e. [PR1381545](#)
- LACP gets stuck in detached/attached state when the interface is configured with native VLAN ID and VXLAN VLAN. [PR1382209](#)
- QFX10008 continuously shows **RPD_KRT_Q_RETRIES: list nexthop ADD: No such file or directory**. [PR1383426](#)
- The DMA failure errors might be seen when the cache is flushed or the cache is full. [PR1383608](#)
- DHCP packets might be dropped on a Junos fusion Data Center scenario (QFX10000 line of devices). [PR1383623](#)
- Last reboot reason is not correct if device is rebooted because of power cycle. [PR1383693](#)
- The Virtual Chassis could not come up after upgrading to QFX5E platforms (TVP-based platforms for QFX5100 or QFX5200 switches). [PR1383876](#)
- A “force host” upgrade is required for QFX5110-48S-4C in Junos OS Release 18.4 if the PTP over IPV6 G.8275.2 feature is configured. [PR1384073](#)
- Tuning issue exist for SFPP-10G-DT-ZRC2 and SFPP-10G-CT50-ZR. [PR1384524](#)
- QFX5120: Occasionally two of the channelized 25-Gigabit Ethernet ports using 4x25G breakout cable will not come up after Junos OS reboot. [PR1384898](#)

- The IPv6 packet might not be routed when the IPv6 packet is encapsulated over IPv4 GRE tunnel on QFX10000. [PR1385723](#)
- The spine EVPN routes might be stuck in a hidden state with next hop as unusable after FPC is offline in the spine. [PR1386147](#)
- DDoS statistics and logging are not working for internal queues such as Q42 and Q4. [PR1387508](#)
- Traffic drop might be seen on QFX10000 platform with EVPN-VXLAN configured. [PR1387593](#)
- QFX5100/QFX5110/QFX5200/QFX5210 Virtual Chassis could not be formed normally. [PR1387730](#)
- Certain log messages might be observed on QFX Series platforms. [PR1388479](#)
- MAC learning might stop working on some LAG interfaces. [PR1389411](#)
- FPC might crash on QFX5100 and EX4600 platforms in a large-scale scenario. [PR1389872](#)
- The vmcore might be seen when routing changes are made on the peer spine in an EVPN VXLAN scenario. [PR1390573](#)
- An incorrect error message might be seen when J-Flow sensors are configured with reporting rate less than 30 seconds. [PR1390740](#)
- Smid core file is seen during sanity script execution on QFX5100 and EX4300. [PR1391909](#)
- Sdk-vmmd might consistently write to the memory. [PR1393044](#)
- 10-Gigabit Ethernet copper link flapping might happen during TISSU operation of QFX5100-48T switches. [PR1393628](#)
- IPV6 next-hop programming issue might be observed on QFX10000/PTX1000/PTX10000 devices. [PR1393937](#)
- L2ALD core file is seen when l2-learning traceoptions were enabled. [PR1394380](#)
- DRAM and buffer utilization fields are not correct for QFX10000 platforms. [PR1394978](#)
- PTP over Ethernet traffic could be dropped if IGMP and PTP TC are configured together. [PR1395186](#)
- DOT1XD core found at `pnac_bd_create_pnac_bdm_handler_knl_async_receive_and_process`. [PR1395384](#)
- Unable to install licenses automatically on QFX Series platforms. [PR1395534](#)
- `BRCM_NH-,brcm_bcm_mpls_tunnel_initiator_clear(),226:bcm_mpls_tunnel_initiator_get failed intf = 4` failure error logs might be seen in syslog. [PR1396014](#)
- If GRES/NSR is enabled on a QFX5100 (single Routing Engine), DHCP subscribers are failing to bind. [PR1396470](#)
- QFX10002-60C: FPC might not be detected after the ukern crashes. [PR1396507](#)
- High jsd or na-grpcd CPU usage might be seen even JET or JTI is not used. [PR1398398](#)
- The DHCPv6 relay packets are dropped when both the UDP source and destination ports are 547. [PR1399067](#)

- CPU hog might be observed on QFX10000 platform. [PR1399369](#)
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. [PR1399683](#)
- ARP refresh functionality might fail in an EVPN scenario. [PR1399873](#)
- PEM I2C failure alarm might be showed incorrectly as failed. [PR1400380](#)
- MAC-limit with persistent MAC is not working after reboot. [PR1400507](#)
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- The authd might crash when issuing **show network-access requests pending** command during the authd restarting. [PR1401249](#)
- File permissions are changed for **/var/db/scripts** files after reboot. [PR1402852](#)
- The VRRP VIP might not work when it is configured on the LAG interface. [PR1404822](#)
- ARP/ND will not be resolved in case of native VLAN ID configured for LAG access interface. [PR1404895](#)
- Commit warning occurs on QFX5100. [PR1405138](#)
- VXLAN transit traffic over tagged underlay L3 Interface gets dropped due to hardware limitation. [PR1406282](#)
- EVPN-VXLAN: QFX10002: With arp-suppression present (enabled by default), packets egressing the QFX Series switch are tagged with 4095 VLAN when using SP-style configurations on the ports. [PR1407059](#)
- DHCP discover packets are getting dropped over VXLAN tunnel on a pure L2 VLAN when DHCP relay is enabled for other VLANs. [PR1408161](#)
- The FPC might crash and could not come up if interface-num or next hop is set to maximum value under vxlan-routing on QFX Series platforms. [PR1409949](#)

Routing Protocols

- QFX5120: The command output **show pfe route summary hw** will show different scale values for the IPv4 and IPv6 LPM routes rather than the supported scale. [PR1366579](#)
- Host-destined packets with filter log action might reach the Routing Engine. [PR1379718](#)
- MMU errors on QFX5200 running Junos OS Release 15.1X53-D234.2. [PR1381790](#)
- BUM packets might get looped if EVPN multihoming interface flaps. [PR1387063](#)
- The next hop in hardware for existing ECMP route might not be updated when **ecmp-resilient-hash** is configured. [PR1387713](#)
- CLI **show evpn igmp-snooping database extensive** output needs to be modified according to the SMET functionality. [PR1391406](#)
- On QFX5110 and QFX5200 switches, the non-collapsed EVPN-VXLAN dcfpe core file is seen at **brcm_pkt_tx_flush, l2alm_mac_ip_timer_handle_expiry_event_loc** after a random event. [PR1397205](#)

SEE ALSO

New and Changed Features	 252
Changes in Behavior and Syntax	 267
Known Behavior	 272
Known Issues	 276
Documentation Updates	 303
Migration, Upgrade, and Downgrade Instructions	 303

Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for the QFX Series.

SEE ALSO

New and Changed Features	 252
Changes in Behavior and Syntax	 267
Known Behavior	 272
Known Issues	 276
Resolved Issues	 282
Migration, Upgrade, and Downgrade Instructions	 303

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches](#) | [304](#)
- [Installing the Software on QFX10002-60C Switches](#) | [306](#)
- [Installing the Software on QFX10002 Switches](#) | [306](#)
- [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](#) | [307](#)
- [Installing the Software on QFX10008 and QFX10016 Switches](#) | [309](#)
- [Performing a Unified ISSU](#) | [313](#)

- Preparing the Switch for Software Installation | 314
- Upgrading the Software Using Unified ISSU | 314
- Upgrade and Downgrade Support Policy for Junos OS Releases | 316

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 **18.3** 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 18.3 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-x86-64-18.3-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 18.3 **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-x86-64-18.1R1.9.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-x86-64-18.1R1.9.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-x86-64-18.3R1.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-x86-64-18.3R1.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-x86-64-18.3R1.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-x86-64-18.3R1.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 314](#)
- [Upgrading the Software Using Unified ISSU on page 314](#)

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-x86-64-18.3R1.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-x86-64-18.3R1.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

[New and Changed Features | 252](#)

[Changes in Behavior and Syntax | 267](#)

[Known Behavior | 272](#)

[Known Issues | 276](#)

[Resolved Issues | 282](#)

[Documentation Updates | 303](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 318](#)
- [What's Changed | 329](#)
- [Known Limitations | 336](#)
- [Open Issues | 339](#)
- [Resolved Issues | 341](#)
- [Documentation Updates | 360](#)
- [Migration, Upgrade, and Downgrade Instructions | 361](#)

These release notes accompany Junos OS Release 19.1R3 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

- [What's New in Release 19.1R3 | 318](#)
- [What's New in Release 19.1R2 | 319](#)
- [What's New in Release 19.1R1 | 319](#)

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

What's New in Release 19.1R3

There are no new features in Junos OS Release 19.1R3 for the SRX Series devices.

What's New in Release 19.1R2

Chassis Clustering

- **Dedicated fabric ports support (SRX4600)**—Starting in Junos OS Release 19.1R2, you can use the built-in dedicated fabric ports as fabric link ports in chassis cluster mode.

[See [Understanding Chassis Cluster Slot Numbering and Physical Port and Logical Interface Naming](#), [SRX Series Chassis Cluster Configuration Overview](#), and [Chassis Cluster Control Plane Interfaces](#).]

What's New in Release 19.1R1

Application Security

- **CLI enhancements to support J-Web in application identification (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, the **show services application-identification** command is enhanced to display application and application group details in J-Web.

The **show services application-identification application** command includes the new **risk** option and the **show services application-identification entries** command is enhanced with the new **category-list** and **subcategory-list** options. These options support and improve the J-Web search mechanism.

[See [show services application-identification application](#).]

- **Support for user source identity in APBR policies (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, you can configure advanced policy-based routing (APBR) policies by defining the user source identity as one of the match criteria along with source addresses, destination addresses, and applications.

If you specify source identity as a match criteria in a policy, then the user and role information are retrieved before policy lookup can proceed. After a successful match, the APBR profile configured with the APBR policy is used for applying the configured rule.

[See [Advanced Policy-Based Routing](#).]

- **Application quality of experience (AppQoE) support in high availability (HA) mode (SRX4100, SRX4200)**—Starting in Junos OS Release 19.1R1, the SRX4100 and SRX4200 support application quality of experience (AppQoE) when these devices operate in chassis cluster mode.

You can configure these SRX Series devices to operate both in active/active and in active/passive modes and deploy the device as spoke device in SD-WAN deployments.

[See [Application Quality of Experience](#).]

- **Application services bypass in APBR (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, you can bypass the application services on a session using advanced policy-based routing (APBR) profile rule. When the APBR profile rule is matched and re-routing is done, you can specify that the traffic matching the APBR profile rule can be bypassed from the application services that are configured on the SRX Series devices.

You can use the APBR profile rule to bypass application services such as security policy, application quality of service (AppQoS), Juniper Sky ATP, IDP, Security Intelligence (SecIntel), and UTM using the APBR rule.

See [\[Advanced Policy-Based Routing.\]](#)

- **AppQoE scaling support (SRX4100 and SRX4200)**—Starting in Junos OS Release 19.1R1, Application quality of experience (AppQoE) enforces the configuration limit for overlay paths, metric profiles, probe parameters, and SLA rules per profile when you configure application-specific SLA rules and associate the SLA rules to an APBR profile. If you configure more parameters than the allowed limit, an error message is displayed after you commit the configuration.

[See [Application Quality of Experience.](#)]

Authentication Access Control

- **Monitoring DHCP session logs (SRX Series)**—Starting in Junos OS Release 19.1R1, you can monitor the Dynamic Host Configuration Protocol (DHCP) session events. Using the session logs generated by the `jdhcp` process, you can observe the session (subscribe) creation, session deletion, and renew events details. You can configure the DHCP session logs by using the `log session` and `log session dhcpv6` options at the `[edit system processes dhcp-service]` hierarchy level for IPv4 and IPv6 addresses, respectively. You can use the session logs for monitoring and troubleshooting purposes.

[See [log.](#)]

Intrusion Detection and Prevention (IDP)

- **Covert channels identification and mitigation for IPv6 extension headers (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, Intrusion Detection and Prevention (IDP) supports the identification and mitigation of covert channels for IPv6 extension headers.

Covert channel is a type of attack in which information is transferred through existing channels that should not be allowed to communicate by the configured security policy. Thus, this kind of communication violates the existing security system.

The IPv6 covert channel anomalies are part of the IDP signature database package. You can configure the anomalies by using the `predefined-attacks` statement under the `idp-policies` hierarchy level.

[See [Attack Objects and Object Groups for IDP Polices.](#)]

- **Deprecation of signatures in IDP (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, certain signatures are marked as deprecated or excluded from the Intrusion Prevention System (IPS).

For dynamic attack groups, two filters—`Excluded` and `no-excluded`—are introduced at the `[edit security idp dynamic-attack-group dynamic-attack-group-name filters]` hierarchy level to check the signatures which are part of the database updates.

The `show security idp attack deprecated-list` and `show security idp policy deprecated attacks` commands are introduced to display the list of deprecated attacks in the signature updates.

[See [IDP Signature Database Overview.](#)]

- **Support for Hyperscan extended parameters in IDP signature-based attacks (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, you can configure signature-based attacks by using Hyperscan extended parameters. By setting optimal values for the Hyperscan extended parameters, you can enhance the attack pattern matching process significantly.

To configure the extended parameters, include the **optional-parameters** option at the **[edit security idp custom-attack *attack-name* attack-type signature]** hierarchy level. You can configure the following parameters under the **optional-parameters** option:

- **min-offset**
- **max-offset**
- **min-length**

[See [Understanding IDP Signature-Based Attacks.](#)]

J-Web

- **Threat Assessment report supports new charts (SRX Series)**—Starting in Junos OS Release 19.1R1, the Threat Assessment report supports the following charts:
 - Top Web Categories for Security High—Displays only high severities and the top 10 Web categories.
 - Top Web Categories—Displays the top 10 Web categories.
 - Top Users Accessing Risky Websites—Displays the top 10 values.
 - Top URL Categories for Security Risk (High and Medium)—Displays both high and medium severities and the top 10 values.
 - Top URL Categories for Productivity Loss—Displays the top 10 values.
 - Top URL Categories for Legal Liability—Displays the top 10 values.

[See [Reports.](#)]

- **IPsec VPN security services support new authentication algorithm and Diffie-Hellman (DH) group values (SRX Series)**—Starting in Junos OS Release 19.1R1, IPsec VPN security services support and display the following new values:
 - IKE (Phase I)—SHA 512-bit authentication algorithm, DH Group 15, 16, and 21
 - IKE (Phase II)—HMAC-SHA-512 authentication algorithm, HMAC-SHA-384 authentication algorithm, DH Group 15, 16, and 21

NOTE: The new authentication algorithms and DH groups support the SRX5000 line of devices with SPC3 upon installation of junos-ike package only. Click **Install** from **Configure>Security Services>IPsec VPN>Global Settings** to install the package.

[See [VPN Global Settings Configuration Page Options](#), [IKE \(Phase I\) Configuration Page Options](#), and [IKE \(Phase II\) Configuration Page Options](#).]

- **Certificate management supports new bit length for the Elliptic Curve Digital Signature Algorithm (ECDSA) key (SRX Series)**—Starting in Junos OS Release 19.1R1, when you create a certificate, Certificate management supports the bit length of the 521 ECDSA key.

[See [Managing Certificates](#).]

- **User management supports new password setting range (SRX Series)**—Starting in Junos OS Release 19.1R1, the user management configuration supports the password settings range as follows:
 - Minimum Reuse: 1-20 old passwords, but these must not be the same as the new password you set.
 - Maximum Lifetime: 30-365 days
 - Minimum Lifetime: 1-30 days

NOTE: Using J-Web, you cannot configure the minimum number of characters required for a new password.

[See [User Management Configuration Page Options](#).]

- **In J-Web, device basic settings can be configured on a single page (SRX Series)**—Starting in Junos OS Release 19.1R1, you can configure the following basic settings for a device on a single page in J-Web:
 - System Identity Details
 - Date & Time
 - Management Access Configuration

NOTE: If the SRX Series device does not have a dedicated management port (fxp0), then **Loopback Address** and **Subnet** are the only options available for configuring management access. For SRX Series devices with the fxp0 port, IPv4 configuration is supported for configuring management access.

- Security Logging—Supports only stream mode.
- SNMP

[See [System Identity Configuration Page Options](#).]

- **Support for monitoring logical system users and tenant users (SRX Series)**—Starting in Junos OS Release 19.1R1, the **Users** option under the **Monitor** tab is available for both logical system users and tenant users.

[See [Monitoring Users](#).]

- **Support for events monitoring configuration for logical system users and tenant users (SRX Series)**—Starting in Junos OS Release 19.1R1, the following events monitoring configurations are supported for logical system users and tenant users:
 - Firewall events are supported for both logical system users and tenant users.
 - All events, Web filtering, content filtering, antispam, antivirus, and IPS events are supported for logical system users.

[See [Monitoring Firewall Events](#) and [Monitoring All Events](#).]

- **Supported reports for logical system users and tenant users (SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 19.1R1:
 - Threat assessment, application and user, talkers, firewall, screen, and source zone reports are supported for logical system users and tenant users.
 - IPS, URL, viruses, antispam, Web applications, roles, botnet, malware, blocked application, and permitted application reports are supported only for logical system users.

[See [Reports](#).]

- **Report generation status when the context is switched (SRX Series)**—Starting in Junos OS Release 19.1R1, you can choose to stop generating a report to switch the context or continue generating the report without switching the context using the confirmation message.

[See [Configuring Multi Tenancy Logical Systems](#).]

- **Support for traffic logging (SRX Series)**—Starting in Junos OS Release 19.1R1, traffic logging is enabled as part of the security logging configuration for logical system users and tenant users. When you enable traffic logging, the existing event mode configuration (if any) is deleted.

[See [Security Logging Configuration Page Options](#).]

- **Firewall security policy rules support source identity for local authentication users (SRX Series)**—Starting in Junos OS Release 19.1R1, a list of local authentication users is available in source identity for logical system users and tenant users.

[See [Configuring Firewall Policy Rules](#).]

- **Local authentication monitoring for logical system users and tenant users (SRX Series)**—Starting in Junos OS Release 19.1R1, the local authentication option (Monitor > Authentication > Local Auth) is enabled for logical system and tenant users.

The Clear All option is not available for either logical system users or tenant users to clear the authentication information.

[See [Monitoring Local Authentication](#).]

- **Autocompletion of logical system names or tenant names (SRX Series)**—Starting in Junos OS Release 19.1R1, when you type the partial name of the logical system name or tenant name, the user interface automatically completes the name.

[See [Interconnecting Interface Ports Configuration Page Options](#).]

- **Multitenancy support is provided for logical system users and tenant users (SRX Series)**—Starting in Junos OS Release 19.1R1, you can have the following maximum number of logical system users and tenant users for multitenancy:

Table 1: Maximum Number of Logical System Users and Tenant Users for Multitenancy

SRX Series	Number of Logical System Users	Number of Tenant Users
SRX5000 line of devices with SPC2	32	100
SRX5000 line of devices with SPC3	32	500
SRX5000 line of devices with mixed SPC2 and SPC3	32	100
SRX4600	32	300
SRX4200	32	200
SRX4100	32	200
SRX1500	32	50

[See [Configuring Multi Tenancy Logical Systems](#) and [Configuring Multi Tenancy Tenants](#).]

- **User configurations available on a single page (SRX Series)**—Starting in Junos OS Release 19.1R1, the following user configurations are available on a single page:
 - User Management
 - Firewall Authentication
 - Access Profiles
 - UAC Settings

[See [User Management Configuration Page Options](#).]

- **Address Pool available as a separate configuration page (SRX Series)**—Starting in Junos OS Release 19.1R1, you can access Address Pool as a separate configuration page from Configure > Security Objects.

[See [Address Pools Configuration Page Options](#).]

- **App Tracking available under Security Objects (SRX Series)**—Starting in Junos OS Release 19.1R1, you can configure application tracking from Configure > Security Objects > App Tracking.

[See [Application Tracking Configuration Page Options](#).]

- **Changes on the Monitoring Events page (SRX Series)**—Starting in Junos OS Release 19.1R1, the Summary View is replaced with the Chart View, and the Detailed View is replaced with the Grid View. These changes are applicable to all the configuration pages (except the System page) under Monitor > Events.

[See [Monitoring All Events](#).]

- **IKE (Phase II) supports new values for the Establish tunnels option (SRX Series)**—Starting in Junos OS Release 19.1R1, the Establish tunnels option supports the **responder-only** and **responder-only-no-rekey** values.

NOTE:

- The **responder-only** option is supported on the SRX5000 line of devices with an SPC3 card only if the junos-ike-package is installed. To install this package from J-Web, navigate to **Configure>Security Services>IPsec VPN>Global Settings**, and click **Install**.
- When you configure the **responder-only** value on multiple VPN objects with a single gateway configuration, ensure that all the VPN objects are configured with this mode.
- The **responder-only** option is supported only on a site-to-site VPN. This option is not supported on AutoVPN.

[See [VPN AutoKey Configuration Page Options](#).]

- **New risk values in application signature (SRX Series)**—Starting in Junos OS Release 19.1R1, when the custom application creates an application signature, it supports the following application signature risk levels:
 - Low
 - Moderate
 - Unsafe
 - High
 - Critical

[See [Application Signature Configuration Page Options](#).]

- **Support for PowerMode IPsec (SRX4100, SRX4200, SRX4600, SRX5000 line with SPC3 card, and vSRX)**—Starting in Junos OS Release 19.1R1, you can enable or disable **PowerMode IPsec (PMI)** in the IPsec VPN Global Settings.

NOTE:

- After the PMI configuration is committed, the Packet Forwarding Engine service restarts automatically. The Packet Forwarding Engine service will not be explicitly restarted.
- You can use the J-Web user interface to enable or disable PMI depending on the configuration required for each of the devices.

[See [VPN Global Settings Configuration Page Options](#).]

Logical Systems and Tenant Systems

- **SSL proxy support for logical systems (SRX Series)**—Starting in Junos OS Release 19.1R1, SRX Series devices that have logical systems configured support the Secure Sockets Layer (SSL) proxy functionality. The logical-system users can configure and view the SSL profiles specific to their own logical systems by using the root certificate. The logical-system users can configure SSL profiles for proxy termination and initiation on logical systems and can also configure the certificate authority (CA), load a CA profile group, and apply an SSL proxy profile to a security policy for logical systems.

[See [SSL Forward Proxy Overview](#).]

- Starting in Junos OS Release 19.1R1, the following features that are supported on the logical systems are now extended to tenant systems:
 - **Logging support for tenant systems (SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 19.1R1, on-box reporting configurations are supported for each tenant system, and logs are handled based on these configurations. Use the **set security log report** and **set security log mode stream** commands to enable the on-box reporting. The on-box reporting feature with stream mode is also supported on tenant systems.

[See [Security Log for Tenant Systems](#).]

- **User firewall enhanced support for tenant systems (SRX Series)**—Starting in Junos OS Release 19.1R1, support for user firewall authentication is enhanced using a shared model. In this model, tenant systems share user firewall configuration and authentication entries with the master logical system. The tenant system shares the authentication data collected from the local authentication, Active Directory authentication, firewall authentication, Juniper Identity Management Service (JIMS), and ClearPass authentication with the master logical system.

[See [Firewall Authentication for Tenant Systems](#).]

Routing Policy and Firewall Filters

- **Optional application configuration in a unified policy (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R1, configuring the **application** statement at the **[edit security policies from-zone zone-name to-zone zone-name policy policy-name match]** hierarchy level is optional if the **dynamic-application** statement is configured at the same hierarchy level.

In releases before Junos OS Release 19.1R1, it is mandatory to configure the **application** statement even if the **dynamic-application** statement is configured.

[See [application \(Security Policies\)](#) and [dynamic-application \(Security Policies\)](#).]

Routing Protocols

- **Support for BGP graceful shutdown (SRX Series)**— Starting in Junos OS Release 19.1R1, graceful traffic migration from one BGP next hop to another is supported, without traffic interruption. Also, BGP administrative shutdown communication can be sent to the BGP peer.

You can configure both **graceful-shutdown** and **shutdown** statements at the **[edit protocols bgp]**, **[edit protocols bgp group group-name]**, and **[edit protocols bgp group group-name neighbor address]** hierarchy levels.

NOTE: Graceful shutdown is disabled by default.

[See: [graceful-shutdown \(Protocols BGP\)](#), [shutdown \(Protocols BGP\)](#).]

Security

- **Juniper Entropy Beacon (SRX Series)**—Starting in Junos OS Release 19.1R1, Juniper Entropy Beacon (JEB) allows authorized devices to request entropy packages from a SRX345 Services Gateway configured as a JEB server. Entropy is a crucial component of all cryptographic security systems because it is used to generate symmetric and asymmetric cryptographic keys. Low entropy leads to predictable keys, which can compromise the security of a system. JEB provides high quality entropy from a trusted source to entropy starved clients securely over the network.

[See [Juniper Entropy Beacon Overview](#)]

Unified Threat Management (UTM)

- **SRX TAP mode support for UTM features (SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 19.1R1, the Unified Threat Management (UTM) module supports TAP (Terminal Access Point) mode. When you configure SRX Series device to operate in TAP mode, the device generates and displays security log information such as threats detected, application usage, and user details. When configured to operate in TAP mode, the SRX Series device receives packets only from the configured TAP interface.

[See [Enhanced Web Filtering](#).]

VPN

- **PowerMode IPsec with SPC3 (SRX5400, SRX5600, and SRX5800)**—Starting in Release 19.1R1, Junos OS on SRX Series devices with SPC3 supports a new mode of IPsec operation called PowerMode IPsec (PMI). PMI uses a small software block inside the Packet Forwarding Engine that bypasses flow processing and utilizes the Intel Advanced Encryption Standard New Instructions (AES-NI) for optimized performance of IPsec processing.

You can enable PMI processing by using the **power-mode-ipsec** statement at the [edit security flow hierarchy level.

With PMI configured, the device supports the following features:

- Internet Key Exchange (IKE) functionality
- AutoVPN with traffic selectors
- High availability
- IPv6
- Stateful firewall
- st0 interface
- Traffic selectors

[See [Understanding PowerMode IPsec.](#)]

- **Cryptographic algorithm support for IPsec and IKE on SRX5K-SPC3 card (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.1R1, SRX5000 line of devices with SRX5K-SPC3 card support cryptographic algorithms to protect classified information.

The following algorithms are supported for IPsec:

- Diffie-Hellman Group 15
- Diffie-Hellman Group 16
- Diffie-Hellman Group 21
- HMAC-SHA-512
- HMAC-SHA-384

The following algorithms are supported for IKE:

- Diffie-Hellman Group 15
- Diffie-Hellman Group 16
- Diffie-Hellman Group 21
- SHA-512
- ECDSA-521 for X509 signatures

[See [IPsec VPN Overview](#) and [Understanding Certificates and PKI](#).]

- **Support for CoS classifier and rewrite functions in PMI on SPC3 (SRX Series)**— Starting in Junos OS Release 19.1R1, class of service (CoS) supports the configuration of behavior aggregate (BA) classifier, multifield (MF) classifier, and rewrite-rule functions in PowerMode IPsec (PMI) on SPC3 cards.

[See [Improving IPsec Performance with PowerMode IPsec](#).]

- **Support for IKE responder-only mode (SRX Series)**—Starting in Junos OS Release 19.1R1, two new options for the establishment of IPsec tunnels are introduced. The **responder-only** and **responder-only-no-rekey** options are added to the **establish-tunnels** statement under the **[edit security ipsec vpn vpn-name]** hierarchy level.

When you use these options, the VPN tunnel is established from the remote peer. In the case of the **responder-only** option, an established tunnel rekeys both Internet Key Exchange (IKE) and IPsec, based on the configured lifetime values. When you use the **responder-only-no-rekey** option, an established tunnel does not initiate rekeying from the device but relies on the remote peer to initiate rekeying.

[See [IPsec VPN Overview](#).]

SEE ALSO

Changes in Behavior and Syntax 329
Known Behavior 336
Known Issues 339
Resolved Issues 341
Documentation Updates 360
Migration, Upgrade, and Downgrade Instructions 361

What's Changed

IN THIS SECTION

- [Changes in Behavior and Syntax: Release 19.1R3 | 330](#)
- [Changes in Behavior and Syntax: Release 19.1R2 | 333](#)
- [Changes in Behavior and Syntax: Release 19.1R1 | 335](#)

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

Changes in Behavior and Syntax: Release 19.1R3

Application Layer Gateways (ALG)

- **Disable the do not fragment flag from packet IP header (SRX Series and vSRX)**—Starting in Junos OS Release 19.1R3, we've introduced the **clear-dont-frag-bit** option at the `[edit security alg alg-manager]` hierarchy level to disable the do not fragment flag from the packet IP header, which allows the packet to be split after NAT is performed.

In Junos OS releases earlier than Release 19.1R3, when the ALG performs payload-NAT, sometimes the size of the packet becomes bigger than the outgoing interface maximum transmission unit (MTU). If the packet IP header has the do not fragment flag, this packet cannot be sent out.

[See [alg-manager](#).]

Application Security

- Starting in Junos OS 19.1R3 release, you can configure maximum memory limit for the deep packet inspection (DPI) by using the following configuration statement:

```
user@host# set services application-identification max-memory memory-value
```

You can set 1 through 200000 MB as memory value.

Once the JDPI memory consumption reaches to 90% of the configured value, then DPI stops processing new sessions.

- In Junos OS Release 19.1R3, 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 2 on page 331 provides the range and default value for configuring the byte limit and the packet limit for TCP and UDP sessions.

**Table 2: 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"> <li>• 6000</li> <li>• For Junos OS Release 15.1X49-D200, the default value is 10000.</li> </ul> |
|         | Packet limit | 0 through 4294967295 | Zero                                                                                                                               |
| UDP     | Byte limit   | 0 through 4294967295 | Zero                                                                                                                               |
|         | Packet limit | 0 through 4294967295 | <ul style="list-style-type: none"> <li>• 10</li> <li>• For Junos OS Release 15.1X49-D200, the default value is 20.</li> </ul>      |

The byte limit excludes the IP header and the TCP/UDP header lengths.

If you set the 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 19.3R1, 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** value, 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](#)]

### *Authentication and Access Control*

- **Enhanced user firewall support**—In Junos OS Release 19.1R3, 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](#).]

### *Ethernet Switching*

- **LLDP support on redundant Ethernet interfaces (SRX Series)**—Starting in Junos OS Release 19.1R3, you can configure the Link Layer Discovery Protocol (LLDP) on redundant Ethernet (reth) interfaces. Use the **set protocol lldp interface *reth-interface*** command to configure LLDP on a reth interface.

[See [Configuring LLDP](#) and [Ethernet Ports Switching Overview for Security Devices](#).]

### *Juniper Sky ATP*

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 19.1R3, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Sky Advanced Threat Prevention (ATP).

### *Unified Threat Management (UTM)*

- **Increase in the UTM scale number (SRX1500, SRX4100, SRX4200, SRX4600, SRX4800, SRX5400, SRX5600, and SRX5800)**—Starting with Junos OS Release 19.1R3, on SRX Series devices, UTM policies, profiles, MIME patterns, filename extensions, protocol commands, and custom messages are increased up to 1500. Custom URL patterns and custom URL categories are increased up to 3000.

[See [Unified Threat Management overview](#).]

### *VPNs*

- **Public key infrastructure warning message (SRX Series)**—Junos OS Release 19.1R3 displays the warning message **ECDSA Keypair not supported with SCEP for cert\_id *certificate id*** when you try to enroll a local certificate using an Elliptic Curve Digital Signature Algorithm (ECDSA) key with Simple Certificate Enrollment Protocol (SCEP). The message appears because the ECDSA key is not supported with SCEP.

In Junos OS Release before Junos OS Release 19.1R3, the warning message is not displayed.

[See [Example: Enrolling a Local Certificate Online Using SCEP](#).]

## **Changes in Behavior and Syntax: Release 19.1R2**

### *Authentication and Access Control*

- **SSH protocol version v1 option deprecated from CLI (SRX Series)**—Starting in Junos OS Release 19.1R2, we've removed the nonsecure SSH protocol version 1 (**v1**) option from the **[edit system services ssh protocol-version]** hierarchy level. You can use the SSH protocol version 2 (**v2**) as the default option to remotely manage systems and applications. With the **v1** option deprecated, Junos OS is compatible with OpenSSH 7.4 and later versions.

Junos OS releases earlier than Release 19.1R2 continue to support the **v1** option to remotely manage systems and applications.

[See [protocol-version](#).]

### **Network Management and Monitoring**

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (SRX Series)**—Starting in Junos OS Release 19.1R2, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.
- **Default system log messages (SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550M)**—Starting in Junos OS Release 19.1R2, the default mode for system log messages is changed from event mode to stream mode.

[See [Understanding System Logging for Security Devices](#) and [mode \(Security Log\)](#).]

## Changes in Behavior and Syntax: Release 19.1R1

### Flow-Based and Packet-Based Processing

- **Change in the maximum number of sessions permitted (SRX340)**—Starting in Junos OS Release 19.1R1, the maximum number of sessions permitted on SRX340 devices increases. [Table 3 on page 335](#) shows the maximum number of sessions permitted on SRX340 devices.

**Table 3: Maximum Number of Sessions Permitted on SRX340 Devices**

| Junos OS Release                | Device                           | Maximum Number of Sessions |
|---------------------------------|----------------------------------|----------------------------|
| Junos OS Release 19.1R1 onward  | SRX340                           | 375000                     |
|                                 | SRX340 configured with a license | 256000                     |
| Junos OS Releases before 19.1R1 | SRX340                           | 256000                     |
|                                 | SRX340 configured with a license | 128000                     |

See [Features Requiring a License on SRX340 Devices](#) for more information about licenses for SRX340 Series Devices.

[See [show security flow session.](#)]

### Platform and Infrastructure

- **Chassis cluster with SPC card (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.1R1, when an SPC is acting as the central point as well as hosting the single control link port, this creates a single point of failure. If the SPC goes down on the primary node, the node is automatically rebooted to avoid a split-brain condition.

[See [Connecting SRX Series Devices to Create a Chassis Cluster.](#)]

### User Interface and Configuration

- **Options for monitor traffic interfaces statement added (SRX Series)**—Starting in Junos OS Release 19.1R1, the options **write-file** and **read-file** under the **monitor traffic** command are included in the visible CLI.

[See [monitor traffic.](#)]

### VPNs

- **Certificate revocation list (SRX Series)**—Local certificates are being validated against certificate revocation list (CRL) even when CRL check is disabled. Starting in Junos OS Release 19.1R1, this can be stopped by disabling the CRL check through the Public Key Infrastructure (PKI) configuration. When CRL check is disabled, PKI will not validate local certificate against CRL.

[See [revocation-check \(Security PKI\)](#) and [Understanding Online Certificate Status Protocol and Certificate Revocation Lists.](#)]

- **Encryption algorithm (SRX Series)**—Starting in Junos OS Release 19.1R1, when AES-GCM 128-bit or AES-GCM 256-bit encryption algorithms are configured in the IPsec proposal, it is not mandatory to configure AES-GCM encryption algorithm in the corresponding IKE proposal.

[See [IPsec VPN Configuration Overview](#) and [encryption-algorithm \(Security IKE\).](#)]

- **Local or remote certificates (SRX Series)**—Starting in Junos OS Release 19.1R1, a commit check is added to prevent user from adding ., /, %, and space in a certificate identifier while generating a local or remote certificates or a key pair.

[See [certificate-id \(Security\)](#) and [Example: Configuring PKI.](#)]

- **Encryption algorithm support for high availability**—Starting in Junos OS Release 19.1R1, on the SRX5000 line of devices, you can configure the **aes-128-cbc** option at **set security ipsec internal security-association manual encryption algorithm**. You configure this option for encrypting the high availability link.

[See [internal \(Security IPsec\).](#)]

## SEE ALSO

[New and Changed Features | 318](#)

[Known Behavior | 336](#)

[Known Issues | 339](#)

[Resolved Issues | 341](#)

[Documentation Updates | 360](#)

[Migration, Upgrade, and Downgrade Instructions | 361](#)

## Known Limitations

### IN THIS SECTION

● [Application Security | 338](#)

● [Installation and Upgrade | 338](#)



- J-Web | 338
- Platform and Infrastructure | 338
- Switching | 338
- VPNs | 338

Learn about known limitations in this release 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 Security

- For any web application that is hosted on a content delivery network (CDN) such as AWS, Akamai, Azure, Fastly, Cloudflare and so on, application identification (AppID) classifies such applications incorrectly as "AMAZON" application instead of actual application.

## Installation and Upgrade

- The USB storage device stops working if it is removed while it is in initialization state. To avoid this issue, wait for a few seconds before removing the USB storage device. [PR1332360](#)

## J-Web

- The CLI terminal does not work in Java version 1.8 because of a security restriction in running the applet. [PR1341956](#)

## Platform and Infrastructure

- The gRPC connection with the gRPC collector resets upon RGO failover. [PR1402149](#)

## 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

- Tunnel debugging configuration is not synchronized to backup node. It needs to be configured again after RGO failover. [PR1450393](#)

## SEE ALSO

[New and Changed Features | 318](#)

[Changes in Behavior and Syntax | 329](#)

[Known Issues | 339](#)

[Resolved Issues | 341](#)

Documentation Updates | 360

Migration, Upgrade, and Downgrade Instructions | 361

## Open Issues

### IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 339](#)
- [Intrusion Detection and Prevention \(IDP\) | 340](#)
- [J-Web | 340](#)
- [Platform and Infrastructure | 340](#)
- [Routing Policy and Firewall Filters | 340](#)
- [SSL Proxy | 340](#)
- [VPNs | 341](#)

Learn about open issues in this release 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.

### Flow-Based and Packet-Based Processing

- On an SRX4600 device, the output of the **show route forwarding-table** command displays the next-hop IP address twice if the next hop is the st0 interface. The routing functionality is not impacted. [PR1290725](#)
- On all SRX Series devices, in a chassis cluster with Z-mode traffic and local (non-reth) interfaces configured, when using ECMP routing between multiple interfaces residing on both node0 and node1, if a session is initiated through one node and the return traffic comes in through the other node, packets might be dropped due to reroute failure. [PR1410233](#)
- Juniper Sky ATP does not escape the \ inside the username before the metadata is sent to the cloud. [PR1416093](#)
- Syslog severity level of msg subtype is end of policy is set to error although this message can be ignored. [PR1435233](#)
- The dynamic applications that relied on SSL proxy for matching have been deprecated. These dynamic applications are HTTPS, SMTPS, POP3S, and IMAPS. Use SSL, SMTP, POP3, and IMAP in all policies moving forward as they will now match for both encrypted and decrypted streams. [PR1444767](#)

- The flowd or srxpfe process might stop when SSL proxy service is used. [PR1450829](#)
- TCP session cannot time out properly upon receiving the TCP RESET packet, and the session timeout does not change to two second. [PR1467654](#)

## Intrusion Detection and Prevention (IDP)

- On the SRX Series devices, the commit or show command for IDP might not work if you keep running SNMP queries when large-scale IDP is used. [PR1444043](#)

## J-Web

- Forming a chassis cluster from J-Web by using the HA cluster wizard is not supported from Junos OS Release 12.1X47 onward for SRX5400 only. [PR1372518](#)

## Platform and Infrastructure

- Multiple monitor failures are seen on the rg1 interface after ISSU from Junos OS Release 17.4R1-S3 to Junos OS Release 18.1R1.9. [PR1354395](#)
- When SRX5K-SPC3 or MX-SPC3 cards are installed in slots 0 or 1 in a SRX5800 or MX960 chassis, EMI-radiated emissions are observed to be higher than regulatory compliance requirements [PR1479001](#)

## Routing Policy and Firewall Filters

- During commit, the nsd\_vrf\_group\_config\_lsys log messages are displayed. [PR1446303](#)
- 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](#)

## SSL Proxy

- Within an SSL-proxy configuration, if **trusted-ca** and **root-ca** have the same name, then in the associated SSL-T and I profiles do not get pushed to the Packet Forwarding Engine and thereby impact the SSL-proxy functionality. As a workaround, ensure to have different IDs or names for **trusted-ca** and **root-ca**. If already in the scenario, do the following to recover:
  - Configure different names for **trusted-ca** and **root-ca**.
  - From the CLI, restart NSD process using command `restart network-security`.

[PR1420859](#)

## 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](#)
- When multiple traffic selectors are configured on a particular VPN, the ikev2 process checks for a maximum of one DPD probe that is sent to the peer for the configured DPD interval. The DPD probe is sent to the peer if traffic flows over even one of the tunnels for the given VPN object. [PR1366585](#)
- VPN tunnels flap after a group is added or deleted a group in edit private mode in a clustered setup. [PR1390831](#)
- On the SRX5000 line of devices with SPC3, sometimes the IPsec tunnel may not come up after the configuration is changed from **responder-only** to **responder-only-no-rekey ikev1**. [PR1441320](#)
- With the NCP remote access solution, where IPsec traffic has to be encapsulated as TCP packets, TCP encapsulation for transit traffic is failing. [PR1442145](#)
- On SRX5400, SRX5600, and SRX5800 devices, when a chassis cluster is configured and IPsec tunnels are set up with DPD, after RG failover on the chassis cluster, some IPsec tunnels flap and there is temporary VPN traffic interruption until it restored automatically. [PR1450217](#)

## SEE ALSO

[New and Changed Features | 318](#)

[Changes in Behavior and Syntax | 329](#)

[Known Behavior | 336](#)

[Resolved Issues | 341](#)

[Documentation Updates | 360](#)

[Migration, Upgrade, and Downgrade Instructions | 361](#)

## 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: 19.1R3

### *Application Layer Gateways (ALGs)*

- The H323 call with NAT64 cannot be established on the SRX5000 line of devices. [PR1462984](#)
- The flowd or srpxfe 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](#)

### *Application Security*

- Introduction of default inspection limits to application identification to optimize CPU usage and improve resistance to evasive applications. [PR1454180](#)

### *Authentication and Access Control*

- The same source IP sessions are cleared when the IP entry is removed from the UAC table. [PR1457570](#)

### *Chassis Clustering*

- The flowd process might stop and restart, with the generation of a flowd core file, while using NAT PBA in a chassis cluster AA mode [PR1443148](#)
- The flowd process might stop on SRX Series devices with chassis cluster and IRB interface configured. [PR1446833](#)
- IP monitoring might fail on the secondary node. [PR1468441](#)
- An unhealthy node might become primary in an SRX4600 chassis cluster scenario. [PR1474233](#)

### *Class of Service (CoS)*

- The CoS rewrite rule does not work for the st0 interface. [PR1439401](#)

### *Flow-Based and Packet-Based Processing*

- The security flow traceoptions files fills in with RTSP ALG-related information. [PR1458578](#)
- Optimizations were made to improve the connections-per-second performance of an SPC3. [PR1458727](#)
- 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](#)
- AAWM policy rules for IMAP traffic sometimes might not get applied when the IMAP traffic passes through an SRX Series device. [PR1450904](#)
- The SRX Series devices stops and generates several core files. [PR1455169](#)
- When you try to reset the system configuration on an SRX1500 device using the reset config button, it does not work properly. [PR1458323](#)
- The SRX Series device might not be reachable when you initiate the **offline** command for a PIC. [PR1459037](#)

- The AAMWD process exceeds 85 percent RLIMIT\_DATA limitation due to memory leak. [PR1460619](#)
- The srxpfe or flowd process might stop if you change the sampling configuration. [PR1462610](#)
- On the SRX300 line of devices, you might encounter slow loading of Authentication-Table while using user identification. [PR1462922](#)
- The EA WAN SerDes gets into a stuck state, leading to continuous **DFE tuning timeout** errors and link staying down. [PR1463015](#)
- A core file is generated when you perform an ISSU on SRX Series devices. [PR1463159](#)
- Fragmented traffic might get looped on the fabric interface between the nodes in a rare case. [PR1465100](#)
- The pkid process keeps leaking memory on SRX Series devices. [PR1465614](#)
- The jbuf process usage may increase up to 99 percent after Junos OS upgrade. [PR1467351](#)
- The rpd process might stop after several changes to the flow-spec routes. [PR1467838](#)
- Packet Forwarding Engine core files are generated because SSL proxy is enabled on SRX Series devices and NFX Series routers. [PR1467856](#)
- 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](#)
- Tail drop on all ports is observed when any switch-side egress port gets congested. [PR1468430](#)
- FTP data connection might be dropped if the SRX Series device sends the FTP connection traffic through the dialer (dl) interface. [PR1468570](#)
- RPM test probe failure due to exceeded round-trip time is not working. [PR1471606](#)
- SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550HM devices may not retrieve the complete users or groups and user-group-mappings if the DC includes more than 20,000 users or groups. Use JIMS solution on SRX300, SRX320, SRX340, SRX345, SRX550, and SRX550HM devices when there are more than 20,000 users or devices or groups in the AD deployment. [PR1472601](#)
- Supports LLDP on reth interfaces. [PR1473456](#)
- Packet drop might be observed on the SRX300 line of devices when an interface is added to or removed from MACsec. [PR1474674](#)
- The flowd or srxpfe process might stop when advanced anti-malware service is used. [PR1480005](#)

### ***Install and Upgrade***

- Certificate error while validating configuration during Junos OS upgrade. [PR1474225](#)

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

- On E1 interface, BERT bits count is not within the range. [PR1445041](#)
- Static route through dl0.0 interface is not active. [PR1465199](#)
- MAC limiting on Layer 3 routing interfaces does not work. [PR1465366](#)

### ***Intrusion Detection and Prevention (IDP)***

- The flowd or srpxfe process stops and generates a core file. [PR1437569](#)
- Rogue .gz files in /var/tmp/sec-download/ might cause an offline security package update to fail. [PR1466283](#)

### ***J-Web***

- Editing a destination NAT rule in J-Web introduces a nonconfigured routing instance field. [PR1461599](#)

### ***Layer 2 Ethernet Services***

- The metric is not changing when configured under DHCP. [PR1461571](#)

### ***Network Address Translation (NAT)***

- The flowd and srpxfe process might stop when traffic is processed by both ALGs and NAT. [PR1471932](#)

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

- The flowd and srpxfe process might stop immediately after you commit the jflowv9 configuration or after you upgrade Junos OS to affected releases. [PR1471524](#)
- SNMP trap coldStart agent-address becomes 0.0.0.0. [PR1473288](#)
- All ingress packets are dropped if the traffic transit network is also the same network for LTE mPIM internal management. [PR1450046](#)

### ***Platform and Infrastructure***

- The SPC card might stop on SRX5000 line of devices. [PR1439744](#)
- Cm errors on certain MPC line cards are classified as major, which should be minor or non-fatal. [PR1449427](#)
- Traffic loss might occur when there are around 80,000 routes in the FIB. [PR1450545](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- A VM core might be observed if you are configuring a sampling rate of more than 65,535. [PR1461487](#)
- The aggregated Ethernet interface cannot be configured on an SRX4600 device. [PR1465159](#)
- Physically disconnecting the cable from the fxp0 interface causes hardware monitor failure. [PR1467376](#)
- The RGx may fail over after RG0 failover in a rare case. [PR1479255](#)

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

- Security policies cannot be synchronized between the Routing Engine and the Packet Forwarding Engine on SRX Series devices. [PR1453852](#)
- Traffic log shows a wrong custom-application name when the **alg ignore** option is used in the application configuration. [PR1457029](#)
- The NSD process might get stuck and cause problems. [PR1458639](#)



- 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 removed. [PR1476530](#)

### ***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](#)

### ***VPNs***

- After RG1 failover, IKE phase 1 SA is getting cleared. [PR1352457](#)
- Displaying an incorrect port number when scale is 1000 on IKEv1 AutoVPN tunnels. [PR1399147](#)
- The IKE and IPsec configuration under groups is not supported in this release. [PR1405840](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, if an existing IKE gateway configuration is changed from AutoVPN to Site-to-Site VPN, the IKE negotiation behavior remains in responder-only mode. [PR1413619](#)
- The VPN tunnel might flap when IKE and IPsec rekey happen simultaneously. [PR1421905](#)
- Old tunnel entries are also seen when a new tunnel negotiation happens from the peer device after a change in the IKE gateway configuration at the peer side. [PR1423821](#)
- The **show security ipsec statistics** command output of encrypted bytes fields count is not properly displayed.
- The P1 configuration delete message is not sent on loading the baseline configuration if there has been a prior change in the VPN configuration. [PR1432434](#)
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. [PR1455951](#)
- Traffic is not sent out through an IPsec VPN after upgrade 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](#)
- IPsec tunnels might lose connectivity after an SRX Series chassis cluster failover when using AutoVPN point-to-multipoint mode. [PR1469172](#)

## Resolved Issues: 19.1R2

### *Application Layer Gateways (ALGs)*

- On all SRX Series platforms, SIP/FTP ALG does not work when SIP traffic with source NAT goes through the SRX Series devices. [PR1398377](#)
- Unexpected forwarding sessions appear for tenant ALG SIP traffic in cross tenant case sometimes. [PR1409748](#)
- When both ALG and **rst-invalidate-session** are enabled, the TCP reset packet will be dropped by the SRX Series devices. This will impact all TCP ALG related traffic. [PR1430685](#)
- The H.323 connection might not be established when the H.323 packet passes SRX Series devices twice through different virtual routers. [PR1436449](#)
- Packet loss happens during cold synchronization from secondary node after rebooting. [PR1448252](#)

### *Application Security*

- Automatic application-identification download stops after going over the year and reboot. [PR1436265](#)
- With a single SPC3 card, AppQoS configured with unified policy can't provide throughput of more than 60 Gbps. [PR1439575](#)
- The flowd process core files might be seen when the traffic hits AppQoS policy. [PR1446080](#)

### *Authentication and Access Control*

- The CPU utilization of the uacd is high, about 100 percent, in the output of **show chassis routing-engine**. [PR1424971](#)

### *Chassis Clustering*

- The SNMP trap sends wrong information with manual failover. [PR1378903](#)
- Mixed mode (SPC3 coexisting with SPC2 cards) high availability (HA) IP monitoring fails on the secondary node with **secondary arp entry not found** error. [PR1407056](#)
- Node 0 stayed in secondary-hold status for long time but cannot change back to secondary status after manual failover in RG0. [PR1421242](#)
- Starting in Junos OS Release 18.4, a maximum of six PDN connects can be contained in PDP context response. Otherwise, the response is dropped. [PR1422877](#)
- Memory leaks might be seen on the jsqsyncd process on SRX Series chassis clusters. [PR1424884](#)
- RG0 failover sometimes causes FPC offline/present status. [PR1428312](#)
- Hardware failure is seen on both nodes in **show chassis cluster status**. [PR1452137](#)
- Chassis cluster control link will remain up even though control link is actually down. [PR1452488](#)

### *Class of Service (CoS)*

- Frequent issuance of the **show class-of-service spu statistics** command causes rtlogd to become busy. [PR1438747](#)

### *Flow-Based and Packet-Based Processing*

- Control traffic loss might be seen on SRX4600 platform. [PR1357591](#)
- On SRX1500 devices, the activity LED (right LED) for 1-Gigabit Ethernet/10-Gigabit Ethernet port is not on although traffic is passing through that interface. [PR1380928](#)
- Password recovery menu is not shown on SRX Series devices. [PR1381653](#)
- Invalid sessions time out over 48 hours with stress TCP traffic in the backup node. [PR1383139](#)
- On SRX4600 platform, the 40-Gigabit Ethernet interface might flap continuously by MAC local fault. [PR1397012](#)
- SRX Series devices might not strip VLAN added by native VLAN ID command. [PR1397443](#)
- CPU is hitting 100 percent with fragmented traffic. [PR1402471](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when PowerMode IPsec is enabled, the **show security flow statistics** and **show security flow session tunnel summary** commands do not count or display the number of packets processed within PowerMode IPsec, because these packets do not go through the regular flow path. [PR1403037](#)
- Throughput or latency performance of TCP traffic is dropped when TCP traffic is passing through from one logical system to another logical system. [PR1403727](#)
- The kernel might stop on the secondary node when committing **set system management-instance** command. [PR1407938](#)
- While PMI is on, IPsec-encrypted statistics on the Routing Engine **show security ipsec statistics** are not working anymore for fragment packets. [PR1411486](#)
- Traffic might be lost on SRX Series devices if IPsec session affinity is configured with **ipsec-performance-acceleration**. [PR1418135](#)
- If the traffic-log feature is configured, logs might incorrectly display IPv4 addresses in IPv6 format and packets might be dropped. [PR1421255](#)
- On PEM 0 or PEM 1 or fan, I2C failure major alarm might be set and cleared multiple times. [PR1413758](#)
- On SRX1500, SRX4100, SRX4200, SRX4600, and SRX5000 line of devices with SPC3 card, if SSL proxy is configured, the firewall FPC CPU might spike above 80 percent and traffic might be lost. [PR1414467](#)
- The input and output bytes or BPS statistic values might not be identical for the same size of packets. [PR1415117](#)
- The reth interfaces are now supported when configuring SSL decryption mirroring (mirror-decrypt-traffic interface). [PR1415352](#)
- Traffic would be dropped if SOF is enabled in a chassis cluster in active/active mode. [PR1415761](#)

- The command **show security firewall-authentication jims statistics** will output statistics of both the primary JIMS server and secondary JIMS server. [PR1415987](#)
- Juniper Sky ATP does not escape the \ inside the username before the metadata is sent to cloud. [PR1416093](#)
- The flowd process stops on the SRX5000 or SRX4000 lines of devices when large-size packets go through IPsec tunnel with the post-fragment check. [PR1417219](#)
- Traffic logging shows service-name **junos-dhcp-server** for UDP destination port 68. [PR1417423](#)
- Best path selected keeps changing at regular intervals even when no violation is reported. [PR1417926](#)
- Blacklist compilation failed syslog message not in later releases. [PR1418980](#)
- Group VPN IKE security associations cannot be established before RGO failover. [PR1419341](#)
- SSL proxy did not correctly warn users about unsupported certificates. [PR1419485](#)
- AAMW diagnostic script gives incorrect error **Error: Platform does not support SkyATP**. [PR1423378](#)
- The **show security flow session session-identifier <sessID>** command is not working if the session ID is bigger than 10 million on SRX4600 platform. [PR1423818](#)
- PIM neighbors might not come up on SRX Series chassis cluster. [PR1425884](#)
- When configuring a GRE tunnel (GRE-over-IPsec-tunnel) or an IPsec tunnel on an SRX Series device, the MTU of the tunnel interface is calculated incorrectly. [PR1426607](#)
- The IPsec traffic going through the SRX5000 line of devices with SPC2 cards installed causes high SPU CPU utilization. [PR1427912](#)
- Packet loss by FPGA backpressure on SPC3. [PR1429899](#)
- The flowd process might stop on the SRX5000 line of devices. [PR1430804](#)
- VPN traffic fails after primary node reboot or power off. [PR1433336](#)
- SRX550M running Junos OS Release 18.4R1 shows PEM 1 output failure message, whereas with Junos OS Release 15.1X49 or Junos OS Release 18.1R3.3 it does not show any alarms. [PR1433577](#)
- Intermittent packet drop might be observed if IPsec is configured. [PR1434757](#)
- The second IPsec ESP tunnel might not be able to establish between two IPv6 IKE peers. [PR1435687](#)
- On an SRX4600 device, core file generation might be observed and SPM might be in present state. [PR1436421](#)
- The ipfd process might crash when SecIntel is used. [PR1436455](#)
- Some webpages cannot be fully rendered. [PR1436813](#)
- SPMC version mismatch errors after Junos OS install using USB method. [PR1437065](#)
- The flowd or srpxfe process might crash when advanced anti-malware service is used. [PR1437270](#)
- Member of dynamically created VLANs information is not displaying on show VLANs. [PR1438153](#)

- Decryption traffic doesn't take PMI path after IPsec rekey (initiated by peer) when loopback interface is configured as external interface. [PR1438847](#)
- The flowd process stops and generates core files. [PR1438445](#)
- Security logs cannot be sent to external syslog server through TCP. [PR1438834](#)
- When lmd is rotating database, there is possibility that a reading access a NULL db at the same time, which generates core files. [PR1439186](#)
- The wmic process might stop and restart when using user firewall with Active Directory. [PR1439538](#)
- The flowd process stops on SRX550 or SRX300 line of devices when an SFP transceiver is plugged in. [PR1440194](#)
- Performance improvements were made to Screens, which benefit multi-socket systems. [PR1440677](#)
- The IKE pass-through packet might be dropped after source has undergone NAT. [PR1440605](#)
- While checking the flow session XML for source NAT under tenant, there is no value identifier for **tenant-name** ( < tenant>< /tenant>). [PR1440652](#)
- New CLI option to show only useful group information for an Active Directory user. [PR1442567](#)
- SPC2 wrongly forwarded packet to SPC3 core0 and core14. [PR1441234](#)
- The flowd or srpxfe process might crash when processing fragmented packets. [PR1443868](#)
- The J-Flow version 5 stops working after changing **input rate** value. [PR1446996](#)
- Packet loss happens during cold sync from secondary node after rebooting. [PR1447122](#)
- SPC3 Talus FPGA stuck on 0x3D or 0x69 golden version. [PR1448722](#)
- FTP data cannot pass through SRX320 4G wireless from FTP server to client. [PR1451122](#)
- Traffic forwarding on Q-in-Q port and VLAN tagging is not observed properly on R0. [PR1451474](#)
- The rpd process might stop and restart with an rpd core file created when committing the configuration. [PR1451860](#)
- Removed commit peers and **peers-synchronize** command from SRX Series devices. [PR1456661](#)

### **Infrastructure**

- Increase in Junos OS image size for Junos OS Release 19.1R1. [PR1423139](#)

### **Installation and Upgrade**

- Junos OS upgrade fails when partition option is used on SRX Series devices. [PR1449728](#)

### **Interfaces and Chassis**

- Both nodes in the SRX Series chassis cluster go into DB mode after downgrading to Junos OS Release 18.1. [PR1407295](#)
- Disabling the interface on the primary node causes traffic to get silently dropped through the secondary node. [PR1424705](#)
- MTU change after a CFM session is up can impact L2 Ethernet ping (loopback messages). If the new change is less than the value in the initial incarnation then L2 Ethernet ping would fail. [PR1427589](#)
- LFM remote loopback is not working as expected. [PR1428780](#)
- The LACP interface might flap if performing a failover. [PR1429712](#)

### **Interfaces and Routing**

- The fxp0 interface might redirect packet not destined to itself. [PR1453154](#)

### **Intrusion Detection and Prevention (IDP)**

- IDP install fails on one node because ApplD process gets stuck. [PR1336145](#)
- IDP might crash with the custom IDP signature. [PR1390205](#)
- Unable to configure **dynamic-attack-group** command. [PR1418754](#)
- NSD fails to push security zone to the Packet Forwarding Engine after reboot, if there is an active IDP rule configured with FQDN. [PR1420787](#)

### **J-Web**

- J-Web now supports defining SSL proxy and redirect (block page) profiles when a policy contains dynamic applications. [PR1376117](#)
- Risk report, when generated in IE browser, appears completely out of alignment and XML tags are displayed. [PR1415767](#)
- J-Web configuration change for an address set using the search function results in a commit error. [PR1426321](#)
- J-Web not working when logged in as read-only user. [PR1428520](#)
- On SRX Series devices, J-Web incorrectly displays port mode access for the link aggregation interfaces despite them being configured with multiple VLAN IDs and port mode trunk. [PR1430414](#)
- IRB interface is not available in zone option of J-Web. [PR1431428](#)

- When J-Web is used, if you log in to J-Web and navigate to multiple pages frequently, some error messages would be seen. It has no impact to service or traffic. This affects only J-Web UI. [PR1446081](#)
- The idle-timeout for J-Web access doesn't work properly. [PR1446990](#)
- J-Web fails to display the traffic log in event mode when stream mode host is configured. [PR1448541](#)

### **Layer 2 Ethernet Services**

- IPv6 address default route might not be installed from the received router advertisement message. [PR1411921](#)
- DHCP request might get dropped in DHCP relay scenario. [PR1435039](#)

### **Network Address Translation (NAT)**

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

### **Network Management and Monitoring**

- The **set system no-redirects** setting does not take effect for the reth interface. [PR894194](#)
- MIB OID **dot3StatsDuplexStatus** shows wrong status. [PR1409979](#)
- Partial traffic might get dropped on an existing LAG. [PR1423989](#)
- SNMPD might generate core files after restarting NSD process by restart **network-security gracefully**. [PR1443675](#)
- Control links are logically down on SRX Series chassis cluster when software version is Junos OS Release 12.3X48. [PR1458314](#)

### **Platform and Infrastructure**

- Memory leak might occur on the data plane during composite next-hop installation failure. [PR1391074](#)
- The **show security flow session** command fails with error messages when SRX4600 has over a million routing entries. [PR1408172](#)
- On SRX1500 platform, traffic is blocked on all interfaces after configuring the **interface-mac-limit** command on one interface. [PR1409018](#)
- Complete device outage might be seen when an SPU VM core file is generated. [PR1417252](#)
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10 (based system). [PR1417321](#)
- On SRX Series device, the flowd process might stop. [PR1417658](#)
- Routing Engine CPU utilization is high and eventd process is consuming a lot of resources. [PR1418444](#)

- On SRX4600 device, commit failed while configuring 2047 VLAN IDs on the reth interface. [PR1420685](#)
- SPC in slot1 of node0 remained in offline state for more than 1 hour after the cluster was upgraded from Junos OS Release 18.2R2-S1.3 to Junos OS Release 18.2X41.1. [PR1423169](#)
- Screen sync cookie causes 100 percent CPU utilization across all SPC3 cards of SRX5800, when packet rate is high. [PR1425332](#)
- The ipfd process might crash if the security intelligence feature is configured. [PR1425366](#)
- Alarms triggered due to high temperature when operating within expected temperatures. [PR1425807](#)
- 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](#)
- Uneven distribution of CPU with high PPS on device. [PR1430721](#)
- Packet Forwarding Engine crashes might be seen on SRX1500 platform. [PR1431380](#)
- The false license alarm might be seen even if there is a valid license. [PR1431609](#)
- The interface using LACP flaps when the Routing Engine is busy. [PR1435955](#)
- LACP traffic is distributed evenly on ingress child links but not on egress links. [PR1437098](#)
- The ksyncd process might crash and restart on SRX Series devices. [PR1440576](#)
- The configured RPM probe server hardware timestamp does not respond with correct timestamp to the RPM client. [PR1441743](#)
- The **show security flow session** command fails, generating an error message, when an SRX4100 or SRX4200 has around 1 million routing entries in the FIB. [PR1445791](#)
- LACP cannot work with the **encapsulation flexible-ethernet-services** configuration. [PR1448161](#)
- REST API process will get non-responsive when a number of requests come at a high rate. [PR1449987](#)

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

- Memory leak in nsd causes configuration change to not take effect after a commit. [PR1414319](#)
- The flowd process stops on SRX Series devices while deleting a lot of policies from Junos Space. [PR1419704](#)
- The NSD process might crash due to a memory corruption issue. [PR1419983](#)
- A commit warning is now presented to the user when a traditional policy is placed below a unified policy. [PR1420471](#)
- The dynamic-address summary's IP entry count does not include IP entries in the root logical system. [PR1422525](#)
- After a new alarm is created, the NSD process fails to restart because subcomponents fail. [PR1422738](#)
- DNS cache entry does not time out from device even after TTL=0. [PR1426186](#)



- The ipfd generates a core file while scaling. [PR1431861](#)
- An SRX1500 device allows only a maximum of 256 policies with counting enabled. [PR1435231](#)
- Two ipfd processes appear in **ps** command and the process pauses. [PR1444472](#)
- On all SRX Series devices that have policy counter configured, there is a potential risk where the network security process (NSD) on the Routing Engine cannot communicate with its Packet and Forwarding Engine counterpart (NSD-PFE) after either a chassis cluster failover, control link down, or Packet Forwarding Engine restart. At that point, it could no longer respond to network-security related commands and will not be able to complete coldsync for a newly joined node in chassis cluster environment. [PR1458639](#)

### ***Services Applications***

- The flowd process stops when SRX5800 devices works at SPC3 mix mode with 1 SPC3 card and 7 SPC2 cards. [PR1448395](#)
- In rare condition, SRX device Platform and Forwarding Engine might generate core file because corrupted or malformed HTTP long (over 64,000 bytes) messages hit security policy that is attached on ICAP redirect policy. [PR1460035](#)

### ***Unified Threat Management (UTM)***

- On SRX Series devices, when using Unified Policies and Web filtering (EWF) without SSL proxy, the Server Name Indication (SNI) might not be identified correctly and the RT\_UTM logs were recording incomplete information. [PR1410981](#)
- The device might not look up the blacklist first in the local Web filtering environment. [PR1417330](#)
- Unable to achieve better Avira antivirus TP on SRX4600 as mbuf high watermark is reached. [PR1419064](#)
- UTM Web filtering status shows down when using Hostname [routing-instance synchronization failure]. [PR1421398](#)
- When using Unified Policies, the base-filter for certain UTM profiles might not be applied correctly. [PR1424633](#)
- The custom-url-categories are now pushed correctly to the Packet Forwarding Engine under all circumstances. [PR1426189](#)
- The command **show security utm web-filtering status** 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](#)

### User Interface and Configuration

- Tenant system administrator cannot view its configuration with empty database message when using groups. [PR1422036](#)

### VPNs

- On SRX1500 device, when configuring IPsec VPN and BGP simultaneously, the kmd process might stop and generate a core file if BGP peers reach approximately 350. All of the VPN tunnels will be disconnected during the pause. [PR1336235](#)
- IPsec SA inconsistent on SPCs of node0 and node1 in chassis cluster. [PR1351646](#)
- Tunnel flapping is seen after doing RG0 failover. [PR1357402](#)
- SPC3 IKE SA detail output is not showing proper traffic statistics. [PR1371638](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, the **show security ike security-association detail** command does not display local IKE-ID field correctly. [PR1388979](#)
- 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](#)
- Idle IPsec VPN tunnels without traffic and with ongoing DPD probes can be affected during RG0 failover. [PR1405515](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when SRX Series device is configured in IKEv1 and NAT traversal is active, after a successful IPsec rekey, 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 with SPC3, when SRX Series device is configured to initiate IKEv2 reauthentication when NAT traversal is active, occasionally reauthentication might fail. [PR1414193](#)
- The iked process does not handle cases and core files might be generated when a remote gateway address is configured as an IPv6 address while the local interface where the tunnel is anchored has an IPv4 address. [PR1416081](#)
- The flowd/srxpfe process might stop when traffic selector is used for IPsec VPN. [PR1418984](#)
- The IKED process might stop when IKE and IPsec SA rekey happens simultaneously. [PR1420762](#)
- The 4G network connection might not be established if LTE mPIM card is in use. [PR1421418](#)
- Tenant system administrator can change VLAN assignment beyond the allocated tenant system. [PR1422058](#)
- The **show security ike sa detail** command shows incorrect values in the IPsec security associations column. [PR1423249](#)
- IPsec packet throughput might be impacted if NAT-T is configured and the fragmentation operation of post fragment happens. [PR1424937](#)

- On SRX Series devices with SPC3, the device does not send IKE delete notification to the peer if the traffic selector configuration is changed. [PR1426714](#)
- The kmd process stops and generates a core file after running the **show security ipsec traffic-selector** command. [PR1428029](#)
- In SPC3 and SPC2 mixed mode, IPsec SA is not getting cleared by executing the **clear security ipsec sa** command. [PR1428082](#)
- On the SRX5000 line of devices with SPC3, with P2MP and IKEv1 configured, if negotiation fails on the peer device, then multiple IPsec SA entries are created on the device if the peer keeps triggering a new negotiation. [PR1432852](#)
- IPsec rekey triggers for when sequence number in AH and ESP packet is about to exhaust is not working. [PR1433343](#)
- The kmd log shows resource temporarily unavailable repeatedly and VPNs might be down. [PR1434137](#)
- On SRX Series devices, fragments exit VPN traffic earlier than required by ingress packet sizes. [PR1435700](#)
- The IKED crashes 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 platforms with NAT-T scenario. [PR1444730](#)
- IPsec tunnels with distribution profile configuration will be renegotiated after perform RGO failover on SRX5000 line of devices with SPC3. [PR1446078](#)
- After a long time (a few hours) of traffic during mini PDT test, the number of IPsec tunnels number is much higher than expected. [PR1449296](#)
- IPsec VPN tunnels are losing routes for traffic selector randomly while tunnel is still up, causing complete outage. [PR1456301](#)

## Resolved Issues: 19.1R1

### *Application Security*

- Fail to match permit rule in AppFW rule set. [PR1404161](#)

### *Application Layer Gateways (ALGs)*

- DNS requests with the EDNS option might be dropped by the DNS ALG. [PR1379433](#)
- The H.323 protocol voice packets might be dropped. [PR1400630](#)

### *Chassis Clustering*

- Traffic loss occurs when the primary node is rebooting. [PR1372862](#)
- If using SRX Series chassis cluster and configuring four 100-Gigabit Ethernet interfaces on PIC 0, all the four interfaces might be down. [PR1387701](#)

- Traffic cannot pass through cross tenants after ISSU from Junos OS Release 18.3 to Junos OS Release 18.4. [PR1382467](#)
- The flowd process might stop if doing an ISSU upgrade. [PR1386522](#)
- The VDSL is not stable if there are sudden noises after configuring VDSL SOS feature. [PR1387133](#)
- ISSU status with error from Junos OS Release 18.2R1-S1 or Junos OS Release 18.2R1-S2 to Junos OS Release 18.2R1-S3. [PR1387947](#)
- The cluster IDs larger than 10 will cause FPCs to remain in offline on SRX4600 chassis cluster. [PR1390202](#)
- The MACsec on a physical port might not initialize properly when a new node is joined to the chassis cluster. [PR1396020](#)
- The flowd process stops if updating or deleting a GTP tunnel. [PR1404317](#)

#### ***Flow-Based and Packet-Based Processing***

- AppID classification logic has been improved for NetBIOS and RPC. [PR1357093](#)
- Control traffic loss may be seen on SRX4600 platform. [PR1357591](#)
- The Application identification (AppID) is supported for HTTP, SMTPS, POP3S, and IMAPS protocols. [PR1365810](#)
- When activating security flow traceoptions, the unfiltered traffic is captured. [PR1367124](#)
- Support for intelligent CLI-based autocomplete is added to secure-wire. [PR1372825](#)
- The pkid process might stop after RGO failover. [PR1379348](#)
- The reth interface flaps after doing an ISSU update. [PR1381475](#)
- Large file downloads slow down for many seconds. [PR1386122](#)
- Traffic might be processed by the VRRP backup when multiple VRRP groups are configured. [PR1386292](#)
- Traffic might be stopped after session created on SRX4600 platform. [PR1388735](#)
- The SRX Series device does not send messages frag needed and DF set back to the source host during path MTU discovery. [PR1389428](#)
- Packet loss might occur on unrelated traffic when AppQoS rate-limiter is applied on SRX4600 and SRX5000 platform using SPC3. [PR1394085](#)
- Request to display **dropped-illegal-packet** and **dropped-icmp-packet** configuration options. [PR1394720](#)
- Switching interface mode between family **ethernet-switching** and family **inet/inet6** might cause traffic loss. [PR1394850](#)
- These messages are seen: /kernel: tcp\_timer\_keep:Local(0x80000004:54652) Foreign(0x80000004:33160). [PR1396584](#)
- SRX Series devices connection to JIMS keeps flapping causes fail over to secondary JIMS. [PR1398140](#)
- On SRX4600 and SRX5000 devices, BGP packets might be dropped under high CPU usage. [PR1398407](#)

- VLAN push might not work on SRX1500. [PR1398877](#)
- Increase DAG feed scale number to 256 from 63. [PR1399314](#)
- The authd process might crash when issuing **show network-access requests pending** command during the authd restarting. [PR1401249](#)
- Unable to access to SRX Series platforms if the messages **kern.maxfiles limit exceeded by uid 65,534, please see tuning(7)** are seen. [PR1402242](#)
- Downloads may stall and/or completely fail when utilizing services that are reliant on TCP proxy. [PR1403412](#)
- Transit UDP 500/4500 traffic might not pass across SRX5000 Series devices when using SPC3/SPC2. [PR1403517](#)
- ISSU failed from Junos OS Release 18.3R1.9 to Junos OS Release 18.4R1.4. [PR1405556](#)
- The flowd process crashes and all cards are brought off. [PR1406210](#)
- The RG1 failover does not happen immediately when the SPC3 card crashes. [PR1407064](#)
- Session capacity of SRX340 is not match SRX345. [PR1410801](#)

#### ***Integrated User Firewall***

- Future group membership updates are not recognized by IUFW after a user's sAMAccountName is changed while the distinguished name (DN) remained the same. [PR1394049](#)

#### ***Interfaces and Routing***

- IPv4 multicast packets might not be broadcasted from the IRB interface on SRX1500 device. [PR1385934](#)
- SRX4600 10-gigabit Interface optics diagnostic access issue. [PR1395806](#)
- The 40-Gigabit and 100-Gigabit Ethernet ports may take a long time (about 30 s) to link up on SRX4600 platform. [PR1397210](#)
- High jsd or na-grpcd CPU usage might be seen when JET or JTI is not used. [PR1398398](#)
- SRX Series device cannot obtain IPv6 address through DHCPv6 when using a PPPoE interface with logical unit number greater than 0. [PR1402066](#)

#### ***Intrusion Detection and Prevention (IDP)***

- Unable to deploy IDP due to the IDP configuration cannot be committed. [PR1374079](#)
- Performance drops are seen in SRX345 and SRX340 platforms for IDP C2S policy. [PR1395592](#)

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

- Junos OS Release 18.3R1 cannot be installed using TFTP in boot loader on SRX300 platforms. [PR1390858](#)

### ***J-Web***

- On SRX Series platforms, the root password configured at first J-Web access (Skip to J-Web) does not work if password length is shorter than eight characters. [PR1371353](#)
- In the J-Web dashboard, the Security Resources widget did not display absolute values. This is now corrected. [PR1372826](#)
- Excluded addresses within J-Web Security Policy editor were not sufficiently differentiated versus normal addresses. They are now highlighted red for ease of identification. [PR1376112](#)
- The next-hop IP address is not displayed in the routing table in the J-Web. [PR1398650](#)
- J-Web page do not load after login with logical-system specific user. [PR1396879](#)
- Special character used in the preshared key is removed silently after a commit operation on J-Web. [PR1399363](#)
- Configuring using the CLI Editor in the J-Web generates an mgd core file. [PR1404946](#)
- The httpd-gk process crashes, leading to dynamic VPN failures and high Routing Engine CPU utilization 100 percent. [PR1414642](#)

### ***Layer 2 Ethernet Services***

- DHCPv6 clients might fail to get addresses on SRX Series platforms. [PR1392723](#)

### ***Multiprotocol Label Switching (MPLS)***

- BGP and OSPF flapped to cause traffic loss with RPD core on SRX550M cluster. [PR1366575](#)

### ***Network Address Translation (NAT)***

- The SRX Series devices might send the **noSuchInstance** value to the SNMP server in get-response during commit. [PR1357840](#)
- NAT64 and traceroute do not work correctly on an SRX Series device. [PR1376890](#)
- SPC3 mix mode NAT core at `../sysdeps/unix/sysv/linux/raise.c:55`. [PR1403583](#)

### ***Platform and Infrastructure***

- High httpd utilization after reboot failover. [PR1352133](#)
- Many chassis commands missing. [PR1363645](#)
- IP monitoring failure resulting in multiple interfaces disappearing from forwarding table. [PR1371500](#)
- Some error messages could be seen when running **show interface extensive** command from CLI or Junos Space. [PR1380439](#)
- Traffic loss seen in Layer 2 VPN with GRE tunnel. [PR1381740](#)

- Redundancy group failover caused by interface monitoring failure is slow to master state at PFE. [PR1385521](#)
- Login class with allowed-days and specific access-start/access-end does not work as expected. [PR1389633](#)
- GW lcores and srpxfe cores at `../src/pfe/usp/rt/applications/ipsec/ipsec_rt_forge_util.c:59` when loading 18.4 image. [PR1392580](#)
- The flowd process crashes if it goes into a dead loop. [PR1403276](#)
- HA failed with the failure code **HW** after loading the image. [PR1406029](#)

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

- When SSL-Forward-Poxy is configured in a unified policy along with the action of Reject+Redirect, a block page was not presented to the user for HTTPS sites. [PR1375823](#)
- The **show security flow session** command now fully supports the dynamic-application construct. [PR1387449](#)
- The nsd process crashes and generates a core file. [PR1388719](#)

### ***Routing Protocols***

- vFPC may continuously crash on vMX platform. [PR1364624](#)

### ***Services Applications***

- SRX5600 HA SPC2, the ICAP redirect objects are in use even after clearing TCP sessions. [PR1390835](#)

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

- Fan speed goes up and down continuously on SRX1500. [PR1335523](#)

### ***Unified Threat Management (UTM)***

- Source and destination zone information are added in the UTM log. [PR1326271](#)
- EWF server status shows UP when 443 is specified as server port. [PR1383695](#)
- Whitelist/Blacklist does not work for HTTPS traffic going through the Web proxy. [PR1401996](#)
- On SRX Series, when configuring Enhanced Web Filtering on the CLI, the autocomplete function did not properly handle or suggest custom categories. [PR1406512](#)
- On SRX Series, when using Unified Policies and Webfiltering (EWF) without SSL-Proxy in Junos OS Release 18.4R1, the Server Name Indication (SNI) may not be identified correctly and the RT\_UTM logs were recording incomplete information. [PR1410981](#)

## VPNs

- ISSU from Junos OS.Release 15.1X49-D120 to Junos OS.Release 15.1X49-D130 seeing KMD core seen at **0x08228b83** in `iked_advpn_timer_cb_delete_inactive_shortcut_tunnel` (`timer_ctx=0x99d8000`) at `../../../../src/usp/usr.sbin/iked/core/iked_advpn.c:227`. [PR1340973](#)
- Dot usage in CA profile name causes issues when the pkid process is restarted. [PR1351727](#)
- A few VPN tunnels do not forward traffic after RG1 failover. [PR1394427](#)
- The kmd process might crash when SNMP polls for the IKE SA. [PR1397897](#)
- VPN does not recover on the high-end standalone SRX Series device when CLI operation **restart ipsec-key-management** is done. [PR1400712](#)
- Syslog is not generated when the ike gateway rejects a duplicate IKE ID connection. [PR1404985](#)
- Not all the tunnels are deleted when the authentication algorithm in IPsec proposal is changed. [PR1406020](#)
- Multiple flowd core files are observed with IPsec acceleration with fragmentation traffic. [PR1407910](#)
- Traffic drops on peer due to bad SPI after first re-authentication. [PR1412316](#)

## SEE ALSO

[New and Changed Features | 318](#)

[Changes in Behavior and Syntax | 329](#)

[Known Behavior | 336](#)

[Known Issues | 339](#)

[Documentation Updates | 360](#)

[Migration, Upgrade, and Downgrade Instructions | 361](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 19.1R3 documentation for the SRX Series.

## SEE ALSO

[New and Changed Features | 318](#)

[Changes in Behavior and Syntax | 329](#)

[Known Behavior | 336](#)

[Known Issues | 339](#)



## 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

Known Behavior | 336

Known Issues | 339

Resolved Issues | 341

Documentation Updates | 360

# 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 we've 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 you 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](https://prsearch.juniper.net).
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](https://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/](https://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](mailto: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

9 February 2023—Revision 13, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

25 November 2022—Revision 13, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

31 March 2022—Revision 12, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

25 November 2021—Revision 11, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

7 October 2021—Revision 11, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

15 July 2021—Revision 10, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

17 June 2021—Revision 9, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

22 April 2021—Revision 8, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

11 March 2021—Revision 7, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

13 January 2021—Revision 6, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

10 July 2020—Revision 5, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

18 June 2020—Revision 4, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

19 May 2020—Revision 3, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

9 April 2020—Revision 2, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

1 April 2020—Revision 1, Junos OS Release 19.1R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

6 December 2019—Revision 2, Junos OS Release 19.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

28 November 2019—Revision 1, Junos OS Release 19.1R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

23 July 2019—Revision 14, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

18 July 2019—Revision 13, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

4 July 2019—Revision 12, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

27 June 2019—Revision 11, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

20 June 2019—Revision 10, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

6 June 2019—Revision 9, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

28 May 2019—Revision 8, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

9 May 2019—Revision 7, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

2 May 2019—Revision 6, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

19 April 2019—Revision 5, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.



11 April 2019—Revision 4, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

9 April 2019—Revision 3, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

3 April 2019—Revision 2, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.

27 March 2019—Revision 1, Junos OS Release 19.1R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos fusion.