

Release Notes

Published
2023-04-21

Junos[®] OS 19.4R1 Release Notes

SUPPORTED ON

- ACX Series, EX Series, JRR Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series

HARDWARE HIGHLIGHTS

- Wi-Fi Mini-Physical Interface Module (SRX320, SRX340, SRX345, and SRX550M)

SOFTWARE HIGHLIGHTS

- Support for EVPN routing policies (ACX5448)
- Inline monitoring services (MX Series with MPCs excluding MPC10E linecards)
- Support for BGP PIC Edge with BGP labeled unicast (MX Series and PTX Series)
- Integrating RIFT protocol into Junos OS (MX Series and VMX virtual routers)
- Support for flexible algorithm in IS-IS for segment routing-traffic engineering (MX Series and PTX Series)
- Junos Multi-Access User Plane (MX240, MX480, MX960)
- Support for Lawful Intercept on Junos Multi-Access User Plane (MX240, MX480, MX960)
- Precision Time Protocol (PTP) transparent clock (QFX5120 and QFX5210)
- Additional support for Bidirectional Forwarding Detection (QFX5110, QFX5120, QFX5200, and QFX5210)
- Selectively disable midstream APBR (SRX Series and vSRX)

- Improved query performance in on-box reporting (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)
- UTM Support for Active/Active Chassis Cluster (SRX Series)

IN FOCUS GUIDE

- [Use this new guide to quickly learn about the most important Junos OS features and how you can deploy them in your network.](#)

Release Notes: Junos[®] OS Release 19.4R1 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, JRR Series, and Junos Fusion

21 April 2023

Contents	Introduction 14
	Junos OS Release Notes for ACX Series 14
	What's New 15
	EVPN 15
	Junos OS XML API and Scripting 16
	MPLS 18
	OAM 18
	Routing Protocols 18
	System Logging 19
	Software Defined Networking (SDN) 19
	What's Changed 20
	General Routing 20
	Junos Telemetry Interface 20
	Routing Protocols 21
	System Logging 21
	Known Limitations 21
	General Routing 22
	Open Issues 23
	General Routing 23

Resolved Issues | 24

General Routing | 24

Layer 2 Ethernet Services | 26

Platform and Infrastructure | 26

Routing Protocols | 26

Documentation Updates | 27

Feature Guides Are Renamed As User Guides | 27

Migration, Upgrade, and Downgrade Instructions | 28

Upgrade and Downgrade Support Policy for Junos OS Releases | 28

Junos OS Release Notes for EX Series Switches | 29

What's New | 29

Authentication, Authorization, and Accounting | 31

Class of Service | 31

EVPN | 31

Junos OS XML, API, and Scripting | 33

Junos Telemetry Interface | 34

Layer 2 Features | 35

MPLS | 37

Multicast | 37

Operation, Administration, and Maintenance (OAM) | 37

Port Security | 38

Routing Policy and Firewall Filters | 39

System Logging | 39

System Management | 39

User Interface and Configuration | 40

What's Changed | 40

What's Changed in Release 19.4R1-S3 | 41

What's Changed in Release 19.4R1 | 41

Known Limitations | 42

Open Issues | 43

Authentication and Access Control | 44

Class of Service (CoS) | 44

EVPN | 44

General Routing | 44

Infrastructure	46
Interfaces and Chassis	46
Junos Fusion Enterprise	46
Junos Fusion Satellite Software	46
Layer 2 Ethernet Services	46
Layer 2 Features	47
Platform and Infrastructure	47
Routing Protocols	47
User Interface and Configuration	48
Resolved Issues	48
Authentication and Access Control	49
EVPN	49
General Routing	49
Infrastructure	51
Interfaces and Chassis	52
Junos Fusion Enterprise	52
Junos Fusion Satellite Software	52
J-Web	52
Layer 2 Ethernet Services	52
Layer 2 Features	52
Platform and Infrastructure	53
Routing Protocols	53
User Interface and Configuration	54
Virtual Chassis	54
Documentation Updates	54
Feature Guides Are Renamed As User Guides	55
Migration, Upgrade, and Downgrade Instructions	55
Upgrade and Downgrade Support Policy for Junos OS Releases	55
Junos OS Release Notes for JRR Series	56
What's New	57
Hardware	57
What's Changed	58
Known Limitations	58

Open Issues | 59

General Routing | 59

Resolved Issues | 59

Documentation Updates | 60

Feature Guides Are Renamed As User Guides | 60

Migration, Upgrade, and Downgrade Instructions | 61

Upgrade and Downgrade Support Policy for Junos OS Releases | 61

Junos OS Release Notes for Junos Fusion Enterprise | 62

What's New | 62

What's Changed | 63

Known Limitations | 63

Open Issues | 64

Junos Fusion for Enterprise | 64

Resolved Issues | 65

Documentation Updates | 65

Feature Guides Are Renamed As User Guides | 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 Provider Edge | 72

What's New | 72

What's Changed | 73

Known Limitations | 73

Open Issues | 74

Junos Fusion for Provider Edge | 74

Resolved Issues | 74

Documentation Updates | 75

Feature Guides Are Renamed As User Guides | 75

Migration, Upgrade, and Downgrade Instructions | 76

Basic Procedure for Upgrading an Aggregation Device | 76

Upgrading an Aggregation Device with Redundant Routing Engines | 79

Preparing the Switch for Satellite Device Conversion | 79

Converting a Satellite Device to a Standalone Device | 81

Upgrading an Aggregation Device | 83

Upgrade and Downgrade Support Policy for Junos OS Releases | 83

Downgrading from Junos OS Release 19.4 | 84

Junos OS Release Notes for MX Series 5G Universal Routing Platform | 84

What's New | 85

Hardware | 86

Class of Service | 89

EVPN | 89

Forwarding and Sampling | 90

General Routing | 91

High Availability (HA) and Resiliency | 91

Interfaces and Chassis | 91

Junos OS, XML, API, and Scripting | 94

Junos Telemetry Interface | 95

Layer 2 Features | 97

Layer 2 VPN | 97

MPLS | 97

Multicast | 98

Network Management and Monitoring | 99

OAM | 99

Routing Policy and Firewall Filters | 99

Routing Protocols | 100

Services Applications | 103

Software-Defined Networking | 103

Software Licensing | 104

Subscriber Management and Services | 104

System Logging	105
What's Changed	106
General Routing	107
Interfaces and Chassis	107
Junos Telemetry Interface	108
MPLS	109
Network Management and Monitoring	109
Routing Protocols	109
Services Applications	109
Software-Defined Networking	109
Subscriber Management and Services	110
System Logging	111
Known Limitations	111
General Routing	112
Interfaces and Chassis	113
MPLS	113
Platform and Infrastructure	113
Routing Protocols	113
Open Issues	114
Application Layer Gateways	115
Class of Service	115
EVPN	115
Forwarding and Sampling	115
General Routing	116
Infrastructure	124
Interfaces and Chassis	124
Layer 2 Features	126
Layer 2 Ethernet Services	126
MPLS	126
Network Management and Monitoring	127
Next Gen Services MX-SPC3 Services Card	127
Platform and Infrastructure	127
Routing Protocols	129
Services Applications	130

Subscriber Access Management	130
VPNs	130
Resolved Issues	131
Resolved Issues: 19.4R1	132
Documentation Updates	152
Feature Guides Are Renamed As User Guides	153
Migration, Upgrade, and Downgrade Instructions	153
Basic Procedure for Upgrading to Release 19.4	154
Procedure to Upgrade to FreeBSD 11.x based Junos OS	154
Procedure to Upgrade to FreeBSD 6.x based Junos OS	157
Upgrade and Downgrade Support Policy for Junos OS Releases	159
Upgrading a Router with Redundant Routing Engines	159
Downgrading from Release 19.4	160
Junos OS Release Notes for NFX Series	160
What's New	161
General routing	162
Hardware	162
Architecture	162
What's Changed	163
System Logging	164
Known Limitations	164
Interfaces	165
Platform and Infrastructure	165
Open Issues	165
Mapping of Address and Port with Encapsulation (MAP-E)	166
Interfaces	166
Platform and Infrastructure	166
Virtual Network Functions (VNFs)	167
Resolved Issues	168
Class of Service	169
High Availability	169
Interfaces	169
Layer 2 Ethernet Services	169
Platform and Infrastructure	169

Routing Protocols	171
SNMP	171
Virtual Network Functions (VNFs)	171
Documentation Updates	172
Feature Guides Are Renamed As User Guides	172
Migration, Upgrade, and Downgrade Instructions	173
Upgrade and Downgrade Support Policy for Junos OS Releases	173
Basic Procedure for Upgrading to Release 19.4	173
Junos OS Release Notes for PTX Series Packet Transport Routers	175
What's New	176
General Routing	177
Hardware	177
High Availability (HA) and Resiliency	177
Junos OS, XML, API, and Scripting	177
Junos Telemetry Interface	178
MPLS	180
Routing Protocols	181
Services Applications	182
Software Defined Networking	183
System Logging	183
What's Changed	184
General Routing	184
Interfaces and Chassis	184
Junos Telemetry Interface	185
Routing Protocols	186
Software-Defined Networking	186
System Logging	186
Known Limitations	187
General Routing	187
Open Issues	188
General Routing	189
Infrastructure	190
Layer 2 Ethernet Services	190
MPLS	190

Routing Protocols | 191

Resolved Issues | 191

Forwarding and Sampling | 192

General Routing | 192

Infrastructure | 194

Interfaces and Chassis | 194

Layer 2 Ethernet Services | 194

MPLS | 194

Platform and Infrastructure | 194

Routing Protocols | 194

VPNs | 195

Documentation Updates | 195

Feature Guides Are Renamed as User Guides | 196

Migration, Upgrade, and Downgrade Instructions | 196

Basic Procedure for Upgrading to Release 19.4 | 196

Upgrade and Downgrade Support Policy for Junos OS Releases | 199

Upgrading a Router with Redundant Routing Engines | 200

Junos OS Release Notes for the QFX Series | 201

What's New | 201

EVPN | 202

General Routing | 206

Interfaces and Chassis | 206

Junos OS XML API and Scripting | 207

Junos Telemetry Interface | 208

Layer 2 Features | 208

MPLS | 209

Routing Protocols | 209

Software Defined Networking (SDN) | 209

System Logging | 210

System Management | 211

VLAN Infrastructure | 212

What's Changed | 212

General Routing | 213

Interfaces and Chassis | 213

Junos Telemetry Interface	213
Management	214
Routing Protocols	214
Software Defined Networking (SDN)	214
System Logging	214
Known Limitations	215
Layer 2 Ethernet Services	216
Layer 2 Features	216
Network Management and Monitoring	216
Platform and Infrastructure	216
Routing Protocols	217
Open Issues	217
Class of Service (CoS)	218
EVPN	218
High Availability (HA) and Resiliency	219
Interfaces and Chassis	219
Junos Fusion for Provider Edge	219
Layer 2 Features	219
Layer 2 Ethernet Services	220
MPLS	220
Platform and Infrastructure	220
Routing Protocols	223
Resolved Issues	224
Class of Service (CoS)	224
EVPN	225
Forwarding and Sampling	225
Interfaces and Chassis	225
Layer 2 Features	225
MPLS	226
Platform and Infrastructure	226
Routing Protocols	230
User Interface and Configuration	231
Documentation Updates	232
Feature Guides Are Renamed As User Guides	232

Migration, Upgrade, and Downgrade Instructions | 232

Upgrading Software on QFX Series Switches | 233

Installing the Software on QFX10002-60C Switches | 235

Installing the Software on QFX10002 Switches | 235

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 | 236

Installing the Software on QFX10008 and QFX10016 Switches | 238

Performing a Unified ISSU | 242

Preparing the Switch for Software Installation | 243

Upgrading the Software Using Unified ISSU | 243

Upgrade and Downgrade Support Policy for Junos OS Releases | 245

Junos OS Release Notes for SRX Series | 246

What's New | 247

Application Security | 247

Chassis Clustering | 248

Flow-Based and Packet-Based Processing | 248

General Packet Radio Switching (GPRS) | 249

Hardware | 249

Interfaces and Chassis | 250

Intrusion Detection and Prevention (IDP) | 250

Junos OS XML API and Scripting | 251

J-Web | 252

Logical Systems and Tenant Systems | 252

Network Management and Monitoring | 253

System Logging | 254

Unified Threat Management (UTM) | 254

VPNs | 254

What's Changed | 255

Application Security | 256

Authentication and Access Control | 259

Class of Service | 259

General Routing | 259

J-Web | 259

Network Management and Monitoring | 259

Port Security	261
Routing Protocols	261
System Logging	261
VPNs	261
Known Limitations	262
Application Layer Gateways (ALGs)	262
Ethernet Switching	262
Flow-Based and Packet-Based Processing	262
J-Web	263
Open Issues	263
ALG	264
Flow-Based and Packet-Based Processing	264
IDP	264
J-Web	264
Platform and Infrastructure	264
Routing Policy and Firewall Filters	264
VPNs	265
Resolved Issues	266
Application Layer Gateways	266
Application Security	266
Authentication and Access Control	266
Chassis Clustering	266
Class of Service	267
Flow-Based and Packet-Based Processing	267
Interfaces and Chassis	269
Intrusion Detection and Prevention (IDP)	269
J-Web	269
Layer 2 Ethernet Services	270
Network Address Translation	270
Network Management and Monitoring	270
Platform and Infrastructure	270
Routing Policy and Firewall Filters	271
Services Applications	271
Unified Threat Management	271

VLAN Infrastructure	272
VPNs	272
Documentation Updates	273
Feature Guides Are Renamed As User Guides	274
Migration, Upgrade, and Downgrade Instructions	274
Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases	274
Upgrading Using ISSU	275
Licensing	276
Compliance Advisor	276
Finding More Information	276
Documentation Feedback	277
Requesting Technical Support	278
Self-Help Online Tools and Resources	278
Creating a Service Request with JTAC	279
Revision History	279

Introduction

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

These release notes accompany Junos OS Release 19.4R1 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, JRR Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

- **Feature Guides Are Renamed As User Guides**—Starting with Junos OS 19.4R1, we renamed our Feature Guides to User Guides to better reflect the purpose of the guides. For example, the BGP Feature Guide is now the BGP User Guide. We didn't change the URLs of the guides, so any existing bookmarks you have will continue to work. To keep the terminology consistent on our documentation product pages, we renamed the Feature Guides section to User Guides. To find documentation for your specific product, check out this link.
- **New In Focus Guide**—Starting on Junos Release 19.4R1, we are introducing a new document called In Focus that provides details on the most important features for the release in one place. We hope this document will quickly get you to the latest information about Junos features. Let us know if you find this information useful by sending email to techpubs-comments@juniper.net.
- **Important Information:**
 - [Upgrading Using ISSU on page 275](#)
 - [Licensing on page 276](#)
 - [Compliance Advisor on page 276](#)
 - [Finding More Information on page 276](#)
 - [Documentation Feedback on page 277](#)
 - [Requesting Technical Support on page 278](#)

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 15](#)
- [What's Changed | 20](#)
- [Known Limitations | 21](#)

- Open Issues | 23
- Resolved Issues | 24
- Documentation Updates | 27
- Migration, Upgrade, and Downgrade Instructions | 28

These release notes accompany Junos OS Release 19.4R1 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

- EVPN | 15
- Junos OS XML API and Scripting | 16
- MPLS | 18
- OAM | 18
- Routing Protocols | 18
- System Logging | 19
- Software Defined Networking (SDN) | 19

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

EVPN

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

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

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

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

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

Junos OS XML API and Scripting

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

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

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

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

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

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

MPLS

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

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

OAM

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

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

Routing Protocols

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

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

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

Following are the SRGB precedences for OSPF protocol:

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

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

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

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:

- Additional information is now logged for MGD-FFP intermodule communication.
- Commit errors that previously were only shown onscreen are now logged.

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

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

Software Defined Networking (SDN)

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

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

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

SEE ALSO

[What's Changed | 20](#)

[Known Limitations | 21](#)

[Open Issues | 23](#)

[Resolved Issues | 24](#)

[Documentation Updates | 27](#)

[Migration, Upgrade, and Downgrade Instructions | 28](#)

What's Changed

IN THIS SECTION

- [General Routing](#) | 20
- [Junos Telemetry Interface](#) | 20
- [Routing Protocols](#) | 21
- [System Logging](#) | 21

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

General Routing

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, MX Series, PTX Series, and SRX Series)**—Starting with Junos OS Release 19.4R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.
[See [commit \(System\)](#).]
- **Support for `gigether-options` statement (ACX5048, ACX5096)**—Junos OS supports the `gigether-options` statement at the **edit interfaces interface-name** hierarchy on the ACX5048 and ACX5096 routers. Previously, support for the `gigether-statement` was deprecated. See [gigether-options](#) and [gigether-statement](#).
- **IPv6 address in the prefix TIEs displayed correctly (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.

Junos Telemetry Interface

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

[See [Guidelines for gRPC and gNMI Sensors](#).]

Routing Protocols

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

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

System Logging

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

SEE ALSO

What's New 15
Known Limitations 21
Open Issues 23
Resolved Issues 24
Documentation Updates 27
Migration, Upgrade, and Downgrade Instructions 28

Known Limitations

IN THIS SECTION

- [General Routing | 22](#)

Learn about known limitations in Junos OS Release 19.4R1 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

- ACX6360-OR Telemetry infrastructure does not support the interface-filtering capability. Therefore, after you enable a particular sensor for telemetry, it is turned on for all the interfaces. [PR1371996](#)
- When the timing configuration and the corresponding interface configuration is flapped for multiple times in iteration, PTP is stuck in initialize state where the ARP for the neighbor is not resolved. In issue state, BCM hardware block get into inconsistency state, where the lookup is failing. [PR1410746](#)
- The port LEDs and the system LED glow during system or vmhost halt state on all ACX Series devices. [PR1430129](#)
- With an asymmetric network connection (for example, 10G MACsec port connected to 10G channelized port), high and asymmetric T1 and T4 time errors are seen, which introduces a high 2 way time error. This introduces different CF updates in forward and reverse paths. [PR1440140](#)
- With MACsec feature enabled and introduction of traffic, the peak-to-peak value varies with the percentage of traffic introduced. [PR1441388](#)
- You cannot directly associate a policer in an IFL on ACX Series devices. It has to be achieved using "filters" only. In this case, "family any" filter can be configured. Also, ACX5448 hardware ASIC does not support "egress policing". Egress shaping (H-QoS at the IFL level) can be used instead. [PR1446376](#)
- ACX Series routers support rate of join of IGMPv3 users only around 900 joins per second. Same is updated to PDD. [PR1448146](#)

SEE ALSO

[What's New | 15](#)

[What's Changed | 20](#)

[Open Issues | 23](#)

[Resolved Issues | 24](#)

[Documentation Updates | 27](#)

[Migration, Upgrade, and Downgrade Instructions | 28](#)

Open Issues

IN THIS SECTION

- [General Routing | 23](#)

Learn about open issues in Junos OS Release 19.4R1 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

- High CPU usage for fxpc processes with class-of-service changes on aggregated Ethernet interfaces. [PR1407098](#)
- On ACX1000, ACX2000, ACX4000, ACX5048, and ACX5096 devices, after a new child logical interface with VLAN and filter is added on an aggregated Ethernet physical interface or after changing the VLAN ID of a child logical interface with filter, traffic over the aggregated Ethernet physical interface might get filtered with that filter on the child logical interface. For example, ae-0/0/0 is a physical interface and ae-0/0/0.100 is a logical interface. [PR1407855](#)
- Layer 2 rewrites happens on regular Bridge domain and VLAN interfaces. Though there are some service dependencies, VPLS in this case, in which the Egress interface map table is not updated properly with the Layer 2 rewrite map id that causes rewrite to stop working. [PR1414414](#)
- CoS table error can sometimes cause traffic outages and SNMP timeouts if the optic is removed and inserted back. [PR1418696](#)
- On the ACX5000 line of devices, high CPU usage might be seen for the fxpc process. [PR1419761](#)
- The **request system reboot** command on the ACX5448 and ACX5448-D box triggers a reboot on the host (Linux) instead of being limited to Junos OS. [PR1426486](#)
- DHCP clients are not able to scale to 96000. [PR1432849](#)
- The time consumed for 1-Gigabit performance is not equal to the time for 10-Gigabit performance. Compensation is done to bring the mean value under class A, but the peak-to-peak variations are high and can go beyond 100ns. It has a latency variation with peak-to-peak variations of around 125ns-250ns without any traffic (for example, 5-10% of the mean latency introduced by each phy, which is around 2.5us). [PR1437175](#)
- Memory leaks are expected. [PR1438358](#)

- ACX can support rate of join of IGMPv3 users only around 900 joins per second. Same will be updated to PDD. [PR1448146](#)
- Drop profile maximum threshold might not reach its limit when the packet size is other than 1000 bytes. This is due to the current design limitation. [PR1448418](#)
- Not possible to form 125000 IGMP groups with ACX5448 receiving 125000 IGMP v2 Reports per second. This is a product limitation from BCM and CPU host path queuing model. [PR1454465](#)

SEE ALSO

[What's New | 15](#)

[What's Changed | 20](#)

[Known Limitations | 21](#)

[Resolved Issues | 24](#)

[Documentation Updates | 27](#)

[Migration, Upgrade, and Downgrade Instructions | 28](#)

Resolved Issues

IN THIS SECTION

- [General Routing | 24](#)
- [Layer 2 Ethernet Services | 26](#)
- [Platform and Infrastructure | 26](#)
- [Routing Protocols | 26](#)

Learn which issues were resolved in Junos OS Release 19.4R1 for ACX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- On ACX5000 **MacDrainTimeOut** and **bcm_port_update failed: Internal error** error is seen. [PR1284590](#)
- bcmDPC task is high even though Interrupt START_BY_START flag is set to 0. [PR1329656](#)

- The AE interface with LACP stays down after the router reboots if link-speed is configured. [PR1357012](#)
- On ACX Series devices, the LED on the GE interface goes down when speed 10M is added. [PR1385855](#)
- Link Fault Signaling (LFS) doesn't work on ACX5448 10-, 40-, and 100-Gigabit Ethernet interfaces. [PR1401718](#)
- The optic comes with Tx enabled by default. Because the port is administratively disabled, the port is stopped. However, because the port has not been started, it does not disable Tx. [PR1411015](#)
- The ACX5448:40G FEC on ACX5448, which is FEC enabled by default, must be aligned with the MX and QFX platforms, where FEC is NONE. [PR1414649](#)
- On the ACX5448-X:SKU and ACX5448-D, 96000 ARPs get populated. However, only 47000 NH entries are present. Around 50 percent of packet drop is observed. [PR1426734](#)
- Chassisd might crash with unsupported hcos configuration when an MX104 is used as a fusion aggregation device. [PR1430076](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- Deviation in traffic rate in the queue is around 8 % to 10% percent in some cases. [PR1436297](#)
- 1PPS performance metrics (class A) of G.8273.2 are not met for 1G interfaces because of the variable latency added by the Vitesse PHY. [PR1439231](#)
- The interface on ACX1100 devices remains down when using SFP-1FE-FX (740-021487). [PR1439384](#)
- Transit DHCP packets are not punted to CPU and are transparently passthrough. [PR1439518](#)
- When the interface is flapped between channelized configurations, 25-Gigabit Ethernet to 100-Gigabit the aggregated Ethernet interface configuration is not cleaned up properly. [PR1441374](#)
- In an ACX5448 platforms, when the PFE failed to allocate packet buffer, portion of packet memories may not be freed. [PR1442901](#)
- RED drops might be seen after link flaps or CoS configuration changes. [PR1443466](#)
- ACX5448/18.3R1-S4.1 is not performing proper dot1p CoS rewrite on interfaces configured with l2circuit/local-switching/family ccc. [PR1445979](#)
- On ACX Series, the auto exported route between VRFs might not respond for icmp echo requests. [PR1446043](#)
- l2circuit with a **backup-neighbor** (hot-standby) configured might stop forwarding traffic after failovers. [PR1449681](#)
- oper-state for et interface does not transition from 'init' to 'Normal'. [PR1449937](#)
- RMPC core files are found after configuration changes done on the network for PTP/Clock Synchronization. [PR1451950](#)
- After disabling 100G and 40G interface Laser output power in **show interfaces diagnostics optics** shows some values. [PR1452323](#)

- ACX5448 FPC crashed due to segmentation fault. [PR1453766](#)
- Incorrect operating state is displayed in snmp trap for fan removal. [PR1455577](#)
- Enable gigether option to configure Ethernet FEC on client ports. [PR1456293](#)
- ACX5448-D and ACX5448-M Devices does not display airflow information and temperature sensors as expected. [PR1456593](#)
- ACX5448 Layer2 VPN with encapsulation-type ethernet stops passing traffic after a random port is added with vlan configuration. [PR1456624](#)
- The rpd crash might be seen if BGP route is resolved over same prefix protocol next-hop in inet.3 table which has both RSVP and LDP routes. [PR1458595](#)
- Route resolve resolution is not happening when the packet size is 10000. [PR1458744](#)
- The traffic might be blackholed during link recovery in an open ethernet access ring with ERPS configured. [PR1459446](#)
- ACX5000: SNMP mib walk for jnxOperatingTemp not returning anything for FPC in new versions. [PR1460391](#)
- ACX5448-M Interfaces and Optics support: on enabling local loopback 10G interface is going down. [PR1460715](#)
- ACX5448-D Interfaces and Optics support: sometimes during the bring up of AE interface there are ARP resolution issues. [PR1461485](#)
- ACX Series routers LLDP neighbor not up on lag after software upgrade to Junos OS Release 18.2R3-S1. [PR1461831](#)
- RED drop on interface, no congestion. [PR1470619](#)

Layer 2 Ethernet Services

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

Platform and Infrastructure

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

Routing Protocols

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

SEE ALSO

What's New		15
What's Changed		20
Known Limitations		21
Open Issues		23
Documentation Updates		27
Migration, Upgrade, and Downgrade Instructions		28

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides](#) | [27](#)

This section lists the errata and changes in Junos OS Release 19.4R1 for the ACX Series documentation.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

What's New		15
What's Changed		20
Known Limitations		21
Open Issues		23
Resolved Issues		24
Migration, Upgrade, and Downgrade Instructions		28

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 28](#)

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

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

SEE ALSO

[What's New | 15](#)

[What's Changed | 20](#)

[Known Limitations | 21](#)

[Open Issues | 23](#)

[Resolved Issues | 24](#)

[Documentation Updates | 27](#)

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- [What's New | 29](#)
- [What's Changed | 40](#)
- [Known Limitations | 42](#)
- [Open Issues | 43](#)
- [Resolved Issues | 48](#)
- [Documentation Updates | 54](#)
- [Migration, Upgrade, and Downgrade Instructions | 55](#)

These release notes accompany Junos OS Release 19.4R1 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

- [Authentication, Authorization, and Accounting | 31](#)
- [Class of Service | 31](#)
- [EVPN | 31](#)
- [Junos OS XML, API, and Scripting | 33](#)
- [Junos Telemetry Interface | 34](#)

- Layer 2 Features | 35
- MPLS | 37
- Multicast | 37
- Operation, Administration, and Maintenance (OAM) | 37
- Port Security | 38
- Routing Policy and Firewall Filters | 39
- System Logging | 39
- System Management | 39
- User Interface and Configuration | 40

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 19.4R1: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

Authentication, Authorization, and Accounting

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

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

Class of Service

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

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

EVPN

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

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

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

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

- DHCPv4 and DHCPv6 snooping. [See [DHCP Snooping](#).]
- Dynamic ARP inspection (DAI). [See [Understanding and Using Dynamic ARP Inspection \(DAI\)](#).]
- Neighbor discovery inspection (NDI). [See [IPv6 Neighbor Discovery Inspection](#).]
- IPv4 and IPv6 source guard. [See [Understanding IP Source Guard for Port Security on Switches](#).]
- Router advertisement (RA) guard. [See [Understanding IPv6 Router Advertisement Guard](#).]

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

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

- MAC limiting, storm control, and port mirroring support in EVPN-VXLAN overlay networks.
[See [MAC Limiting, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.](#)]
- Multihomed proxy advertisement.
[See [EVPN Multihoming Implementation.](#)]
- Selective multicast forwarding and SMET route support in EVPN-VXLAN.
[See [Overview of Selective Multicast Forwarding.](#)]
- Standard class-of-service (CoS) features—classifiers, rewrite rules, and schedulers—are supported on VXLAN interfaces.
[See [Understanding CoS on OVSD-Managed VXLAN Interfaces.](#)]
- VMTO for ingress traffic.
[See [Ingress Virtual Machine Traffic Optimization.](#)]

Junos OS XML, API, and Scripting

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

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

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

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

- **ncclient** (0.6.4)
- **packaging** (19.0)
- **paho.mqtt** (1.4.0)
- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

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

Junos Telemetry Interface

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

The following Routing Engine statistics are supported:

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

The following Packet Forwarding Engine statistics are supported:

- Congestion and latency monitoring
- Logical interface
- Filter
- Physical interface
- NPU/LC memory
- Network discovery NDP table state

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

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

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

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

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

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

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

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

Layer 2 Features

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

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

- **Ethernet ring protection switching (ERPS)(EX4300-MP)**—Starting in Junos OS Release 19.4R1, the EX4300-MP supports Ethernet ring protection switching (ERPS) to reliably achieve carrier-class network requirements for Ethernet topologies forming a closed loop. The ITU-T Recommendation is G.8032 version 1.

ERPS version 1 comprises the following features:

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

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

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

ERPS version 1 comprises the following features:

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

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

MPLS

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

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

Multicast

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

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

Operation, Administration, and Maintenance (OAM)

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

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

- **Support for Ethernet CFM (EX4650)**—Starting with Junos OS Release 19.4R1, the EX4650 switch supports Ethernet connectivity fault management (CFM). You can use Ethernet CFM to:
 - Monitor faults, using the continuity check messages (CCMs) to discover and maintain adjacencies at the VLAN or link level.

- Discover paths and verify faults, using the Link Trace protocol to map the path taken to a destination MAC address.
- Isolate and troubleshoot faults, using loopback messages. .

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

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

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

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

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

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

Port Security

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

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

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

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

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

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

Routing Policy and Firewall Filters

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

[See [erac1-ip6-match](#) and [Configuring an Egress Filter Based on IPv6 Source or Destination IP Addresses](#).]

System Logging

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

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

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

System Management

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

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

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

- After a specified number of minutes (1 through 120)

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

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

User Interface and Configuration

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

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

SEE ALSO

What's Changed	 40
Known Limitations	 42
Open Issues	 43
Resolved Issues	 48
Documentation Updates	 54
Migration, Upgrade, and Downgrade Instructions	 55

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.4R1-S3](#) | 41
- [What's Changed in Release 19.4R1](#) | 41

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

What's Changed in Release 19.4R1-S3

General Routing

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

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

What's Changed in Release 19.4R1

General Routing

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

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

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

- **IPv6 address in the prefix TIEs displayed correctly (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.

Interfaces and Chassis

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

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

Junos Telemetry Interface

- **LLDP ON_CHANGE statistics support with JTI (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—Enhanced telemetry ON_CHANGE event support provides the following LLDP attributes:

- When LLDP is enabled on interfaces, LLDP interface counters are notified along with other interface-level attributes.
- ON_CHANGE event reports LLDP neighbor age and custom TLVs, as well as when a neighbor is initially discovered.

[See [Guidelines for gRPC and gNMI Sensors](#).]

Routing Protocols

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

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

System Logging

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

SEE ALSO

[What's New | 29](#)

[Known Limitations | 42](#)

[Open Issues | 43](#)

[Resolved Issues | 48](#)

[Documentation Updates | 54](#)

[Migration, Upgrade, and Downgrade Instructions | 55](#)

Known Limitations

There are no known limitations for the EX Series switches in Junos OS Release 19.4R1.

SEE ALSO

What's New		29
What's Changed		40
Open Issues		43
Resolved Issues		48
Documentation Updates		54
Migration, Upgrade, and Downgrade Instructions		55

Open Issues

IN THIS SECTION

- [Authentication and Access Control](#) | [44](#)
- [Class of Service \(CoS\)](#) | [44](#)
- [EVPN](#) | [44](#)
- [General Routing](#) | [44](#)
- [Infrastructure](#) | [46](#)
- [Interfaces and Chassis](#) | [46](#)
- [Junos Fusion Enterprise](#) | [46](#)
- [Junos Fusion Satellite Software](#) | [46](#)
- [Layer 2 Ethernet Services](#) | [46](#)
- [Layer 2 Features](#) | [47](#)
- [Platform and Infrastructure](#) | [47](#)
- [Routing Protocols](#) | [47](#)
- [User Interface and Configuration](#) | [48](#)

Learn about open issues in Junos OS Release 19.4R1 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 command **load ssl-certificate path PATHNAME**, now we have to configure path using CLI **set protocols dot1x ssl-certificate-path PATHNAME**, if pathname is not default path **/var/tmp/**. [PR1431086](#)

Class of Service (CoS)

- On EX4600 platforms, if **shaping-rate** is configured, the shaping feature might not work after a reboot. The service might be impacted as the traffic cannot be rate limited. [PR1432078](#)
- In Junos Fusion scenario, when traffic from aggregation device (AD) to satellite device (SD) is exported with different DSCP marking, it might be changed into network-control queue on extended port of SD. [PR1433252](#)

EVPN

- In an EVPN environment, proxy ARP and ARP suppression is enabled on the PE device by default for reducing the flooding of ARP packets; however, in the case of ARP probe packets used in the process of Duplicate Address Detection (DAD), the client might treat the IP address that it is in use as duplicated address after receiving the proxied packets from PE device. [PR1427109](#)

General Routing

- On EX Series or QFX Series switches, if the switch is power cycled, then some processes (such as jdhcp, lacp, and lldpd) might stop working after the switch reboots. [PR1222504](#)
- If Non-NEBS (Network Equipment-Building System, a design guideline applied to telecommunications equipment) compliant optics is used on MX MPC and chassis temperature exceeds non-nebs-optics-overheat-trigger (default: 50 degrees), the fan might not change to high speed because the temperature hasn't reached fan speed "High" threshold (default: 60 degrees). If the temperature remains over non-nebs-optics-overheat-trigger for about 10 minutes, the non-NEBS compliant optics might be disabled. [PR1331186](#)
- On an EX2300 switch, the output of the **show chassis routing-engine** command might display an incorrect value of **mac reset** for the last reboot reason field. [PR1331264](#)
- EX4300 virtual-chassis systems may fail to register some jnxOperating SNMP OIDs related to the routing-engines. This behavior is more likely if virtual-chassis members 0 and 1 (FPC0 and FPC1) are not selected as routing-engines. [PR1368845](#)
- Traffic flooding occurs instead of routing, when VRRP is scaled more than 150. [PR1371520](#)
- On the EX9208, a few xe-interfaces go down with the following error message:
if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error. [PR1377840](#)

- On EX Series platforms, the DHCP/PPP subscribers might fail to bind. The reason is that when installing new software images, the shared memory (created by the previously running image) might not to be cleared out. As a result, the previous values in the shared memory are removed and the daemons affected by the data in the shared memory might continue to crash and generate core files. [PR1396470](#)
- On the EX4650, uRPF check in strict mode might not work properly. [PR1417546](#)
- When the chassisd process receives incorrect values from LCMD for the RPM values, it changes the Fan status to "Failed" from "Ok", and vice versa. [PR1417839](#)
- On EX2300/EX3400 Virtual Chassis platform with storm-control enabled, when multiple filter entries get exhausted, VC becomes unstable and FXPC cores. [PR1422132](#)
- In Virtual Chassis scenario, when the interfaces flaps or VLAN configuration is changed frequently, the network topology will be changed accordingly, then CPU utilization will be dramatically increased to very high within a short time, which might cause the failure of essential communications between VC master and members. When the failure happens, FPC will automatically restart. As a result, VC is split and traffic is lost. [PR1427075](#)
- EX4300-48MP switch cannot learn MAC address through some access ports that are directly connected to a host when autonegotiation is used. [PR1430109](#)
- The time taken to install or delete IPv4 or IPv6 routes into the FIB is slowed down in Junos OS Release 19.3. Analysis shows that rpd learning rates are not degraded but RIB to FIB download rate is degraded. [PR1441737](#)
- On the EX9214 device, if the MACsec-enabled link flaps after reboot, the error `errorlib_set_error_log(): err_id(-1718026239)` is observed. [PR1448368](#)
- The sFlow sample packets might stop on one aggregated Ethernet member link if ingress sFlow is configured on the member link. This might cause inaccurate monitoring on the network traffic. [PR1449568](#)
- Traffic is dropped from SD to AD. The loss is intermediate and is not seen regularly. This occurs because few packets that are transmitted from the egress of AD1 is short of FCS of data. It is observed that the normal data packets are of size 128 bytes while the corrupted data packet is of size 122 byte. [PR1450373](#)
- 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](#)
- In VXLAN setup containing a large number of child interfaces, significant link up delay was seen when one of FPC in EX4600-40F rebooted. [PR1456336](#)
- When tunnel-services are configured on a PIC, the optics measurements that subscribed through gRPC might not be streamed. [PR1468435](#)
- On EX Series platforms, the shaping of CoS does not work after reboot when the shaping rate is configured with an absolute value. The line rate of traffic is sent out, no shaping occurs. This issue has traffic or service impact. [PR1472223](#)

Infrastructure

- When xSTP/RTG is not configured in the network and there is a traffic loop, after the network loop is broken, sometimes MAC address learning might not be learned. [PR473454](#)
- On EX3400 and EX2300 line of switches during ZTP with configuration and image upgrade with FTP as file transfer, image upgrade is successful but sometimes VM core files might be generated. [PR1377721](#)
- On EX Series platforms, when you configure a large number of firewall filters on some interfaces, the FPC crashes generating core files. [PR1434927](#)
- Packet Forwarding Engine sometimes does not come up after system reboot. Timeout is required to handle the fifo tx/rx error. Debug sysctls are been removed. [PR1454950](#)
- On EX4300-48MP device acting as a leaf in Layer 2 IP fabric EVPN VXLAN environment, the 100 percent traffic drop might be seen if unplug cable is connected to the et-interface and plug it back. [PR1463318](#)

Interfaces and Chassis

- When dynamic DHCP sessions exist in the device, if multiple commits in parallel are performed, the commit might hang up. [PR1470622](#)

Junos Fusion Enterprise

- In a Junos Fusion Enterprise environment, when traffic originates from a peer device connected to the aggregation device and the ICL is a LAG, there might be a reachability issue if the cascade port is disabled and traffic has to flow through the ICL LAG to reach the satellite device. As a workaround, use single interface as the ICL instead of a LAG. [PR1447873](#)

Junos Fusion Satellite Software

- In Junos Fusion dpd might crash on satellite devices running SNOS. [PR1460607](#)

Layer 2 Ethernet Services

- The jdhcpd_era log files constantly consume 121M of space out of 170M, resulting into file system full and traffic impact. [PR1431201](#)
- In EVPN multihomed ACTIVE-ACTIVE scenario when LACP is enable on PE-CE child member links and after recovering from a core-isolation on PE device, the PE-CE child member links might be stuck in DETACHED state if LACP sync-reset feature is enabled on CE device. The child links on the CE device might show LACP state as "Collecting Distributing", but on the PE device the LACP state might be "DETACHED". [PR1463791](#)

Layer 2 Features

- On EX4600 platforms, if copper base SFP-T is used, it might not get up on physical layer and the MAC/ARP learning might not work if it gets up. [PR1437577](#)

Platform and Infrastructure

- The **commit synchronize** command fails because the kernel socket gets stuck. [PR1177692](#)
- On EX4300-32F platform, when SFP-T is used in a port earlier and SFP is inserted and then removed from the same port, the pfex process might crash and core dump. [PR1421257](#)
- On all Junos OS platform, in some rare conditions, there might be packet drops, replication failures or ksynchd crashes on the logical system. This issue might appear at the time of Routing Engine switchover if the system is running for a long time and lot of configuration changes have been made over the time. [PR1427842](#)
- On EX9208, the traffic loss is observed if ingress and egress ports are in different FPC. [PR1429714](#)
- On EX4300 platform, if FBF filters are applied on IRB with LAG configuration also existing on the box, the firewall filters cannot be created and function correctly due to TCAM programming issues. [PR1447012](#)
- On EX9208, 33 percent degradation with MAC learning rate in Junos OS Release 19.3R1 while comparing with Junos OS Release 18.4R1. [PR1450729](#)
- In the Virtual Chassis scenario, the IRB traffic might get dropped after master switchover. [PR1453025](#)
- On EX4300 platforms in IGMP snooping scenario, if both SP/EP styles of Ethernet encapsulations are used in one VLAN, IGMP reports are not forwarded from group member interfaces towards multicast routers interface in the VLAN. [PR1466075](#)

Routing Protocols

- On EX4300 and EX4600 Series switches, if host destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log/syslog action (for example, 'filter <> term <> then log/syslog'), such packets should not be dropped and reach the Routing Engine. [PR1379718](#)
- BGP IPv4 or IPv6 convergence and RIB install/delete time is degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- On EX4600 platforms, the received traffic will be dropped if the destination UDP port is 520/521 though the device runs pure layer 2 switching. [PR1429543](#)
- On EX4600 with service provider (SP) style VLAN configuration (in this method, each VLAN-ID is locally significant to a physical interface), if **interface-mac-limit** or **mac-table-size** is configured (that is, software MAC learning is enabled) and the scale of MAC addresses on the box is more than 2000, traffic might be dropped after QinQ enabled interface is flapped or a change is made to the **vlan-id-list**. [PR1441402](#)

User Interface and Configuration

- On EX Series switches, while checking add or delete interface-range configuration, **couldn't find end-range in deleted tree for member-range** message is logged. [PR1357574](#)

SEE ALSO

What's New	 29
What's Changed	 40
Known Limitations	 42
Resolved Issues	 48
Documentation Updates	 54
Migration, Upgrade, and Downgrade Instructions	 55

Resolved Issues

IN THIS SECTION

- [Authentication and Access Control](#) | [49](#)
- [EVPN](#) | [49](#)
- [General Routing](#) | [49](#)
- [Infrastructure](#) | [51](#)
- [Interfaces and Chassis](#) | [52](#)
- [Junos Fusion Enterprise](#) | [52](#)
- [Junos Fusion Satellite Software](#) | [52](#)
- [J-Web](#) | [52](#)
- [Layer 2 Ethernet Services](#) | [52](#)
- [Layer 2 Features](#) | [52](#)
- [Platform and Infrastructure](#) | [53](#)
- [Routing Protocols](#) | [53](#)
- [User Interface and Configuration](#) | [54](#)
- [Virtual Chassis](#) | [54](#)

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

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

Authentication and Access Control

- After rebooting the preloaded box, the SSL certificate is not displayed. [PR1431086](#)

EVPN

- In EVPN scenario, the IRB logical interface might not come up when the local Layer 2 interface is down. [PR1436207](#)
- ARP request or Neighbor Solicitation (NS) message might be sent back to the local segment by the DF router. [PR1459830](#)
- The rpd might crash after changing EVPN related configuration. [PR1467309](#)

General Routing

- On the EX3400, when me0 ports are connected between two EX3400 switches, the link does not come up. [PR1351757](#)
- Transit OSPF traffic over Q-in-Q tunneling might be dropped if a firewall filter is applied to the LoO interface. [PR1355111](#)
- The l2ald process might crash and generate a core file on EX2300 Virtual Chassis when converted a trunk port is converted to a dot1x access port with tagged traffic flowing. [PR1362587](#)
- The interface on the failed member FPC of EX2300 and EX3400 Virtual Chassis might stay up for 120 seconds. [PR1422507](#)
- IPv6 multicast traffic received on one Virtual Chassis member might be dropped when egressing on another Virtual Chassis member if MLD snooping is enabled. [PR1423310](#)
- MAC addresses overlaps between different switches. [PR1425123](#)
- The delay in transmission of BPDUs after GRES might result in loss of traffic on EX2300 and EX3400 Virtual Chassis. [PR1428935](#)
- Erroneous log messages and chassis environment output related to the fan tray in EX4300MP and EX4300-48P Virtual Chassis. [PR1431263](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- Packet drop might be seen if native VLAN is configured along with flexible VLAN tagging. [PR1434646](#)
- Micro BFD session might flap upon inserting a QSFP to other port. [PR1435221](#)

- The mc-ae interface might get stuck in waiting state in a dual mc-ae scenario. [PR1435874](#)
- Commit check error for VSTP on the EX9200 line of switches **xSTP:Trying to configure too many interfaces for given protocol**. [PR1438195](#)
- LED turns on even after the Virtual Chassis 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 rpd process might generate a core file when the router boots up because of a file pointer issue because there are two code paths that can close the file. [PR1438597](#)
- The dot1x might not work when **captive-port** is also configured on the interface on the backup or nonmaster FPC. [PR1439200](#)
- DHCPv6 relay binding is not up while verifying DHCP snooping along with DHCPv6 relay. [PR1439844](#)
- EX4600 Virtual Chassis does not come up after replacing the Virtual Chassis port from fiber connection to DAC cable. [PR1440062](#)
- CPU might hang or interface might get stuck on a particular 100-Gigabit Ethernet 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 isolated VLAN might not get IP addresses after completing authentication when both **dhcp-security** and **dot1x** are configured. [PR1442078](#)
- On the EX3400, the fan alarm **Fan X not spinning** appears and disappears repeatedly after the fan tray is removed. [PR1442134](#)
- The rpd might crash when the BGP sends a notification message. [PR1442786](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- Non-designated port does not move to the backup port role. [PR1443489](#)
- The **/var/host/motd does not exist** message is flooded every 5 seconds in chassisd logs. [PR1444903](#)
- On the EX4300-MP, the following log messages is generated continuously: **rpdd[6550]: task_connect: task AGENTD I/O.128.0.0.1+9500 addr 128.0.0.1+9500: Connection refused**. [PR1445618](#)
- On the EX3400 dot1xd core file is found at **macsec_update_intf macsec_destroy_ca**. [PR1445764](#)
- Major alarm log messages for temperature conditions are generated for the EX4600 at 56 degrees Celsius. [PR1446363](#)
- Traffic might be dropped when a firewall filter rule uses 'then vlan' as the action in a Virtual Chassis scenario. [PR1446844](#)
- The phone-home feature might fail on EX3400 switches because sysctl cannot read the device serial number. [PR1447291](#)
- On EX3400, Virtual Chassis might hang when a disk error occurs. [PR1447853](#)

- Unicast ARP requests do not receive a reply with the **no-arp-trap** option. [PR1448071](#)
- On EX3400, IPv6 routes received through BGP do not show the correct age time. [PR1449305](#)
- Except one aggregated Ethernet member link, the other links do not send out sFlow sample packets for ingress traffic. [PR1449568](#)
- DHCP snooping static binding does not take effect after deleting and readding the entries. [PR1451688](#)
- The l2ald and eventd processes are hogging 100 percent after issuing the **clear ethernet-switching table** command. [PR1452738](#)
- Configuration change in the **VLAN all** option might affect the **per-VLAN** configuration. [PR1453505](#)
- Version compare in PHC might fail and the same image might be downloaded. [PR1453535](#)
- Packet drops might be seen after removing and reinserting the SFP transceiver of the 40-Gigabit Ethernet uplink module ports. [PR1456039](#)
- Syslog message **Timeout connecting to peer database-replication** is generated when the command **show version detail** is issued. [PR1457284](#)
- SNMP trap messages are generated after an upgrade even though the temperature is 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 because the BGP IPv6 session flaps. [PR1459759](#)
- Storage space limitation leads to image installation failure when the phone-home client is used on EX2300 and EX3400 devices. [PR1460087](#)
- Configure any combination of VLANs and interfaces under VSTP/MSTP might cause VSTP/MSTP related configuration cannot be committed. [PR1463251](#)
- The Virtual Chassis function might brake after an upgrade on EX2300 and EX3400 devices. [PR1463635](#)
- On the EX2300, FXPC core file is generated after mastership election based on user priority. [PR1465526](#)

Infrastructure

- The operations on the console might not work if the **system ports console log-out-on-disconnect** statement is configured. [PR1433224](#)
- The recovery snapshot cannot be created after system zeroization. [PR1439189](#)
- On EX4300 CLI configuration **on-disk-failure** is not supported in Junos OS Release 18.2R3-S2. [PR1450093](#)
- Certain EX Series platforms might generate VM core files by panic and reboot. [PR1456668](#)

- Error messages related to soft reset of port because the queue buffers are stuck might be seen on EX4600-EX4300 VC. [PR1462106](#)
- The traffic is dropped on EX4300-48MP device acting as a leaf in Layer 2 IP fabric EVPN VXLAN environment. [PR1463318](#)

Interfaces and Chassis

- VRRP-V6 state flaps with init and idle states after configuring **vlan-tagging**. [PR1445370](#)
- The traffic might be forwarded to incorrect interfaces in MC-LAG scenario. [PR1465077](#)

Junos Fusion Enterprise

- Reachability issue of the host connected to the SD might be affected in a Junos Fusion Enterprise environment with EX9200 Series devices as AD. [PR1447873](#)

Junos Fusion Satellite Software

- The dpd might crash on satellite devices in a Junos Fusion Enterprise environment. [PR1460607](#)

J-Web

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

Layer 2 Ethernet Services

- The jdhcpd_era log files constantly consume 121M of space out of 170M, resulting in a full file system and affecting traffic. [PR1431201](#)
- DHCP request might get dropped in DHCP relay scenario. [PR1435039](#)
- On EX9200, the DHCP relay strips the 'GIADDR' field in messages towards the DHCP clients. [PR1443516](#)

Layer 2 Features

- Ethernet Ring Protection Switching (ERPS) nodes might not converge to IDLE state after failure recovery or reboot. [PR1431262](#)
- The MAC/ARP learning might not work for copper base SFP-T on EX4600. [PR1437577](#)
- The fxpc core files might be generated when committing the configuration. [PR1467763](#)

Platform and Infrastructure

- LACP DDoS policer is incorrectly triggered by other protocols traffic on all EX92XX Series platforms. [PR1409626](#)
- Over temperature SNMP trap is generated incorrectly for LC (EX4300-48P) based on master Routing Engine (EX4300-48MP) temperature threshold value. [PR1419300](#)
- Packet drops, replication failure, or ksyncd crash might be seen on the logical system of a Junos OS device after Routing Engine switchover. [PR1427842](#)
- IPv6 traffic might be dropped when static /64 IPv6 routes are configured. [PR1427866](#)
- Unicast ARP requests are not replied to with **no-arp-trap** option. [PR1429964](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- The FPC/pfec might crash due to DMA buffer leaking. [PR1436642](#)
- The laser TX might be enabled while the interface is disabled. [PR1445626](#)
- The PoE might not work after upgrading the PoE firmware on EX4300 platforms. [PR1446915](#)
- The firewall filters might not be created due to TCAM issues. [PR1447012](#)
- NSSU causes traffic loss again after the backup to master transitions. [PR1448607](#)
- On certain MPC line cards, cm errors need to be reclassified. [PR1449427](#)
- The REST service might become nonresponsive when the REST API receives several continuous HTTP requests. [PR1449987](#)
- The traffic for some VLANs might not be forwarded when **vlan-id-list** is configured. [PR1456879](#)
- ERP might not revert to idle state after reload or reboot of multiple switches. [PR1461434](#)

Routing Protocols

- Host-destined packets with filter log action might not reach to the Routing Engine if log/syslog is enabled. [PR1379718](#)
- On EX9208, BGP IPv4/IPv6 convergence and RIB install/delete time is degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The fxpc core file might be generated during the reboot of EX4600 switches. [PR1432023](#)
- Error message **RPD_DYN_CFG_GET_PROF_NAME_FAILED: Get profile name for session XXX failed: -7** might be seen in syslog after restarting the routing daemon. [PR1439514](#)
- Traffic might be dropped after the Q-in-Q enabled interface flaps 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 flaps after a certain time. [PR1448649](#)
- Loopback address exported into other VRF instance might not work on EX Series platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even when the LDP neighbor's MAC changes. [PR1451217](#)
- Changing "other querier present interval" timer is not working on IGMP/MLD snooping device in the existing bridge domain (BD) or listener domain (LD). [PR1461590](#)

User Interface and Configuration

- EX4600 switches are unable to commit baseline configuration after zeroization. [PR1426341](#)
- Problem with access to J-Web after updating from Junos OS Release 18.2R2 to 18.2R3. [PR1454150](#)

Virtual Chassis

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

SEE ALSO

What's New	 29
What's Changed	 40
Known Limitations	 42
Open Issues	 43
Documentation Updates	 54
Migration, Upgrade, and Downgrade Instructions	 55

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides](#) | 55

This section lists the errata and changes in Junos OS Release 19.4R1 for the EX Series switches documentation.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

What's New	 29
What's Changed	 40
Known Limitations	 42
Open Issues	 43
Resolved Issues	 48
Migration, Upgrade, and Downgrade Instructions	 55

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 55

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 | 29](#)

[What's Changed | 40](#)

[Known Limitations | 42](#)

[Open Issues | 43](#)

[Resolved Issues | 48](#)

[Documentation Updates | 54](#)

Junos OS Release Notes for JRR Series

IN THIS SECTION

- [What's New | 57](#)
- [What's Changed | 58](#)
- [Known Limitations | 58](#)
- [Open Issues | 59](#)
- [Resolved Issues | 59](#)
- [Documentation Updates | 60](#)
- [Migration, Upgrade, and Downgrade Instructions | 61](#)

These release notes accompany Junos OS Release 19.4R1 for JRR Series. They describe new and changed features, limitations, and known and resolved problems in the hardware.

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

What's New

IN THIS SECTION

- [Hardware](#) | 57

Learn about new features introduced in Junos OS Release 19.4R1 for JRR Series.

Hardware

- **JRR200 Route Reflector**—Starting with Junos OS Release 19.4R1, JRR200 Route Reflector a 1U form factor appliance with a multicore x86 CPU and preinstalled vRR software that can host one route reflector instance is available. JRR200 is suitable for large enterprises, data centers and service providers for hosting vRR software to scale up to 30 million routing information base (RIB) entries.

The JRR200 route reflector comes with eight 1/10 Gigabit Ethernet SFP+ ports, 64 GB of DDR4 memory, and two 240 GB solid-state drives (SSDs) in a RAID1 configuration. It is available in both AC and DC models which support Zero Touch Provisioning mode (ZTP) to ensure seamless insertion into the network and provide operational simplicity.

[See [JRR200 Route Reflector Hardware Guide](#) and [JRR200 Route Reflector Quick Start](#)]

- **ZTP Support for JRR200 Route Reflector**—Starting in Junos OS Release 19.4R1, ZTP can automate the provisioning of the device configuration and software image on JRR200 Route Reflector. ZTP supports self image upgrades and automatic configuration updates using ZTP DHCP options. In this release, ZTP supports revenue ports em2 thru em9, in addition to management port em0 which is supported in Junos OS Releases before 19.4R1.

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

SEE ALSO

| [What's Changed](#) | 58

Known Limitations	58
Open Issues	59
Resolved Issues	59
Documentation Updates	60
Migration, Upgrade, and Downgrade Instructions	61

What's Changed

There are no changes in behavior and syntax for JRR Series in Junos OS Release 19.4R1.

SEE ALSO

What's New	57
Known Limitations	58
Open Issues	59
Resolved Issues	59
Documentation Updates	60
Migration, Upgrade, and Downgrade Instructions	61

Known Limitations

There are no known limitations for JRR Series in Junos OS Release 19.4R1.

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	57
What's Changed	58
Open Issues	59
Resolved Issues	59
Documentation Updates	60

Open Issues

IN THIS SECTION

- [General Routing | 59](#)

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

General Routing

- USB install image does not work for JRR200 platform. [PR1471986](#)

SEE ALSO

What's New 57
What's Changed 58
Known Limitations 58
Resolved Issues 59
Documentation Updates 60
Migration, Upgrade, and Downgrade Instructions 61

Resolved Issues

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

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 57
What's Changed 58
Known Limitations 58
Open Issues 59
Documentation Updates 60
Migration, Upgrade, and Downgrade Instructions 61

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 60](#)

This section lists the errata and changes in Junos OS Release 19.4R1 documentation for JRR Series.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

What's New 57
What's Changed 58
Known Limitations 58
Open Issues 59
Resolved Issues 59
Migration, Upgrade, and Downgrade Instructions 61

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 61

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

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

SEE ALSO

[What's New](#) | 57

[What's Changed](#) | 58

[Known Limitations](#) | 58

[Resolved Issues | 59](#)

[Open Issues | 59](#)

[Documentation Updates | 60](#)

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- [What's New | 62](#)
- [What's Changed | 63](#)
- [Known Limitations | 63](#)
- [Open Issues | 64](#)
- [Resolved Issues | 65](#)
- [Documentation Updates | 65](#)
- [Migration, Upgrade, and Downgrade Instructions | 66](#)

These release notes accompany Junos OS Release 19.4R1 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

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

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

What's New

There are no new features or enhancements to existing features in Junos OS Release 19.4R1 for Junos fusion for enterprise.

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

SEE ALSO

What's Changed	 63
Known Limitations	 63
Open Issues	 64
Resolved Issues	 65
Documentation Updates	 65
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 19.4R1 for Junos fusion for enterprise.

SEE ALSO

What's New	 62
Known Limitations	 63
Open Issues	 64
Resolved Issues	 65
Documentation Updates	 65
Migration, Upgrade, and Downgrade Instructions	 66

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 19.4R1 for Junos fusion for enterprise.

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

SEE ALSO

[What's New | 62](#)

[What's Changed | 63](#)

[Open Issues | 64](#)

[Resolved Issues | 65](#)

[Documentation Updates | 65](#)

[Migration, Upgrade, and Downgrade Instructions | 66](#)

Open Issues

IN THIS SECTION

- [Junos Fusion for Enterprise | 64](#)

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

Junos Fusion for Enterprise

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

SEE ALSO

What's New	 62
What's Changed	 63
Known Limitations	 63
Open Issues	 64
Documentation Updates	 65
Migration, Upgrade, and Downgrade Instructions	 66

Resolved Issues

This section lists the issues fixed in Junos OS Release 19.4R1.

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	 62
What's Changed	 63
Known Limitations	 63
Open Issues	 64
Documentation Updates	 65
Migration, Upgrade, and Downgrade Instructions	 66

Documentation Updates

This section lists the errata and changes in Junos OS Release 19.4R1 documentation for Junos Fusion for enterprise.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

[What's New](#) | 62

[What's Changed](#) | 63

[Known Limitations](#) | 63

[Open Issues](#) | 64

[Resolved Issues](#) | 65

[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 Enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the **junos-install** package. Use other packages (such as the **jbundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **junos-install** package and details of the installation process, see the [Installation and Upgrade Guide](#).

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

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

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

To download and install Junos OS:

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

9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **junos-install** package on the aggregation device.

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

Customers in the United States and Canada, use the following commands:

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

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

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

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname** (available only for Canada and U.S. version)

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Switch

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

Downgrading from Junos OS

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

NOTE: You cannot downgrade more than three releases.

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

To downgrade a Junos Fusion Enterprise from Junos OS Release 19.4R1, follow the procedure for upgrading, but replace the 19.2 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

[What's New | 62](#)

[What's Changed | 63](#)

[Known Limitations | 63](#)

[Open Issues | 64](#)

[Resolved Issues | 65](#)

[Documentation Updates | 65](#)

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- [What's New | 72](#)
- [What's Changed | 73](#)
- [Known Limitations | 73](#)
- [Open Issues | 74](#)
- [Resolved Issues | 74](#)
- [Documentation Updates | 75](#)
- [Migration, Upgrade, and Downgrade Instructions | 76](#)

These release notes accompany Junos OS Release 19.4R1 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.4R1.

SEE ALSO

- | |
|--|
| What's Changed 73 |
| Known Limitations 73 |
| Open Issues 74 |
| Resolved Issues 74 |
| Documentation Updates 75 |

What's Changed

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 19.4R1 for Junos fusion for provider edge.

SEE ALSO

What's New 72
Known Limitations 73
Open Issues 74
Resolved Issues 74
Documentation Updates 75
Migration, Upgrade, and Downgrade Instructions 76

Known Limitations

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 73
Open Issues 74
Resolved Issues 74
Documentation Updates 75
Migration, Upgrade, and Downgrade Instructions 76

Open Issues

IN THIS SECTION

- [Junos Fusion for Provider Edge | 74](#)

Learn about open issues in this release for Junos Fusion Provider Edge.

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

Junos Fusion for Provider Edge

- The sdpd process might continuously crash if there are more than 12 cascade-ports configured to a satellite device. [PR1437387](#)
- The aggregated Ethernet interface might flap whenever a new logical interface is added. [PR1441869](#)

SEE ALSO

[What's New | 72](#)

[What's Changed | 73](#)

[Known Limitations | 73](#)

[Resolved Issues | 74](#)

[Documentation Updates | 75](#)

[Migration, Upgrade, and Downgrade Instructions | 76](#)

Resolved Issues

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

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

SEE ALSO

What's New		72
What's Changed		73
Known Limitations		73
Open Issues		74
Documentation Updates		75
Migration, Upgrade, and Downgrade Instructions		76

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides](#) | [75](#)

This section lists the errata and changes in Junos OS Release 19.4R1 for Junos Fusion Provider Edge.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

What's New		72
What's Changed		73
Known Limitations		73
Open Issues		74
Resolved Issues		74
Migration, Upgrade, and Downgrade Instructions		76

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 76
- Upgrading an Aggregation Device with Redundant Routing Engines | 79
- Preparing the Switch for Satellite Device Conversion | 79
- Converting a Satellite Device to a Standalone Device | 81
- Upgrading an Aggregation Device | 83
- Upgrade and Downgrade Support Policy for Junos OS Releases | 83
- Downgrading from Junos OS Release 19.4 | 84

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 19.4R1 is different from that for earlier Junos OS releases.

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

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

Customers in the United States and Canada, use the following commands.

- For 64-bit software:

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

```
user@host> request system software add validate reboot  
source/jinstall64-19.4R1.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-19.4R1.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

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

```
user@host> request system software add validate reboot
source/jinstall64-19.4R1.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-19.4R1.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 19.4R1 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see [Understanding Junos Fusion Software and Hardware Requirements](#)

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

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

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

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

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

8. Remove the satellite device from the automatic satellite conversion configuration.

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

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

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

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

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

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

```
[edit]
```

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

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

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

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

Upgrading an Aggregation Device

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

Downgrading from Junos OS Release 19.4

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

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

What's New		72
What's Changed		73
Known Limitations		73
Open Issues		74
Resolved Issues		74
Documentation Updates		75

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

IN THIS SECTION

- [What's New](#) | [85](#)
- [What's Changed](#) | [106](#)
- [Known Limitations](#) | [111](#)
- [Open Issues](#) | [114](#)
- [Resolved Issues](#) | [131](#)
- [Documentation Updates](#) | [152](#)
- [Migration, Upgrade, and Downgrade Instructions](#) | [153](#)

These release notes accompany Junos OS Release 19.4R1 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

- Hardware | 86
- Class of Service | 89
- EVPN | 89
- Forwarding and Sampling | 90
- General Routing | 91
- High Availability (HA) and Resiliency | 91
- Interfaces and Chassis | 91
- Junos OS, XML, API, and Scripting | 94
- Junos Telemetry Interface | 95
- Layer 2 Features | 97
- Layer 2 VPN | 97
- MPLS | 97
- Multicast | 98
- Network Management and Monitoring | 99
- OAM | 99
- Routing Policy and Firewall Filters | 99
- Routing Protocols | 100
- Services Applications | 103
- Software-Defined Networking | 103
- Software Licensing | 104
- Subscriber Management and Services | 104
- System Logging | 105

Learn about new features introduced in Junos OS Release 19.4R1 for MX Series routers.

Hardware

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

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

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

[Table 1 on page 87](#) summarizes the Next Gen Services supported in this release.

Table 1: Next Gen Services Summary

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

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

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

The MX-SPC3 services card is compatible end-to-end with the MX Series Switch Fabrics, Routing Engines and MS-MPC line cards. See [Table 2 on page 88](#):

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

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

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

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

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

Class of Service

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

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

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

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

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

EVPN

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

- Ethernet segment identifier
- MAC address on EVPN Type 2 routes
- IP address on EVPN Type 2 and EVPN Type 5 routes
- Extended community

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

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

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

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

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

Forwarding and Sampling

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

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

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

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

is active. The improved failover and traffic rebalancing introduced in this release minimize the traffic redistribution when multiple servers in the ECMP path fail.

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

General Routing

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

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

High Availability (HA) and Resiliency

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

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

Interfaces and Chassis

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

The transceivers encapsulate PDH (E1 or T1 or DS3) packets as Ethernet frames while transporting legacy time division multiplexing (TDM) traffic over packet switched networks (PSNs). At the receiver

end of the emulated circuit, another smart transceiver, paired with the first one and preconfigured to carry packets that are in the same multicast MAC address group, de-encapsulates the Ethernet frames, rebuilds the TDM data stream, and forwards it onto the local TDM interface.

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

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

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

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

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

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

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

- MPC7E (MPC7E-MRATE and MPC7E-10G (non-MACsec mode))
- MPC8E (MIC-MRATE MICs)
- MPC9E (MIC-MRATE MICs)

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

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

- **Optimize fabric path to prevent traffic hop (MX2008, MX2010, and MX2020 with MPC9E)**—Starting in Junos OS Release 19.4R1, you can optimize the fabric path of the traffic flowing over abstracted fabric

(af) interfaces between two guest network functions (GNFs) by configuring a fabric optimization mode. This feature reduces fabric bandwidth consumption by preventing any additional fabric hop (switching of traffic flows from one Packet Forwarding Engine to another because of abstracted fabric interface load balancing) before the packets eventually reach the destination Packet Forwarding Engine.

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

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

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

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

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

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

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

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

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

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

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

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

[See [Alarm Overview](#).]

Junos OS, XML, API, and Scripting

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

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

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

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

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

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

- **Support for 64-bit architecture added for use of management interface in a nondefault routing instance in op scripts and JET applications (MX Series)**—Junos OS Release 19.4R1 supports 64-bit architecture for Junos operating scripts and on-box JET applications being able to use the function

`set_routing_instance()` to program the protocol software (TCP/UDP) to use a nondefault routing instance instead of the default management routing interface.

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

Junos Telemetry Interface

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

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

- `/components/components/transceiver/`

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

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

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

- `/interfaces/interface/ethernet/state/counters/in-8021q-frames`
- `/interfaces/interface/ethernet/state/counters/in-crc-errors`
- `/interfaces/interface/ethernet/state/counters/in-block-errors`
- `/interfaces/interface/ethernet/state/counters/out-mac-control-frames`
- `/interfaces/interface/ethernet/state/counters/out-mac-pause-frames`

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

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

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

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

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

Layer 2 Features

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

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

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

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

Layer 2 VPN

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

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

MPLS

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

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

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

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

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

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

- **Support for static adjacency segment identifier for aggregated Ethernet member links on MPC10E-15C-MRATE line cards (MX240, MX480, and MX960)**—Starting with Junos OS Release 19.4R1, you can configure a transit single-hop static label-switched path (LSP) for a specific member link of an aggregated Ethernet (ae) interface. The label for this route comes from the segment routing local block

(SRLB) pool of the configured static label range. Configure the ae member interface name using the **member-interface** statement option at the **[edit protocols mpls static-label-switched-path *name* transit *name*]** hierarchy level. This feature is supported for ae interfaces only.

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

Multicast

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

[See [Multicast Overview](#).]

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

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

Network Management and Monitoring

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

OAM

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

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

Routing Policy and Firewall Filters

- **Support for firewall forwarding on MPC10E line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.4R1, the following traffic policers are fully supported on MX240, MX480, and MX960 routers with MPC10E line cards:
 - GRE tunnels, including encapsulation (**family any**), decapsulation, GRE-in-UDP over IPv6, and the following sub-options: sample, forwarding class, interface group, and no-ttl-decrement
 - Input and output filter chains
 - Actions, including policy-map filters, do-not-fragment, and prefix
 - Layer 2 policers
 - Policer overhead adjustment
 - Hierarchical policers
 - Shared bandwidth
 - Percentages
 - Logical interfaces

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

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

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

Routing Protocols

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

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

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

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

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

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

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

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

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

- **Support for BGP RIB Sharding (MX Series and VRR)**—Starting in Junos OS Release 19.4R1, the BGP process is split into different threads so that they can run concurrently on a multicore routing engine through RIB sharding which results in reduced convergence time and faster performance. BGP RIB sharding splits a BGP RIB into several sub RIBs and each sub RIB handles a subset of BGP routes. Each sub RIB is served by a separate RPD thread to achieve parallel processing.

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

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

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

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

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

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

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

[See [Understanding VRRP.](#)]

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

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

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

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

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

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

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

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

NOTE:

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

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

- **Support for BGP PIC Edge with BGP labeled unicast (MX Series and PTX Series)**—Starting with Junos OS Release 19.4R1, MX Series and PTX Series routers support BGP PIC Edge with BGP labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect

traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

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

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

Services Applications

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

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

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

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

[dns-filter](#)

Software-Defined Networking

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

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

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

- **Tunnel templates for PCE-initiated segment routing LSPs (MX Series)**—Starting in Junos OS Release 19.4R1, you can configure a tunnel template for Path Computation Element (PCE)-initiated segment routing LSPs and apply it through policy configuration. These templates enable dynamic creation of

segment routing tunnels with two additional parameters – Bidirectional forwarding detection (BFD) and LDP tunneling.

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

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

Software Licensing

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

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

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

Subscriber Management and Services

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

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

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

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

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

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

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

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

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

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

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

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

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

System Logging

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

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

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

SEE ALSO

[What's Changed | 106](#)

[Known Limitations | 111](#)

[Open Issues | 114](#)

[Resolved Issues | 131](#)

[Documentation Updates | 152](#)

[Migration, Upgrade, and Downgrade Instructions | 153](#)

What's Changed

IN THIS SECTION

- [General Routing | 107](#)
- [Interfaces and Chassis | 107](#)
- [Junos Telemetry Interface | 108](#)
- [MPLS | 109](#)
- [Network Management and Monitoring | 109](#)
- [Routing Protocols | 109](#)
- [Services Applications | 109](#)
- [Software-Defined Networking | 109](#)
- [Subscriber Management and Services | 110](#)
- [System Logging | 111](#)

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

General Routing

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

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

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

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

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

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

- **Automatic installation of YANG-based CLI for RIFT protocol (MX Series, QFX Series, and vMX with 64-bit and x86-based servers)**—In RIFT 1.2 Release, installation of the CLI for RIFT protocol occurs automatically along with the installation of the junos-rift package. In the pre-1.0 releases of the junos-rift package, the RIFT CLI had to be installed separately using **request system yang** command after installation of the junos-rift package.
- **IPv6 address in the prefix TIEs displayed correctly (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.
- **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

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

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

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

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

Junos Telemetry Interface

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

[See [Guidelines for gRPC and gNMI Sensors.](#)]

MPLS

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

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

Network Management and Monitoring

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

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

Routing Protocols

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

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

Services Applications

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

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

Software-Defined Networking

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

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

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

Subscriber Management and Services

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

- The Routing Engines have different amounts of DRAM.
- A 64-bit Junos OS image is loaded on the chassis.

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

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

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

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

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

System Logging

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

SEE ALSO

[What's New | 85](#)

[Known Limitations | 111](#)

[Open Issues | 114](#)

[Resolved Issues | 131](#)

[Documentation Updates | 152](#)

[Migration, Upgrade, and Downgrade Instructions | 153](#)

Known Limitations

IN THIS SECTION

- [General Routing | 112](#)
- [Interfaces and Chassis | 113](#)
- [MPLS | 113](#)
- [Platform and Infrastructure | 113](#)
- [Routing Protocols | 113](#)

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

General Routing

- First packet pertaining to J-Flow Packet Forwarding Engine sensor in UDP mode is missing after line card reboot on PORTER-R platform. [PR1344755](#)
- Traffic on GRE interface both ingress and egress cannot be Layer 2 mirrored. [PR1462375](#)
- Applying and removal of 1G speed results in channel being down. [PR1456105](#)
- The control peer PFCP heartbeat request time out window must be greater than 90 seconds. [PR1459135](#)
- Load balancing does not work as expected when tested with NAPT44 case twice. [PR1477670](#)
- afd hogged on executing clear VPLS table and MACs are not learned for less than 5 minutes. [PR1473334](#)
- If MTU is configured to a value higher than 9500, which is the maximum permissible value, configuration is done successfully. However, the actual value is set back to 1518 without any error. [PR1372690](#)
- The MIC-MACSEC-20G supports 10-Gigabit speed through the **set chassis fpc x pic y pic-mode 10G** configuration that is applied to both the PICs in that MIC. Any other PIC mode configuration must be removed and the 10-Gigabit PIC mode configuration must be applied. [PR1374680](#)
- In USF and non-USF cases, the monitor interface is MS (or) VMS interface. When chassisd restarts, all FPCs restart. SRD also restarts and ICCP connection goes down. If the FPC hosting the ICL goes down first before SRD receives physical interface IFD down for the monitored interface, the switchover does not happen immediately. [PR1416064](#)
- JSD generates core files when aggressively subscribing and unsubscribing both gRPC and gNMI subscriptions from multiple sessions. [PR1433744](#)
- The SPC3 cards are not supported with RE-2000. Even if the RE-2000 is the backup RE. [PR1435790](#)
- In a large-scale setup such as large number of routing-instances or interfaces, if there are frequent changes in the configuration and interface flapping when the rpd restarts through deactivate or activate of logical-system or restart routing, the rpd might crash. [PR1438049](#)
- Whenever the primary path goes down for the SRTE-tunnel, dynamic tunnel module (DTM) starts an expiry timer of 15 minutes. If the primary path comes up within this timer period, the tunnel comes up again. After the timer expires and the primary path is still not up, DTM asks SR-TE to remove the tunnel. Also, if there are multiple paths to reach the tunnel endpoint, BGP routes resolve over the other route, for example L-ISIS path. Later, even if the primary path comes up, BGP routes remain resolved over the other secondary route and do not change. No re-resolution happens because the SRTE-tunnel is being resolved with more than one indirection. For example, SR-TE over MPLS over IS-IS in this case. The same issue occurs in RSVP tunnels. The issue is applicable to uncolored tunnels only. [PR1439557](#)
- Sampling applications like **port-mirror** and **inline-jflow** are not supported on VPLS tunnel interfaces in ingress direction where ingress packets are sent to the IRB interface for routing. Configuration of sampling application on VPLS tunnel interfaces in such scenario causes packet to drop in ingress direction. [PR1444849](#)

- If Sx Modification-Request has an Update FAR Apply Action that has the DUPL and DROP bits set, the traffic is dropped as expected. However, the packets are not duplicated to the SX3LIF/MD. This happens for both upstream and downstream traffic. [PR1450859](#)
- When 32000 inetcolor and 32000 inet6color are programmed together, the jsd process is hit. [PR1452464](#)
- In a scaled scenario where the Routing Engine pushes a lot of routes to the Packet Forwarding Engine in the presence of the dynamic tunnel configuration, FIB convergence might take more time, leading to traffic drops. [PR1454817](#)
- Member of lt interface of a rlt interface must have same bandwidth configured. Bandwidth mismatch might lead to unexpected behavior. Changes to lt or rlt interfaces must not be done if a ps interface is anchored over these tunnel interfaces. [PR1458951](#)
- The lt interface Scheduler remains in the invalid state under egress IFD list after changing the lt tunnel to a different Packet Forwarding Engine. [PR1458955](#)
- Changes to rlt interface with ps anchored over is not recommended. For more information, refer to the following Juniper documentation: [\[Unresolved xref\]](#). [PR1460898](#), [PR1460910](#)
- The traffic on GRE interface on both ingress and egress cannot be Layer 2 mirrored. [PR1462375](#)

Interfaces and Chassis

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

MPLS

- The device might use the locally computed path for the PCE-controlled LSPs after the link or node fails. [PR1465902](#)

Platform and Infrastructure

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

Routing Protocols

- Three BGP replication flaps are seen on a new master Routing Engine after GRES. The route synchronization issue is also seen between Routing Engines without GRES. [PR1441925](#)

- When you scale RIB to 80 million after FPC restarts, it is not able to scale on the backup Routing Engine.
[PR1444073](#)

SEE ALSO

[What's New | 85](#)

[What's Changed | 106](#)

[Open Issues | 114](#)

[Resolved Issues | 131](#)

[Documentation Updates | 152](#)

[Migration, Upgrade, and Downgrade Instructions | 153](#)

Open Issues

IN THIS SECTION

- [Application Layer Gateways | 115](#)
- [Class of Service | 115](#)
- [EVPN | 115](#)
- [Forwarding and Sampling | 115](#)
- [General Routing | 116](#)
- [Infrastructure | 124](#)
- [Interfaces and Chassis | 124](#)
- [Layer 2 Features | 126](#)
- [Layer 2 Ethernet Services | 126](#)
- [MPLS | 126](#)
- [Network Management and Monitoring | 127](#)
- [Next Gen Services MX-SPC3 Services Card | 127](#)
- [Platform and Infrastructure | 127](#)
- [Routing Protocols | 129](#)
- [Services Applications | 130](#)
- [Subscriber Access Management | 130](#)
- [VPNs | 130](#)

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.

Application Layer Gateways

- When you use the SIP ALG after payload changed by ALG, some SIP messages size might be bigger than the outgoing MTU interface and it might need to be fragmented. Else, the SIP messages might be dropped by SIP ALG. [PR1475031](#)

Class of Service

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

EVPN

- With Junos OS Release 19.3R1, VXLAN OAM host-bound packets are not throttled with DDoS policers. [PR1435228](#)
- When DHCP is used with EVPN, Layer 2 learning daemon adds a destination route to the kernel with the permanent remote flag while dhcp process adds a destination route with a permanent flag. There might be a race condition where the Layer 2 learning destination route is overwritten by the DHCP route, causing the remote flag to get deleted. This subsequently leads to the ARP route to age out in the kernel. To ensure that DHCP routes are not added to the kernel, you must configure the **forward-only** command under **forwarding-options dhcp-relay**. [PR1439568](#)

Forwarding and Sampling

- For Junos OS Releases 18.4R1 and 18.3R2, if IPv4 prefix is added on a prefix-list referred by IPv6 firewall filter, the following log message is not seen: **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized**. [PR1395923](#)
- On the MX Series routers with MPC line card (except DPC line card) used, if an input firewall filter is configured at the ingress VPLS interface, the packet with a VLAN priority of 5 with three or more VLAN tags might be forwarded into the wrong queue. When this occurs, it might cause traffic loss due to congestion as all traffic is forwarded into the default queue. [PR1473093](#)
- Error of traffic does not get policed as expected after locally switched for VLAN 100 and 101, while verifying the selective local-switching functionality with 4000 VLANs. [PR1436343](#)
- After restart routing, the remote mask, which indicates from which remote PE devices MAC IP addresses are learned, that the routing daemon sends might be different from the existing remote mask that the Layer 2 learning daemon had prior to restart. This causes a mismatch between Layer 2 learning and

routing daemons interpretation as to where the MAC IP address entries are learned, either local or remote, leading to the MAP IP table being out of synchronization. [PR1452990](#)

General Routing

- On the MX104 platform, when using `snmpbulkget` or `snmpbulkwalk` (for example, used by the SNMP server) on a chassisd-related component (for example, `jnxOperatingEntry`), chassis process (chassisd) high CPU usage and slow response might be seen because of a hardware limitation, which might also lead to a query timeout on the SNMP client. In addition, the issue might not be seen while using an SNMP query for interface statistics. As a workaround, to avoid the issue, use either of the following approaches:
 - Use `snmpget` or `snmpwalk` instead of `snmpbulkget` or `snmpbulkwalk` and include the `-t 30` option when doing the SNMP query. For example, `snmpget -v2c -c XX -t 30`.
 - Use the `-t 30` option with `snmpbulkget` or `snmpbulkwalk`. For example, `snmpbulkget -v2c -c XX -t 30`. [PR1103870](#)
- On the MX Series devices, if the **reauthenticate lease-renewal** statement is enabled for DHCP, when the DHCP authentication and re-authenticate lease-renewal occurs, the SDB might go down very frequently. [PR1473063](#)
- In subscriber scenario, when there is a configuration change of the firewall filter used by the subscriber service, the RADIUS accounting updates of service session might have incorrect statistic data. The abnormal accounting data might have impact on billing system, so this issue has service impact. [PR1475729](#)
- If redundant APFes simultaneously fails or reboots while sessions are bound, inconsistencies might occur between the APFes. This inconsistency can occur in rare situations that lead to an `rmrpsd` to generate core files on the backup Routing Engine with additional subsequent APFE failovers. [PR1471580](#)
- Error message are observed during loading of the RLI configurations. [PR1451213](#)
- On Junos OS Release from 16.2R1 onwards, if `commit` is executed after `commit check`, the daemon (for example, `dhcpcd` and `sampled`) might not get started even after the related configuration is successfully committed. [PR1468119](#)
- PPP IPv6 NCP fails to negotiate during the PPP login. [PR1468414](#)
- Traceroute generates ICMP error message like destination host unreachable and time exceeded that helps in identifying the intermediate hops. Code logic for handling ICMP errors was not there as part of asymmetric processing. [PR1466135](#)
- When MS-MIC becomes unreachable or SPD restart, the next hop used by `tcp-log` connection are set to discard. However the SPD does not delete this next hop and incorrectly continue using this next hop in the Packet Forwarding Engine. This causes the MS-MIC not able to establish the TCP connection to the syslog server. [PR1469575](#)
- Error messages are observed that do not impact the functionality and can be ignored. [PR1475187](#)

- If a GRES is performed while the **mobile-edge** sessions are logged in with URR enabled, they cannot be removed by the PFCP session deletion request and a portion of these requests are rejected. Sessions gets stuck in the delete state with the use of the **show services mobile-edge sessions summary** command. A Routing Engine reboot through the **request system reboot both-routing-engines** statement is the only way to recover from this state. [PR1478424](#)
- All the **mobile-edge** sessions are lost when you perform a GRES while sessions with URR are logged in. Sessions that attempt to login after the GRES will also be rejected in this state. It is necessary to reboot the router using the **request system reboot both-routing-engines** statement to recover from this state. [PR1478985](#)
- When Layer 2 bridge domain is configured and traffic is flowing only on one particular interface, the MACsec statistics might be updated incorrectly on other channelized MACsec interfaces on the same port group. [PR1472464](#)
- Some of the Demux VLAN over aggregated Ethernet configured statically from CLI configuration are not programmed with the child legs. All the traffic on these logical interfaces are dropped. [PR1476465](#)
- The MX router acting as LNS does not get to program the PFE with I2tp services that causes forwarding issues for the I2tp subscribers. [PR1476786](#)
- When enhanced subscriber management feature are enabled or Junos OS running at Junos OS Release 18.4R1 or later with the **nextgen-stats** enabled and with XL or EA based line cards (MPC2E-NG/MPC3E-NG/MPC5/MPC6/MPC7/MPC8/MPC9) inserted, the Packet Forwarding Engine might be disabled due to major error under very specific and very rare scenarios. [PR1478028](#)
- During simultaneous scale login of default and dedicated bearers, the router might require the control plane to send retries in order to login all the bearers. In rare situations, the router might reject a small number of requests during the stated scale login procedure. As a workaround, the control plane can send new requests in order to eventually login all the bearers on the router. [PR1478191](#)
- In rare situations, the router is unable to process deletion requests from the control plane for URR sessions. In these rare cases, all sessions are stuck in the delete state. This router state can be resolved by rebooting the router with the **request system reboot both-routing-engines** command. [PR1478220](#)
- Traffic loss is seen for 10 seconds when switching from secondary to primary path, even with disabled SBFDD configuration. [PR1478299](#)
- DHCP-server : RADIUS given mask is being reversed. [PR1474097](#)
- Traffic stops after the volume limit is reached but the traffic resumes after APFE fails. [PR1463723](#)
- This issue occurs only with GRES when both the Routing Engines are rebooted together. During chassis init time, the kernel does not allow any GENCFG to be added before the Routing Engines mastership transition is complete, if GRES is active. If ingress multicast replication configuration is changed after GRES is enabled, before rebooting both the Routing Engines, you must disable the GRES configurations. [PR1474094](#)
- During host ping with gr-tunnel endpoint and lt-interface termination, gr-interface input and lt-interface output counters comes as a host path with transit counter. [PR1461593](#)

- VMCORE-../src/junos/bsd/sys/netjsr/jsr_prl.c:2128 [PR1472519](#)
- lke version 2 tunnel flaps with DPD if initiator is not behind NAT. [PR1477483](#)
- The following error messages keep on continuously flooding in the backup Routing Engine: (
**JTASK_IO_CONNECT_FAILED: RPDTM./var/run/rpdtmd_control: Connecting to
128.0,255.255,255.255,0.0.0.0,0.0.0.0, failed: No such file or directory**) [PR1473846](#)
- [firewall] [filter_installation] Output chain filter counters are not correct. [PR1478358](#)
- [firewall] [filter_installation] chain_filters_negative: output chain filter counter is 0. [PR1478371](#)
- The core files are generated at **cassis_alloc_list_timed_free** in **cassis_free_thread_entry**. [PR1478392](#)
- The mustd core files are generated at **dbm_malloc** (**dbmp=< optimized out>**, **size=< optimized out>**) at, **in cdg_add_path** (**cdbmp=0x30000000**, **sidents=< optimized out>**, **didents=< optimized out>**) at **../..../src/ui/lib/constraint/constraint_dependency_graph.c:934**. [PR1475141](#)
- Expected number of 512000 MACs are not re-learned in the bridge table after clearing 512000 MACs from the table. [PR1475205](#)
- Dark window size is more than expected. 31.0872721524375 seconds of traffic lost is observed. [PR1476505](#)
- With the traffic-manager enhanced-priority mode configuration on ZT-based line cards (MPC10), Routing Engine might not be able to send packets after sometime. [PR1476683](#)
- The following error message continuously appears in the backup Routing Engine: (
**JTASK_IO_CONNECT_FAILED: RPDTM./var/run/rpdtmd_control: Connecting to
128.0,255.255,255.255,0.0.0.0,0.0.0.0, failed: No such file or directory**) [PR1473846](#)
- The clksyncd core files are generated after GRES. [PR1474987](#)
- ALG-SIP64: 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](#)
- Bandwidth percent with shaping rate is not working on aggregated Ethernet after deactivating and activating the class of service. [PR1465766](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- With BGP rib-sharding and update-threading, traffic drops 100 percent in the BGP Layer 3 VPN streams, post the removal or restoration configuration. [PR1469873](#)
- Unable to setup 26M sessions (NAPT44) at 900Kpps/s. [PR1470833](#)
- DHCP relay with forward-only fails to send OFFER when the client is terminated on the lt-0/0/0.2 logical tunnel interface. [PR1471161](#)
- Support of **del_path** for the LLDP neighbor changes at various levels. [PR1460621](#)
- JDI_MMX_REGRESSIONS:[MPC10E][LT Tunnel] More number of output packets are seen than expected when the ping function is performed. [PR1461593](#)

- On dual Routing Engines with graceful routing engine switchover (GRES) enabled, after performing GRES, if the configuration synchronization on the backup Routing Engine fails when it becomes the new master Routing Engine, then in rare conditions, some interfaces cannot be deleted or configuration changes cannot be committed. [PR1179324](#)
- When the scale configurations are applied, chassisd CLI command might delay response or might time out for 10 minutes. [PR1454638](#)
- If a vmhost snapshot is taken on an alternate disk and there is no further vmhost software image upgrade, the expectation is that if the current vmhost image gets corrupted, the system boots with the alternate disk so the user can recover the primary disk to restore the state. However, host root file system, the node boots with the previous vmhost software instead of the alternate disk. [PR1281554](#)
- In some scenarios with MPC, major alarm and following messages are generated: **messages log: fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) alarmd[3158]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC x Major Errors Major alarm set, FPC x Major Errors fpcx XQCHIP(46):XQ-chip[0]: DROP protect_regs error (status=0x8) cli> show chassis alarms 1 alarms currently active Alarm time Class Description 2019-01-25 15:18:03 UTC Major FPC x Major.** Despite major alarm set, this error is due to Unknown Error Address logged in hardware to DQ underrun. This message is harmless and has no service impact. [PR1303489](#)
- As a vendor does not use chained CNH, using the feature does not bring in a lot of gain because TCNH is based on an ingress rewrite premise. Without this feature, things work just fine. [PR1318984](#)
- In Message Queuing Telemetry Transport (MQTT) scenario, about 4000 memory leakage every 30 seconds might be seen. However, on very long runs, this uses up high memory, which can indirectly impact other daemons running. [PR1324531](#)
- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infra relies on the integrity of the TCP connections. The reactions to failure situations might not be handled gracefully, resulting in TCP connection timeouts because of jlock hog crossing the boundary value (5 seconds), which causes bad consequences in MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling marker infra in the MX Series Virtual Chassis setup. Unfortunately, there is no immediate plan on enabling marker, because doing so caused a lot of issues in MX Series Virtual Chassis when we tried to enable it. [PR1332765](#)
- In some cases, online insertion and removal (OIR) of a MIC on an FPC can lead to traffic destined to the FPC being silently dropped or discarded. The only way to recover from this is to restart the FPC. The issue is not be seen if you use the corresponding CLI commands to offline and then online the MIC. [PR1350103](#)
- For configurations of bridging routing instances with aggregated Ethernet logical interfaces (6400 IFLs) 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 configurations on MPC7E takes settling time of more than 4 minutes. [PR1359286](#)

- In rare circumstances, a faulty SFP transceiver installed in an MX104 might cause the AFEB to go offline. The backup routing Engine and fan tray will also show alarm. [PR1360426](#)
- If any of the log messages continue to appear in the MPC console, it indicates the presence of a faulty SFP/SFP+ transceiver, which causes the I2C transaction from main board CPU. There is no software recovery available to recover from this situation. These logs also indicate potential I2C transaction failure with any of the 10 ports available with GMIC2 in PIC 0 resulting in unexpected behavior. The following is an example of the error message: **link not coming up or the MIC itself not booting up on restart. I2C Failed device: group 0xa0 address 0x70 Failed to enable PCA9548(0x70):grp(0xa0)->channel(0) mic_sfp_select_link:MIC(0/0) - Failed to enable PCA9548 channel, PCA9548 unit:0, channel ID: 0, SFP link: 0 mic_sfp_id_read: Failed to select link 0.** The only way to recover from these failures is to detect and replace the faulty SFP/SFP+ transceiver plugged into the GMIC2 ports. [PR1375674](#)
- On the EX9208 device, a few XE interfaces go down with the following error message: **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error.** [PR1377840](#)
- In a subscriber management environment, multiple error messages of shmlog: argcnt 309 not enough memory might be generated every hour. [PR1387690](#)
- NAPT66 pool split is not supported with AMS; thus commit must fail with V6 pool in AMS. [PR1396634](#)
- On MX10003 platform, after removing the FPC from a slot, when a new FPC is plugged in, not only does the chassis shows old serial number for this new FPC, but the entire FPC ID EEPROM data is retained. All the fields show old values. [PR1409930](#)
- The MX104 router has the following limitations in error management: The show chassis fpc error CLI command is not available for MX104 in Junos OS Releases 13.3R7, 15.1R2,14.1R5,14.2R4, 13.3R8, and later. Junos OS does not initiate restart of the system on encountering a fatal error. Although you can configure the action Disable PFE for major errors, Junos OS does not disable its only Packet Forwarding Engine on encountering a major error. [PR1413314](#)
- On MX Series routers with Trio chip set based MPCs, unicast traffic might drop when the destination is reachable over an integrated routing and bridging (IRB) interface and a label-switched interface (LSI) with two next hops. [PR1420626](#)
- Certain JNP10008-SF and JNP10016-SF Switch Interface Boards (SIBs) manufactured between July 2018 and March 2019 might have incorrect core voltage setting. As a workaround, you can correct the issue by reprogramming the core voltage and updating the setting in NVRAM memory. [PR1420864](#)
- The following syslog error message is seen: **"Err] dfw_abstract_issu_stats_counters_restore:2222 Failed to find Index = 4613734? during ISSU with 19.3I-20190409_dev_common.0.2212.** [PR1429879](#)
- In gRIBI, programmed routes reference a next-hop group ID, which in turn points to one or more next-hop IDs. Each next-hop ID contains details of the actual next hop. Next-hop group ID and next-hop ID are mapped to an IPv6 prefix (for example, FC01::<GRPID>). In the case of an IPv4 indirect next hop, gRIBI needs to resolve IPv6 via IPv4 next hop over three levels of indirection. Junos OS does not support IPv6 over IPv4 multilevel next-hop resolution. Therefore, gRIBI cannot resolve nexthop GRPID FC01::<grpId> and nexthop ID <FC01> through an actual indirect IPv4 gateway address. [PR1434050](#)

- On dual routing engines MX Series platforms with subscriber management, the replication daemon (repd process) might crash after booting the first time with a newly installed Junos OS release. The repd process synchronizes subscriber information across Routing Engines, so normally the repd crash has no impact on the live service. [PR1434363](#)
- MPC10E 3D MRATE-15xQSFP : Layer 2 over GRE is not supported in Junos OS Release 19.3R1. Even though the configuration gets committed, the feature will not work. [PR1435855](#)
- ZF interrupts for out-of-range Dest PFE INTR for Gnt seen during MPC6/9 linecard bringup. [PR1436148](#)
- Multiple interfaces on specific FPC's go down on MX480 after baseline profile configuration verification. [PR1437221](#)
- With DAC cable used between the EX4600 or QFX5100 lines of switches and an EX Series device and EX device, during reboot of the EX46XX or QFX51XX device, the ports on EX Series device might still be up and running. [PR1441035](#)
- On routers running Junos OS and serving as EVPN gateways, FPC core files available at **heap_block_log** due to NULL entries are also seen in the ifbd level list, which are typically added for flush list. This occurs because of the relink logic failure flush logic for MACs when there is ifbd/bd delete. [PR1441824](#)
- The interface might go into the down state after the FPC restarts with the PTP configuration enabled. [PR1442665](#)
- Push label is missing in the **show route** command output for colored tunnels. [PR1447900](#)
- The **show ddos-protection protocols arp statistics |display xml** command does not show APR violation packets and the packets are not incremented. [PR1449968](#)
- On Junos fusion system, intermediate traffic drop is sometimes seen between AD and SD when sFlow is enabled on the ingress interface. When sFlow technology is enabled, the original packet is getting corrupted for those packets that hit the sFlow filter This due to few packets transmitted from the egress of AD1 is short of FCS (4 bytes) + 2 bytes of data due to which the drops occur. it is seen that the normal data packets are of size 128 bytes while the corrupted packet is 122 byte. [PR1450373](#)
- Chassisd main thread stalls might be seen at a JNS GNF upon GNF SNMP polling for HW-related OIDs (for example, the ones from jnxBoaAnatomy MIB). When the issue is ongoing, the following messages are logged into the GNF /var/log/mastership log if the stall duration is longer than 60 seconds: **> main chassis-control thread stalled for 60 sec - If the stall duration is longer than 200 seconds, then the GNF chassisd will crash and dump a core, and the following message will be logged into the GNF /var/log/messages file: > chassisd[PID]: %DAEMON-3-CHASSISD_MAIN_THREAD_STALLED: main chassis-control thread stalled for 200 sec ? exiting - Once chassisd crashed, it will restart automatically; - These GNF chassisd main thread stalls and GNF chassisd crashes do not cause GNF-assigned FPC restarts/reconnects to chassisd since a JNS GNF does not manage any hardware component; ISSUE-2: ***** - If a GNF chassisd main thread stalls are ongoing and the GNF is restarted, then a service MGD process at the BSYS could start spinning at 100% CPU. This MGD process won't terminate by itself and will be consuming 100% CPU even when the GNF is back online. This condition could be seen at the BSYS JUNOS root shell as follows: > root@BSYS-re0:~ # ps wuaxd | grep mgd | grep -v grep > root 60221 0.0 0.0 733764 7768 - l 09:31 0:00.02 | |-- /usr/sbin/mgd-api -N > root 60223 0.0 0.1 792196**

```
13672 - I 09:31 0:00.05 | |-- /usr/libexec32/bbe-smgd -b -N > root 60225 0.0 0.2 1410708 37740 - S
09:31 0:32.57 | `-- /usr/sbin/mgd -N > root 9954 100.0 0.3 1413260 49528 - Rs 04:11 66:35.37 | |--
mgd: (mgd) (root) (mgd) <<<--- > root 18029 0.0 0.2 1413260 38508 - Is 04:33 0:00.37 | |-- mgd: (mgd)
(root)/dev/pts/1 (mgd) > root 35331 0.0 0.2 1413260 38516 - Is 05:21 0:00.01 | |-- mgd: (mgd)
(root)/dev/pts/0 (mgd) > root 35392 0.0 0.2 1413260 38516 - Is 05:21 0:00.01 | |-- mgd: (mgd)
(root)/dev/pts/0 (mgd) > root 35414 0.0 0.2 1413260 38516 - Is 05:21 0:00.01 | |-- mgd: (mgd)
(root)/dev/pts/0 (mgd) PR1451215
```

- OIDs-related service-set module might not work because the service-set database for SNMP module is not created yet when the following command is performed: **show snmp mib walk enterprises.2636.3.32.1.3.1.4 show snmp mib walk enterprises.2636.3.32.1.3.1.6 show snmp mib walk enterprises.2636.3.32.1.3.1.8** This is expected behavior. "show snmp mib walk 1.3.6.1.4.1.2636.3.32"/ "show snmp mib walk jnxSpSvcSet" OID access (one time good enough) would result in creating the service-set SNMP data base needed. Once "show snmp mib walk 1.3.6.1.4.1.2636.3.32"/ "show snmp mib walk jnxSpSvcSet" is accessed, above OIDs would successfully return proper values. [PR1452153](#)
- Issues with CLI command is observed after ANCP restarts, before ANCP neighbor reestablishes, and before receiving the port-ups. [PR1453837](#)
- Behavior has been modified to display the correct protocol number instead of 255 whenever unknown protocol is encountered. [PR1454792](#)
- IPv6 accounting stop attributes are not correct for MLPPP subscribers. [PR1455175](#)
- When you enable the **persist-groups-inheritance** command and execute a delete operation to delete the entire configuration, if the user selects no and then later tries to commit the configuration changes related to groups, multiple daemons might crash. [PR1455960](#)
- With logical system configuration, filter-based GRE encapsulation does not work. [PR1456762](#)
- After more than 2 million multicast subscribers are activated without performing GRES or bbe-smgd restart, further multicast subscribers might be unable to log in. [PR1458419](#)
- With the scale filter-based forwarding (FBF) configuration, two instances seem to be unable to forward the traffic to respective routing instances. It appears that the FBF programming is incorrect for these two FBF instances. [PR1459340](#)
- 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 18.4 to Release 18.4 or later. [PR1459961](#)
- NAT performance is impacted with remote syslog enabled. [PR1460211](#)
- If a NETCONF session is initiated over an inband connection, the CPU utilization on mgd daemon might be 100 percent after the NETCONF, which executes an RPC call for some commands and gets interrupted by flapping interface. There is no impact observed to the control plane or the forwarding plane, the subsequent NETCONF session continues to function. [PR1464439](#)
- The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline/online scenarios: [Oct 3 08:48:35.836 LOG: Err] ifl ps240.1 (1712): child ifl lt-1/0/0.32767 (7709) already

there [Oct 3 08:48:35.836 LOG: Err] IFRT: 'Aggregate interface ifl add req' (opcode 87) failed [Oct 3 08:48:35.836 LOG: Err] ifl 1712, child ifl 7709; agg add failed [PR1464524](#)

- The following syslog error messages are harmless and expected during FPC offline/restart scenarios with PS-RLT with or without link protection configuration: Nov 12 15:02:00 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x0 delete failed for ifl lt-3/0/0.32767 with err=2 Nov 12 15:02:00 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x1 delete failed for ifl lt-3/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x1 delete failed for ifl lt-5/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_remove_link_from_stack_bundle: vid 0x0 delete failed for ifl lt-5/0/0.32767 with err=2 Nov 12 15:02:43 cleansing kernel: lag_lp_handle_event: LP event = 6, child lt-5/0/0 err = 22 The following syslog error messages are harmless and expected during ISSU or GRES or FPC offline/online scenarios. Nov 12 15:08:37 cleansing fpc3 user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIfName:ps1.0 memberIfName:lt-3/0/0.32767 Nov 12 15:08:37 cleansing fpc3 user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, aggIfName:ps1.0 memberIfName:lt-5/0/0.32767 [PR1466531](#)
- In the PPPoE subscriber management environment, due to the PPPoE inline keepalives timeout, events might be dropped by the Routing Engine and the PPPoE subscribers might get stuck. This issue might cause the PPPoE subscribers to be unable to reconnect. [PR1467125](#)
- SNMP interface-mib stops working for PPPoE clients. In this scenario SNMP works fine for standard queries on the MX Series router; however, for subscriber statistics, it always returns zero value. [PR1470664](#)
- For MPC10E card line, the IS-IS and micro BFD sessions does not come up during baseline. [PR1474146](#)
- On multicore next-generation Routing Engines on MX960, MX240, and MX480 with USF mode enabled and USF based services configuration, the subsequent Junos vmhost upgrade fails with the following error message: **Validation failed ERROR: Failed to add /var/tmp/junos-vmhost-install-mx-x86-64-20.11-20191112_dev_common.0.1229.tgz z warning: Host software installation has failed.** As a workaround you can use the **no-validate** argument to the **request vmhost software add <>** command. For example, **request vmhost software add junos-vmhost-install-mx-x86-64-20.11-20191112_dev_common.0.1229.tgz no-validate**. You can also move the chassis to the baseline configuration and commit, and then perform a software upgrade. After software upgrade the original configuration can be reapplied. [PR1472287](#)
- Adding 100000 CPS IPv6 SFW traffic on top of 12 million PPS/50-Gbps IPv6 SFW traffic results in PPS traffic reduction to 10 million PPS/42-Gbps due to the latency that the CPS traffic processing creates. [PR1472314](#)
- When the same objects are used in both inet and inet6 services of the same subscriber session, deactivation of the first session causes conditions that avoid releasing the UID entry after deactivation of the second service session. This leads to having a stale UID entry and can cause a subscriber connection problem in the future when the UID pool might be completely exhausted. The probability of hitting the issue increases if the amount of subscribers to the amount of unique services ratio is approaching 1, which occurs when almost every subscriber has a service with unique service objects. [PR1188434](#)

- In a subscriber management environment, multiple error messages **shmlog: argcnt 309 not enough memory** might be generated every hour. [PR1384371](#)
- For ATM interfaces configurations, if any logical interface has **allow-any-vci** configuration, then the commit operation might fail. [PR1479153](#)

Infrastructure

- Slow Response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1. [PR1462986](#)

Interfaces and Chassis

- After GRES, the 1-Gigabit Ethernet interface changes to 10-Gigabit Ethernet. [PR1326316](#)
- When priority is increased for all 4000 VRRP sessions, some of the VRRP sessions does not take over the mastership. Peer router continues to become master for those sessions. There are no traffic impact as one of the router is still a master. As a workaround, deactivate and activate VRRP sessions in the backup state. [PR1478349](#)
- On EVPN active or active software design, disabling the ESI logical interface might effect the designated forwarder election of EVPN when the physical interface has ESI configured. In such configuration, disabling the ESI logical interface, type-1 routes (AD/EVI and AD/ES) are not generated from this PE. With ESI configured at the IFD level, as one of the logical interface in the IFD is down, DF election can not happen for the ESI. Also, AD/EVI and AD/ESI routes are deleted. The following warning message occurs upon commit, where this configuration might cause DF election issues and undesired unicast or BUM traffic drop: **DCD_PARSE_CFG_WARNING: aex.y : Disabling the IFL might affect the Designated Forwarder election of EVPN when IFD is having ESI configured.** [PR1467855](#)
- When dynamic DHCP sessions exists in the device and if multiple commits in parallel are performed, the commit might become nonresponsive. [PR1470622](#)
- Commit error was not thrown when member link was added to multiple aggregation group with different interface specific options. When member interface added to bundle with both the **ether** and **gig-ether** interface specific options, the **gig-ether** option takes precedence over the other. [PR1475634](#)
- For ATM interfaces configuration, if any logical interface has the **allow-any-vci** configuration, then the commit operation might fail. [PR1479153](#)
- MPC10 line cards run on newer version of FPC software. Currently , the convergence number for MPC10 is not at par with the legacy MPC card lines for a high scale. Following are few recommendation to achieve better convergence numbers with VRRP on MPC10 card lines:
 - Configure failover delay to 5 seconds. This will help in quick action without impacting the Routing Engine CPU usage.

- Set the VRRP **failover-delay** protocols to 5000. If system has inherit sessions configured then it is recommended to reduce the inherit advertisement interval timer to 6 seconds. Default value of the inherit time is 120 seconds that causes slower convergence for inherit sessions.
- Set VRRP **inherit-advertisement-interval** protocols to 6 seconds. With the above two statements configured, the worst case convergence for a scale of 8000 can be expected to be at 38 seconds. [PR1474656](#)
- When FPC is restarted all VRRP transmit sessions anchored on that FPC gets redistributed to other available FPC. Tx gets disrupted during this time causing flap at peer end. Because peer router takes mastership traffic that gets redirected to peer for some time. But when transmit sessions are up again on the DUT, the peer router moves to backup. For active VRRP sessions, traffic revert back to original master quickly. Its primarily because advertisement interval of active VRRP session is 1 second. But traffic for inherit session does not revert back to original router quickly. This is because advertisement interval of inherit session is 120 seconds. Even though the peer router has moved to backup intermediate switch still point towards peer for VMAC. It changes only after getting packet from original master. This might take up to 120 seconds for inherit sessions that causes silent discarding of the sessions for 120+ seconds. Similar traffic drop can be seen when FPC4 is restarted on rubles when its has mastership. This can be avoided by reducing advertisement interval of inherit sessions with following command set protocols VRRP inherit-advertisement-interval to 6 seconds. After configuring this worst case, loss were observed to be around 8.5 seconds. [PR1474694](#)
- When there are three VRRP routers (for example, R1, R2, and R3), the VRRP priority on R1 is larger than R2 and R2 is larger than R3. Additionally, a firewall filter on R3 interface input direction is configured to drop all VRRP packets. Then, continuous VRRP state transition (VRRP master or backup flaps) might be seen. It might affect the service. [PR1446390](#)
- The voltage high alarm might not be cleared when voltage level comes back to normal for MIC on MPC5. [PR1467712](#)
- Traffic hit can be as high as 129 seconds when the track route recovers with active or inherit configuration. [PR1475140](#)
- When the addition and the deletion of an logical interface (both logical interfaces with same VLAN ID) is performed in a single commit configuration, the check fails with the following error message: **duplicate VLAN-ID** [PR1477060](#)
- Traffic is seen for 248 seconds when an aggregated Ethernet member link is brought down with minimum link configuration. [PR1477821](#)

Layer 2 Features

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

Layer 2 Ethernet Services

- In EVPN multi-homed active/active scenario, when LACP is enabled on PE-CE child member links and after recovering from a core-isolation on the PE device, the PE-CE child member links might be stuck in Detached state if LACP sync-reset feature is enabled on the CE device. The child links on the CE device might show the LACP state as Collecting Distributing. However, on the PE devices, the LACP state might be shown as Detached. [PR1463791](#)
- EVPN-VXLAN core isolation is not working when the system is rebooted or the routing is restarted. [PR1461795](#)

MPLS

- In RSVP LSP with loose or undefined path, the LSP might stay in a down state due to loop detection after the link in the path flaps. [PR1384929](#)
- RPD crashes on the backup Routing Engine when LDP tries to create LDP p2mp tunnel upon receiving corrupted data from the master Routing Engine. [PR1479249](#)
- In a corner case on Junos OS platform, where the family circuit cross-connect is configured along with any other existing family within the same interface such as, inet and inet6, which Junos OS never allows to do so, but somehow a customer did it, and if the family circuit cross-connect is deleted from the interface, that causes kernel to crash and the device reboot automatically, all the traffic will be interrupted. [PR1478806](#)
- The rpd crash might be seen after some commit operations, which might affect the RSVP ingress routes. [PR1471281](#)
- With LDP-tunneling over RSVP LSP where RSVP LSP has link-protection, LDP route might flap when the interface on the bypass is brought down. [PR1450516](#)
- The traffic might be silently discarded after the LACP time outs. [PR1452866](#)

Network Management and Monitoring

- Junos OS used to send a cold start trap from the new master just after the first GRES. This was because the cold_start timestamp file was not present or updated after the reboot. [PR1461839](#)

Next Gen Services MX-SPC3 Services Card

- NAPT66 pool split is not supported with AMS. As a result, commit fails with V6 pool in AMS. [PR1396634](#)
- IPv6 throughput is not at par with IPv4. [PR1439917](#)
- IPv6 throughput numbers for NAT with HTTPs traffic is not at par with IPv4. [PR1449435](#)
- NAT performance is impacted with remote syslog enabled. [PR1460211](#)
- Drop in session setup rate for sfw is seen with syslog enabled. [PR1462049](#)
- "TALUS(number) PCIe(number) DMA RX interrupt received. Queue stuck status 0xe0000000" are spurious messages which are triggered in system logs due to queue-back pressure or FPGA drops. [PR1465888](#)
- On MX960, MX240 and MX480 with USF mode enabled and USF based services configuration, subsequent Junos vmhost upgrades will fail with an error. Moving to the baseline configuration and then upgrading works without error. [PR1472287](#)
- Unable to setup 26M sessions (NAPT44) at 900Kpps per second. [PR1470833](#)
- **show services sessions** and **show services sessions extensive output** do not display member interface of the AMS where the session got landed. It displays AMS interface name only. [PR1474313](#)
- Adding 100,000 CPS IPv6 SFW traffic on top 12 million PPS/50 gbps IPv6 SFW traffic resulting in PPS traffic getting reduced to 10 million PPS/42 gbps due to the latency created by CPS traffic processing. [PR1472314](#)
- The Next-Gen Services MX-SPC3 services card does not come online automatically when the junos-vmhost image is installed on the Next-Generation Routing Engine (NG-RE): RE-S-X6-64G-UB. [PR1482334](#)

Platform and Infrastructure

- The jcrypto syslog help package and events are not packaged even when errmsg is compiled. [PR1290089](#)
- MAC address does not learn on the correct interface when irb- logical interface is moved from an existing bridge domain to another new bridge domain. [PR1459121](#)
- On Junos OS platforms, if **dot1x** and **interface-mac-limit** are configured when sending traffic continuously to the interfaces, the switch might not be able to learn MAC address. Hence, traffic drop might be seen. [PR1470424](#)

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the following error: **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system re-converging on the expected state. [PR1054798](#)
- Few OAM sessions are not established with scale EVPN ETREE and CFM configurations. [PR1478875](#)
- When traffic is received from 1000 different VRF instance on PF from CE, few flows are dropped at PE. [PR1460471](#)
- On the MX Series devices with chained composite next hop (CNH) for labeled BGP configured, the MPLS COS rewrite does not work for 6 PE traffic. This issue has service or traffic impact. [PR1436872](#)
- In MVPN instance, the traffic drops on multicast receivers within range of 0.1 to 0.9 percent. [PR1460471](#)
- Sometime high CPU utilization is observed in MPC 3D 16x 10GE after ISSU. [PR1461715](#)
- In some cases, the PS interfaces over RLT might appear up but not pass traffic. Log messages reporting ASIC error and a chassis alarm reporting hard FPC errors might also be seen. [PR1400269](#)
- On the EX9208 devices, traffic loss is observed if ingress and egress ports are in different FPCs. [PR1429714](#)
- 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. As a result, the storm-control does not effect the ARP packets on the ports under such bridge-domain. [PR1438326](#)
- A dual Routing Engine Junos node slicing GNF with no GRES configured and with **system internet-options no-tcp-reset drop-all-tcp** configuration might enter dual backup Routing Engine state upon manual GNF Routing Engine mastership switchover attempt with **request chassis routing-engine master [acquire|release|switch]** command from either GNF Routing Engine CLI. [PR1456565](#)
- While SNMP-Agent polls round-trip time (RTT) related to OIDs from router running Junos OS, such as **pingResultsAverageRtt**, the router might respond with zero (0) value even there is no RPM ping failure. The following objects might be impacted: **iso.3.6.1.2.1.80.1.3.1.4 -> pingResultsMinRtt iso.3.6.1.2.1.80.1.3.1.5 -> pingResultsMaxRtt iso.3.6.1.2.1.80.1.3.1.6 -> pingResultsAverageRtt iso.3.6.1.2.1.80.1.3.1.7 -> pingResultsProbeResponses iso.3.6.1.2.1.80.1.3.1.9 -> pingResultsRttSumOfSquares**. [PR1458983](#)
- The Layer 2 traffic sent from one member to another member is corrupted on MX Series Virtual Chassis. [PR1467764](#)
- On the MX150 devices, the default subscriber management license does not include Layer 2 TP. [PR1467368](#)

Routing Protocols

- LDP OSPF are in synchronization state because the IGP interface is down with ldp-synchronization enabled for OSPF. `user@host> show ospf interface ae100.0 extensive` Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. As per the current analysis, the IGP interface goes down because although LDP notified OSPF that LDP synchronization was achieved, OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up yet. [PR1256434](#)
- BGP graceful restart might have some traffic loss when sharding is enabled. [PR1475773](#)
- It is possible for a GNF with Rosen 6 multicast to display stuck KRT queue entries after recovery from a dual Routing Engine reboot at the BSYS. [PR1367849](#)
- TI-LFA backup path for adj-sids is broken in OSPF, where the shortest path to the node opposite the adj-sid is not the one-hop path over the interface indicated by the adj-sid. [PR1452118](#)
- Multiple nonstop attempts to clear IGP database might result in routing daemon generating a core file when locally computed SR-TE LSPs are configured in the order of thousands. [PR1456212](#)
- Removal of cluster from the BGP group might cause prolonged convergence times. [PR1473351](#)
- sftp does not connect properly and the following error message is seen: **Received message too long.** [PR1475255](#)
- Even when `protocols mpls traffic-engineering bgp-igp` command is configured, the UDP tunnel routes are not added to inet.0. The UDP tunnel routes are added only to inet.3 table whether the command is configured or not. [PR1457426](#)
- Consider the case where the backup nexthop for a route in inet6.3 has all valid labels except for the last label. While it is not possible to install a working backup path in inet6.3, it is possible to install a working backup path for inet6.0. This is because the inet6.0 backup path is derived from the inet6.3 backup path by removing the last label. Removing the last label leaves a label stack with all valid labels. However, the current implementation does not install the inet6.0 backup path. [PR1458791](#)
- When Bidirectional Forwarding Detection (BFD) configuration is removed, a BFD packet with session state set to AdminDown and diagnostic code set to some appropriate value must be sent to the peer end. However, the RFC does not mandate what diagnostic code must be sent and what action should be taken if a different diagnostic code is received. Currently, if a BFD packet with session state set to AdminDown is received by the Juniper device, the Juniper device checks both the session state and the diagnostic code in the packet. If the session state is AdminDown and the diagnostic code is 7, which means diag AdminDown, the BFD session is set to Down and the BFD client (that is, the service that is protected by BFD) is notified with AdminDown and the BFD client session does not flap. However, if the BFD packet with session state set to AdminDown along with a diagnostic code other than 7 is received, the BFD client is notified with Down and the BFD client session flaps. Juniper device sets the

diagnostic code to 7 for AdminDown packet, so no issue occurs between Juniper devices. If Juniper device is interworking with other vendor device (for example, Huawei device) that does not set the diagnostic code to 7, the BFD client session might flap on Juniper device side when the BFD configuration is removed from the peer end of the BFD session. The fix now only checks the session state and takes action but does not depend on diagnostic code if BFD AdminDown packet is received. [PR1470603](#)

Services Applications

- Memory corruption causes to L2TP process to crash. [PR1407885](#)
- In L2TP subscriber environment with Juniper LTS (L2TP Tunnel Switch) and LNS (L2TP Network Server), if client negotiates LCP (Link Control Protocol) with no ppp-options to LAC (L2TP Access Concentrator), it might cause some problems but it has no service impact. These ppp-options are ACFC (Address and Control Field Compression), PFC (PPP Protocol Field Compression) and ACCM (Async Control Character Map). The reason is that when the MX Series router functions as L2TP LTS or LNS, it will initiate LCP renegotiation (ppp-options) if Last Received LCP CONFREQ AVP (attribute-value pair) is not included in ICCN (Incoming-Call-Connected) message received from LAC. This might cause some problems for peers, which do not support these options and do not want to negotiate with them. [PR1426164](#)

Subscriber Access Management

- CoA request fails to standard attribute proxy state value [33]. [PR1479697](#)

VPNs

- The p2mp lsp replication to backup Routing Engine is not correct. [PR1453900](#)
- MPC10E: Next-generation MVPN for remote IPv6 source is not working. [PR1454163](#)
- MVPN: Traffic loss is observed while verifying multicast route with VT for VPNA. [PR1460480](#)
- The Layer 2 circuit connections might become stuck in the OL state after changing the Layer 2 circuit community and flapping the primary LSP path. [PR1464194](#)
- Traffic loss is observed while verifying multicast route with VT for VPNA. [PR1460480](#)
- After NSR switch overs, sometimes the selective tunnel on the new master Routing Engine might fall back to the inclusive tunnel. After sometime, the traffic gets migrated to the selective tunnel. Some traffic loss are seen during this migration. [PR1475204](#)
- When ingress PE has duplicate selective tunnel for IPv4 and IPv6, where one is a wildcard, the other is specific (s, g). If the ingress replication configuration is deleted on the egress PE, sometimes it is observed that the ingress replication entries in ingress PE (DUT) are not properly flushing out for IPv6, but it got flushed out for IPv4. No traffic loss is observed. All PIM state and multicast traffic are not impacted due to this issue. [PR1475834](#)

- In the NG-MVPN setup, using MPC10 on egress PE with load balance join of multiple groups in C_VPN, the egress PE might not receive multi-cast traffic. [PR1476969](#)
- In MVPN scenario with ingress replication selective provider tunnel being used, if the **ink-protection** statement is added or deleted from the LSP for MVPN, rpd crash might be seen. The reason is that when link-protection is deleted, the ingress tunnel is not deleted, and when link link-protection is added back, it tries to add same tunnel. Due to which, the rpd asserts as same tunnel exists and the rpd generates core files. [PR1469028](#)

SEE ALSO

[What's New | 85](#)

[What's Changed | 106](#)

[Known Limitations | 111](#)

[Resolved Issues | 131](#)

[Documentation Updates | 152](#)

[Migration, Upgrade, and Downgrade Instructions | 153](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.4R1 | 132](#)

This section lists the issues fixed in Junos OS Release 19.4R1 for MX Series routers.

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

Resolved Issues: 19.4R1

Class of Service (CoS)

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

EVPN

- Asynchronous results between ARP table and Ethernet switching table occurs if EVPN ESI link flaps multiple times. [PR1435306](#)
- EVPN or MPLS IRB logical interface might not come up when local Layer 2 interface is down. [PR1436207](#)
- The specific source ports of UDP packet are dropped on EVPN or VXLAN setup. [PR1441047](#)
- The rpd might crash or consume 100 percent of CPU after flapping the routes. [PR1441550](#)
- Restarting Layer 2 learning might cause some remote MAC addresses to move into forwarding dead state. [PR1441565](#)
- Traffic drop might be seen in EVPN Layer 3 Gateway. [PR1442319](#)
- Core-isolation feature does not work after you set or delete the **no-core-isolation** command on MX Series router. [PR1442973](#)
- The EVPN type 2 routes might not have advertised properly in the logical systems. [PR1443798](#)
- The local host address is missing from the EVPN database and mac-ip-table. [PR1443933](#)
- The bridge mac-table age timer does not expire for rbeb interfaces. [PR1453203](#)
- Instance type is changed from VPLS to EVPN, resulting in loss of packet. [PR1455973](#)
- Preference-based DF Election algorithm does not work on LT interface. [PR1458056](#)
- The rpd crash might be seen if BGP route is resolved over the same prefix protocol next-hop in inet.3 table that has both RSVP and LDP routes. [PR1458595](#)
- The DF router might send ARP request or NS to the local segment. [PR1459830](#)
- In EVPN scenario, memory leak might be observed when **proxy-macip-advertisement** is configured. [PR1461677](#)
- Traffic received from VTEP is dropped if the VNI value used for type-5 routes is greater than 65535. [PR1461860](#)
- Rpd might crash with EVPN-related configuration changes in static VXLAN to MPLS stitching scenario. [PR1467309](#)

Forwarding and Sampling

- You might not be able to apply the firewall filter configuration change after ISSU upgrades to release 16.1R1 or later. [PR1419438](#)
- The following syslog error messages are seen at **pfed: rtplib: ERROR received async message with no handler: 28** [PR1458008](#)

- On the MX Series and QFX Series devices, the Layer 2 ald process might leak memory. [PR1455034](#)
- The rt-delay-threshold can be set below 1 second but rt-marker-interval is limited to 1 second. [PR1425544](#)
- The high CPU utilization of Layer 2 ald is seen after replacing EVPN configuration. [PR1446568](#)
- On MX Series routers with MPC10 line cards, the incoming packets might get dropped. [PR1446736](#)
- On MX204, input/output counters of aggregated Ethernet bundle or member links configured on non-default logical systems are not updated. [PR1446762](#)
- ARP packets gets dropped by Packet Forwarding Engine after chassis-control in the MX Series routers. [PR1450928](#)
- Commit error and dfwd core file might be observed when you apply a firewall filter with the **then traffic-class** or **then dscp** action. [PR1452435](#)
- The following false warning message is seen on commit (commit check) after upgrading to Junos OS Release 19.2R2-S1.4: **warning: vxlan-overlay-load-balance configuration for forwarding options has been changed.....** [PR1459833](#)
- On MX Series router, the following logs are seen: **L2ALD_MAC_IP_LIMIT_REACHED_IF: Limit on learned MAC+IP bindings reached for .local.1048605; current count is 1024.** [PR1462642](#)
- The EA WAN SerDes gets into a stuck state, leading to continuous "DFE tuning timeout' errors and link staying down. [PR1463015](#)
- An output bandwidth-percent policer with logical-bandwidth-policer applied to an aggregated Ethernet bundle along with an output-traffic-control-profile has incorrect effective policing rate. [PR1466698](#)
- Type 1 ESI/ or AD route are not generated locally on EVPN PE in all-active mode. [PR1464778](#)

General Routing

- Load balancing is uneven across aggregated Ethernet member links when the aggregated Ethernet bundle is part of an ECMP path. The aggregated Ethernet member links must span the Virtual Chassis members. [PR1255542](#)
- Unable to configure **pic-mode** when MPC10E is inserted. [PR1452467](#)
- Basic circuit cross-connect traffic flow does not occur with the logical systems. [PR1474983](#)
- Service accounting statistics do not get updated after changes are made to the firewall. filters [PR1472334](#)
- System reboot is required when GRES is enabled or disabled with the **mobile-edge** configuration. [PR1444406](#)
- Agentd memory might leak and crash when the RPD session closes without releasing memory. [PR1455384](#)
- Active error counts are not increasing for Layer 3 circuit in SYNCE cards. [PR1472660](#)
- The PTP function might consume the kernel CPU for a long time. [PR1461031](#)
- Not able to get the service sessions when configure NAT64 with destination-prefix length is 32. [PR1468058](#)

- Inner-list functionality with dual tag does not work. Traffic gets dropped at the ingress port. [PR1469396](#)
- Memory leak on Layer 2 cpd process causes Layer 2 cpd to crash. [PR1469635](#)
- On MPC10 interfaces, certain configuration steps might cause traffic to not get policed properly. [PR1470629](#)
- The interfaces on MPC-3D-16XGE-SFPP card does not get created after upgrading the system to Junos OS Release 18.1 and later. [PR1471429](#)
- In cRPD platform, license violations are captured as nagging log messages and no alarm is raised. [PR1471455](#)
- PCC tries to send a report to PCE but the connection between PCC and PCE is not in the up state especially in the case of MBB in PCE provisioned or controlled LSP. [PR1472051](#)
- Active error counts are not increasing for I2C in the SYNCE cards. [PR1472660](#)
- MX10000 QSA adapter lane 0 port goes in the down state when disabling one of the other lanes. [PR1474231](#)
- The **show services sessions** and **show services sessions extensive output** command does not display the member interface of the AMS where the session got landed. It displays only the AMS interface name. [PR1474313](#)
- **request system [halt | power-off]** reboots the system instead of halting the system. [PR1474985](#)
- The physical interface of aggregated Ethernet might take time to come up after disabling or enabling the interface. [PR1465302](#)
- Observing **bbe-smgd-core (0x000000000088488c** in **bbe_autoconf_delete_vlan_session_only (session_id=918)** at
 `../..../src/junos/usr.sbin/bbe-svcs/smd/plugins/autoconf/bbe_autoconf_plugin.c:3115`). [PR1464371](#)
- ZT VPLS: The **native-vlan-id** functionality does not work and an untagged traffic does not pass with the **native-vlan-id** configuration. [PR1463544](#)
- Traffic might be impacted due to fabric hardening being stuck. [PR1461356](#)
- The **SmiHelperd** process is not initialized in the Junos OS PPC Releases. [PR1455667](#)
- Queue data might be missing from the following path: **/interfaces/interface/state** [PR1456275](#)
- Interface with Tri-rate Copper SFP (P/N:740-01311) in MIC 3D 20x 1GE(LAN)-E,SFP stops forwarding traffic after unified ISSU. [PR1379398](#)
- The **vehostd** application fails to generate a minor alarm. [PR1448413](#)
- IPv6 throughput numbers for NAT with HTTP traffic is not at par with IPv4. [PR1449435](#)
- JFLOW: reducing the maximum flow table size when you use **Flex-flow-sizing**. [PR1413513](#)
- The severity of the following error is reduced from fatal to major:
XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR. [PR1390333](#)

- The **high-cos-queue-threshold** range is changed to [uint 0 .. 90;]. [PR1390424](#)
- The PPPoE subscribers are not able to reconnect after FPC reboots. [PR1397628](#)
- The rpd generates the following core files: **cmgr_if_route_exists_condition_init**, **ctx_handle_node**, **task_reconfigure_complete**. [PR1401396](#)
- Change the default parameters for resource-monitor rtt parameters. [PR1407021](#)
- When you use the inline J-Flow application, the FPC crashes and slows the convergence upon HMC fatal error condition. [PR1407506](#)
- For the initial packet, which is specific to MPC10 and onward, the ICMP redirect s are not seen at the source and packets are sent to the better next hop. [PR1409346](#)
- On MX150, the log severity level changes. [PR1411846](#)
- On platforms running Junos OS Evolved, the redirect IP supports BGP flowspec filters. [PR1413371](#)
- Behavior issues occur with SR-TE Junos telemetry interface sensors when IS-IS sensors are also enabled and the route nexthops are aggregated Ethernet interfaces. [PR1413680](#)
- On PowerPC based MX Series platforms, the DHCP/DHCPv6 subscribers might fail to establish sessions. [PR1414333](#)
- cRPD does not restrict the number of simultaneous JET API sessions. [PR1415802](#)
- The JSU package installation might fail. [PR1417345](#)
- The rpd core files are seen when you restart the rpd or when the logical system is deactivated. [PR1418192](#)
- Changing CAK and CKN multiple times within a short interval (around 5 minutes) sometimes show the security MACsec connection's inbound and outbound channel display with more than one AN active. But on the Packet Forwarding Engine hardware side, the correct AN and SAK is programmed and MKA protocol from both ends transmit correct and latest AN on each hello packet. You should not see any traffic drop due to this display issue. [PR1418448](#)
- The **ROUTING_LOOP_DETECTED** subcode is not generated under **PATHERR_RECV** code when a strict path loop is created for LSP event telemetry notifications. [PR1420763](#)
- The jnxFruState shows value as 10 for Routing Engine instead of 6 in response to .1.3.6.1.4.1.2636.3.1.15.1.8.9.1.0.0. [PR1420906](#)
- MX Series router LNS might fail to forward the traffic on the subscriber access route. [PR1421314](#)
- After the control plane event, a few IPsec tunnels fail to send traffic through the tunnel. [PR1421843](#)
- RSI bloat occurs due to VM host-based log collection. [PR1422354](#)
- The XML output might be not hierarchically structured if you use the **show security group-vpn member ipsec statistics** command. [PR1422496](#)
- The **show system subscriber-management summary** command should include the failure reason for standby disconnect when primary and back Routing Engine memories are not matched. [PR1422976](#)

- Ports might get incorrectly channelized if they are already of 10-Gigabit Ethernet and they are channelized to 10-Gigabit Ethernet again. [PR1423496](#)
- Configuration commit might fail when the file system gets into full state. [PR1423500](#)
- Even when disk-failure-action reboot or disk-failure-action halt are configured, the system does not reboot or halt when disk error is encountered. [PR1424187](#)
- The rpd keeps crashing after changing configuration. [PR1424819](#)
- The mspmand process might crash and restart with a mspmand core file created after doing a commit change to deactivate and activate the service set. [PR1425405](#)
- One hundred percent of CPU usage is seen on route monitor of static routes after the client is disconnected from prpd server. [PR1425559](#)
- On MX204 or MX10003, MPC reboot or Routing Engine mastership switchover might occur. [PR1426120](#)
- Observing NPC core at trinity_rtt_hw_bulk_helper, trinity_rt_delete, rt_entry_delete_msg_proc (rt_params=0x48803bd8) at ../../../../src/pfe/common/applications/route/hal/rt_entry.c:5210. [PR1427825](#)
- On MX Series platforms with PPP configuration, when something abnormal happens such as the user dialup router is abnormally powered off or the keepalive packet is dropped due to network problem, the PPP session ages out. In a rare case, the PPP session does not get deleted, which prevents the new session from being created. So the new session is not able to log in. The PPP traffic might be dropped because of the duplicate-protection feature on the interface. And the IP address of the PPP interface cannot be pinged. [PR1428212](#)
- Incorrect display of MAC/MAC+IP and count values are seen after setting **global-mac-limit** and **global-mac-ip-limit**. [PR1428572](#)
- On MX10003 platform, fabric drops might be seen when two FPCs come online together. [PR1428854](#)
- The aggregated Ethernet interface does not come up after rebooting the FPC or device although the physical member link is up. [PR1429917](#)
- The routers that are configured with protect core might send ipfix sampling packets with the wrong next-hop information. [PR1430244](#)
- Performance degradation is observed for about 20 seconds after the fabric board on MX10008 or MX100016 is taken offline. [PR1430739](#)
- Error might occur when you use a script to load the configuration. [PR1431198](#)
- The l2cpd process might crash and generate a core dump file when interfaces are flapping. [PR1431355](#)
- Dual stack subscriber accounting statistics are not baselined when one stack logs out. [PR1432163](#)
- Traffic might be sent on the standby link of an aggregated Ethernet bundle and get lost with LACP fast-failover enabled. [PR1432449](#)
- After you delete the CLI configuration chassis license bandwidth, the bandwidth value does not default to maximum bandwidth value. [PR1433157](#)

- The rpd generates core files during the route flash when the policy is removed. [PR1434243](#)
- Packet Forwarding Engine memory leak might be seen if MLPPP links are flapped. [PR1434980](#)
- MicroBFD 3x100ms flap is seen upon inserting a QSFP in another port. [PR1435221](#)
- Traffic drops when session key rolls over between primary and fallback nodes for more than 10 times. [PR1435277](#)
- The mc aggregated Ethernet interface might get stuck in the waiting state after a device reboot. [PR1435874](#)
- The local route in the secondary routing table gets stuck in the KRT. [PR1436080](#)
- The ifHCInOctets counter on aggregated Ethernet interface shows the zero value when SNMP MIB walk is executed. [PR1436201](#)
- When you reboot or power off the backup Routing Engine, a trap message is displayed. [PR1436212](#)
- A few static PPP subscribers are stuck in the initialization state permanently and the **Failed to create client session, err=SDB data corrupted** error is seen. [PR1436350](#)
- The subscriber interim statistics might reset to zero and idle-timeout might not work in the MX Series Virtual Chassis setup. [PR1436419](#)
- Not able to reach the router after downgrading from Junos OS Release 18.2-20190513.0 to 18.2R2.6. [PR1436832](#)
- On MPC10, the micro-BFD sessions do not come up in centralized mode. [PR1436937](#)
- Ping fails on logical interfaces with dual tag. [PR1437608](#)
- The CPU utilization on a daemon might be around 100 percent or the backup Routing Engine might crash in race conditions. [PR1437762](#)
- ISSU fails from 19.1R1 legacy Junos release images. [PR1438144](#)
- RPD might generate a core file during router boot up due to file pointer issue because there are two code paths that can close the file. [PR1438597](#)
- On MX Series Virtual Chassis platforms, subscriber flows might not be synchronized between aggregated Ethernet members. [PR1438621](#)
- The syslog server over TCP-based-syslog does not receive carrier-grade NAT logs when data traffic is sent at 10,000 sessions/sec. [PR1438928](#)
- Incorrect values are observed in the **JUNIPER-TIMING-NOTFNS-MIB** table. [PR1439025](#)
- On platforms running Junos OS Evolved, the **show jdaf service cmd statistics/clients** command is not available. [PR1439118](#)
- In an MX Series Virtual Chassis, FPC on Virtual Chassis backup router might reboot. [PR1439170](#)
- Interface-specific filters do not have any effect on MPC10E line cards. [PR1439327](#)

- When a group is applied at non-root level, updating commands inside the group does not update the hierarchies where they are applied. [PR1439805](#)
- IPv6 throughput is not on par with IPv4. [PR1439917](#)
- PRPD flexible tunnel profile queries do not return DMAC when set to all zeros by client. [PR1439940](#)
- The following syslog error message might appear: **UI_SCHEMA_MISMATCH_SEQUENCE: Schema header sequence numbers**. [PR1440141](#)
- On VMware/ESXi in a multiple FPCs chassis, the interfaces assignment is incorrect and some physical interfaces are not visible. [PR1440360](#)
- CoS-related errors are seen and subscribers are not able to get service. [PR1440381](#)
- On MX Series, CPU might hang or interface might stop working on 100-Gigabit Ethernet port. [PR1440526](#)
- In some situations when too many statistics need to be collected from the Packet Forwarding Engine level at the same time, the bulk manager thread of the FPC microkernel level might be continuously busy and cause permanent 100 percent FPC CPU utilization. [PR1440676](#)
- DHCP offer packets toward IRB over LT interface are getting dropped in DHCP relay environment. [PR1440696](#)
- The Layer 2 dynamic VLANs miss when an interface is added to or removed from an aggregated Ethernet bundle. [PR1440872](#)
- When laser receiver power gets -inf , the telemetry value corresponding to -infinity should be equivalent to IEEE 754, which is a single-precision float and the 32-bit value should be 0xff800000. [PR1441015](#)
- New OID is added that calculates the buffer utilization where inactive memory is not considered as free memory. [PR1441680](#)
- Egress stream flush failure and traffic black hole might occur. [PR1441816](#)
- LINUX:SNMP trap comes twice for FRU removal in MX10000, with one trap with FRU name as FPC: JNP10K-LC2101 and second with FRU name as FPC @ 1/*/* . [PR1441857](#)
- The packets originating from the IRB interface might get dropped in a VPLS scenario. [PR1442121](#)
- The chassisd is unable to power off a faulty FPC after Routing Engine switchover, leading to chassisd restart loop. [PR1442138](#)
- The operational status of the interface in hardware and software might be out of synchronization in EVPN setup with arp-proxy feature enabled. [PR1442310](#)
- In the enhanced-ip or enhanced-ethernet mode with DCU (destination-class-usage) accounting enabled, MS-DPC might drop all traffic that should egress through aggregated Ethernet interface. [PR1442527](#)
- EVENT UpDown interface logs are partially collected in syslog messages. [PR1442542](#)
- Different formats of the B4 addresses might be observed in the SERVICES_PORT_BLOCK_ALLOC/RELEASE/ACTIVE log messages. [PR1442552](#)

- A few Path Computation Element Protocol (PCEP) logs are marked as error even though they are not an error. The severity of those logs is now marked as INFO. [PR1442598](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- On MX Series platforms, the bbe-smgd might crash. [PR1443109](#)
- The BGP session fails to be establish when you use the firewall filter to de-capsulate BGP packets from the GRE tunnel. [PR1443238](#)
- The kmd process might crash and restart with a kmd core file created if IP of NAT mapping address for IPsec-VPN remote peer is changed. [PR1444183](#)
- MX204: GRE data packets with size greater than the MTU get dropped when sampling is enabled on the egress interface. [PR1444186](#)
- For eventd, you might observe high CPU utilization along with error logs. [PR1444462](#)
- Inline-keepalive might stop working for LNS subscribers if the routing-services statement is enabled. [PR1444696](#)
- MX:EAPoL: MACsec sessions are down with unicast EAPOL destination address. [PR1445052](#)
- Access route might be stuck in bbe-smgd and rpd might not be cleared. [PR1445155](#)
- The CPCDD process continuously generates core files and stops the process in ServicesMgr::ServicesManager::cpcddSmdInterface::processInputMsg. [PR1445382](#)
- ECMP-FRR might not work for BGP multipath ECMP routes. [PR1445391](#)
- Detached LACP member link gets LACP state as enabled in Packet Forwarding Engine when switchover occurs because of device reboot. [PR1445428](#)
- The 1-Gbps interface on MX204 might stay down after the device reboots. [PR1445508](#)
- Junos OS Release 19.2 group level uses wildcard <*>. [PR1445651](#)
- The Layer 2 ald might crash when FPC restarts. [PR1445720](#)
- The mspmand process might crash if URL filtering is configured and one blacklisted domain name is a substring of another blacklisted domain name in URL filter database file. [PR1445751](#)
- On Ex3400, DOT1XD core file is found at `macsec_update_intf macsec_destroy_ca` directory. [PR1445764](#)
- The jdncpd process might crash after issuing the **show access-security router-advertisement-guard** command. [PR1446034](#)
- When you use a converged CPCD, MX Series router rewrites the HTTPS request with destination-port 80. [PR1446085](#)
- When switchover happens with MX Series router with service interface that has NAT and GR configuration, the static route for NAT never comes up. [PR1446267](#)
- The following rpd core file appears: `task_block_verify(task_io_hook_block, hook),jtask_jthr_endpoint_internal_sanity ,jtask_jthr_endpoint_sanity`. [PR1446320](#)

- Accurate statistics might not include the forwarded packets during the last 2 seconds before subscriber termination. [PR1446546](#)
- NAT service set in certain scale might fail to get programmed. [PR1446931](#)
- ISSU: Core-RMPC3.gz.core.0 and ISSU failure are seen for MPC5. [PR1446993](#)
- The J-Flow version 5 stops working after input rate values are changed. [PR1446996](#)
- Sonet option is enabled for the xe interface. [PR1447487](#)
- DT_BNG: bbe-smgd core file on backup Routing Engine in bbe_ifd_add_vlan (ifd=0x8c3e835, ifl=0xcaf59f18) at ../../../../src/junos/usr/sbin/bbe-svcs/smd/infra/bbe_ifd.c:6374. [PR1447493](#)
- On MX Series routers, when you use ps interface over redundant logical tunnel in Layer 2 circuit, the pseudowire traffic gets dropped or discarded if **no-control-word** is enabled. [PR1447917](#)
- The rpd process might crash if BGP is activated or deactivated multiple times. [PR1448325](#)
- PCEP: PCE-initiated SR LSP in the first PCE tears down when PCInitiate LSP is brought up and brought down in the second PCE. [PR1448665](#)
- DCD CPU spike is observed after Junos OS upgrade from Junos OS Release 14.2 to Release 16.1. [PR1448858](#)
- Unexpected behavior might occur when you use the **load override** command. [PR1448965](#)
- IPv6 packets might get dropped when vMX acts as a VRRPv3 gateway. [PR1449014](#)
- FPC reboots when PIC 0 is taken offline. [PR1449067](#)
- The DHCP relay feature might not work as expected with **helpers bootp** configured. [PR1449201](#)
- The packets might get dropped when the usage of CPU Core 0 on the host is high. [PR1449289](#)
- There might be an increase in the maximum value of **delegation-cleanup-timeout**. [PR1449468](#)
- Changing the hostname triggers LSP on-change notification and not the adjacency on-change notification. [PR1449837](#)
- The following error message is changed: **Failed to fetch JDM software version from <other_server_full_name>**. If authentication of peer server is not done yet, run **request server authenticate-peer-server from the earlier message: Failed to fetch software version from <other_server_full_name>** to make the error message meaningful. [PR1449871](#)
- On MX Series router running Junos OS enhanced subscriber management feature, no localhost logical interface for rtt 65535 is observed. [PR1450057](#)
- The power that supplies LED on the status panel remains green while one or more PEMs have FAULT LED turned on. [PR1450090](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)
- MoFRR: Issue with MLD plus IGMP scale. [PR1450803](#)

- On VLAN configuration changes with Layer 2 ald, restart might cause kernel synchronization issues and impact forwarding. [PR1450832](#)
- On MPC10E, dcd is unable to clean stale mt- logical interfaces while reloading rosen configuration on the DUT. [PR1450953](#)
- When you use the Standard_D5_v2, which has 16 vCPUs and 56 GB of memory, the deployment fails. [PR1450975](#)
- JNP10000-LC2101 FPC generates **Voltage Tolerance Exceeded** major alarm for EACHIP 2V5 sensors. [PR1451011](#)
- The burst size is not updated when the dynamic profile uses the static traffic control profile. [PR1451033](#)
- SNMP query for IPsec with decrypted or encrypted packets does not fetch the correct values. The following error is observed: **KMD_SNMP_FATAL_ERROR** [PR1451324](#)
- The VFP external static IP configuration is not persistent after rebooting the VFP instance. [PR1451709](#)
- RMPC core files are found after the configuration changes are done on the network for PTP or clock synchronization. [PR1451950](#)
- On MX Series, the dropped packets are seen on MQ/XM-based MPCs, although there is no traffic flowing through the system. [PR1451958](#)
- The mgd might crash when you use the **replace pattern** command. [PR1452136](#)
- On the MX10000 and PTX10000 lines of routers with Routing Engine redundancy configuration enabled, the firmware upgrade for PSU (JNP10000-AC2) and JNP10000-DC2) might fail due to lcmd being disabled by the firmware upgrade command. [PR1452324](#)
- PLL errors might be seen after FPC reboots or restarts. [PR1452604](#)
- On MX10003, MACsec framing errors are seen whenever the sequence number exceeds 2 power 32 with extended packet numbering (XPN). [PR1452851](#)
- Hide the **drop-flow** command under tcp-non-syn configuration. [PR1452902](#)
- On MPC10E, inconsistency between AFT and non-AFT line cards occurs while displaying ldp p2mp traffic-statistics on bud node. [PR1453130](#)
- The values displayed in the output of the **show snmp mib walk jnxTimingNotfnsMIB.3** command are not correct. This MIB table is responsible for timing feature defect or event notification. [PR1453436](#)
- PTP can go out of synchronization due to Layer 2 ald hwdb access failure. [PR1453531](#)
- On MX10003 platform, alarms are not sent to syslog. [PR1453533](#)
- Delay in freeing processed defragment buffers lead to prolonged flow control and might crash. [PR1453811](#)
- The ANCP interface-set QoS adjusts might not be processed. [PR1453826](#)
- The FPC might crash when the severity of error is modified. [PR1453871](#)

- Timestamp is not shown with count option after changing the match condition for the **show <> | mathc <> | count** command. [PR1454387](#)
- On the MX204 platform, the **radius-acct-interim** statistics are not populated for subscribers. Statistics are properly populated in the **radius-acct-stop** packets. [PR1454541](#)
- The 100-Gigabit Ethernet interfaces might not come up again after going down on MPC3E-NG. [PR1454595](#)
- The access request for Layer 2 BSA port up might not be retransmitted if the RADIUS server was unreachable. [PR1454975](#)
- JNS/GNF: CRAFTD syslog fatal errors along with junk characters appear upon startup and exits after four startup attempts. [PR1454985](#)
- JET/JSD RPC tag handling bug. [PR1455426](#)
- Device chooses incorrect source address for locally originated IPv6 packets in routing-instance when destination address is reachable through static route with next-table command. [PR1455893](#)
- Excessive fragmentations of IKE packets might cause failure in the tunnel establishment. [PR1455896](#)
- The BgpRouteInitialize API exits with error code 2. [PR1455967](#)
- The rpd crashes at `__mem_assert func=0x2266f3a "free_jemalloc"`, while adding and deleting the sensors. [PR1456049](#)
- High temperature from the **show chassis environment** output is observed after MPC4E is inserted to slot 5. [PR1456457](#)
- CLI command with invoke-on and display xml rpc results in unexpected multiple RPC commands. [PR1456578](#)
- All the IPsec tunnels might be cleared when the **clear** command is executed for only one IPsec tunnel with specified service-set name. [PR1456749](#)
- The bbe-statsd process might continuously crash if any parameter is set to 0 in the **mx_large.xml** file. [PR1457257](#)
- Default value of 2^32 replay-window size results in framing errors at an average of one in 2^32 frames received. [PR1457555](#)
- The chassisd process and all FPCs might restart after Routing Engine switchover. [PR1457657](#)
- The **show subscriber extensive** command incorrectly displays DNS address provided to the DHCP clients. [PR1457949](#)
- The subscriber routes are not cleared from the backup Routing Engine when the session is aborted. [PR1458369](#)
- Traffic black hole or MPC crash might be seen on MPC10E during firewall filter terms change. [PR1458499](#)
- If you use dynamic VoIP VLAN assignment, the correct VoIP VLAN information in LLDP-MED packets might not be sent after you commit. [PR1458559](#)

- The FPC X major errors alarm might be raised after committing the PTP configuration change. [PR1458581](#)
- When you perform delete operations, the gRPC updates on_change does not work. [PR1459038](#)
- After you set interface <> is disabled with QSA, the link still remains up. [PR1459093](#)
- The traffic might be stuck on MS-MPC or MS-MIC with sessions receiving a huge number of affinity packets. [PR1459306](#)
- The following error message might be seen after the chassisd restarts: **create_pseudos: unable to create interface device for pip0 (File exists)** [PR1459373](#)
- The **show ancp subscriber access-aggregation-circuit-id < access aggregation circuit ID>** command displays incomplete output. [PR1459386](#)
- Telemetry streaming of mandatory TLV 'ttl' learned from LLDP neighbor is missing. [PR1459441](#)
- The traffic might be silently dropped or discarded during link recovery in an open Ethernet access ring with ERPS configured. [PR1459446](#)
- In MC-LAG scenario, the traffic destined to VRRP-virtual MAC gets dropped. [PR1459692](#)
- After the DRD auto recovery, the traffic blackholing upon interface flaps. [PR1459698](#)
- CPCDD core file is found at
**ServicesMgr::ServicesManager::cpcddSmdlInterface::processServiceNotifyMsg
, SmdlInterface::cbStateSyncServiceNotifyMsgHandler ,statesync_consumer_poll_new_state_cb.**
[PR1459904](#)
- Initial synchronization for **OpenConfig** event sensors are streamed only from producers supporting event paths. [PR1459927](#)
- The PPTP does not work with destination NAT. [PR1460027](#)
- If **vlan-offload** is configured on the VMX platform, **input-vlan-map** might not work. [PR1460544](#)
- The bbe-smgd generates a core file when all RADIUS servers are unreachable. [PR1461340](#)
- When you receive IPv6 over IPv4 IBGP session, the IPv6 prefix is hidden. [PR1460786](#)
- The repd generates a core file during system startup. [PR1461796](#)
- During the BBE statistics collection and management process, issues with the **bbe-statsd** memory on backup Routing Engine occurs. [PR1461821](#)
- JET RIB API RouteRemove and RouteRemoveMatching RPCs do not work as the first RIB API call. [PR1461974](#)
- The rpd might crash after committing **dynamic-tunnel-anchor-pfe** command. [PR1461980](#)
- The **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed"** message appears when both DIP switches and power switch are turned off. [PR1462065](#)
- The flow stuck and flowd watchdog generate core files while trying to ping DNS server 8.8.8.8 on the internet through DUT configured with NAPT44. [PR1462277](#)

- On MX204, RADIUS interim accounting statistics are not populated. [PR1462325](#)
- The vty remote MAC addresses are not learned with correct age if vty is from a line card without Juniper Trio 5 silicon. [PR1463040](#)
- MAC-learning is broken for vlan-id all scenario. [PR1463078](#)
- The subscribers might not pass traffic after making some changes to the dynamic-profiles filter. [PR1463420](#)
- The MPC2E-NG or MPC3E-NG card with specific MIC might crash after a high rate of interface flaps. [PR1463859](#)
- RPC ALG causes MSPMAND core files when MX Series router is used as a stateful firewall with the MS-MIC or MS-MPC service cards. [PR1464020](#)
- The PPP IPv6CP might fail if the **routing-services** command is enabled. [PR1464415](#)
- The **show task memory detail** command shows incorrect cookie information. [PR1464659](#)
- The PPPoE session goes in to terminated state and the accounting stops for the session that is delayed. [PR1464804](#)
- MPC5E or MPC6E might crash due to internal thread hogging of the CPU. [PR1464820](#)
- DNS sinkhole server results in multiple core files. [PR1466567](#)
- Layer 2 wholesale does not forward all the client requests with stacked VLAN. [PR1467468](#)
- Crypto code might cause high CPU utilization. [PR1467874](#)
- The process rpd might crash after making several changes to the flow-spec routes. [PR1467838](#)

Infrastructure

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

Interfaces and Chassis

- Unrelated aggregated Ethernet interfaces might go down if the committing configuration changes. [PR1409535](#)
- The demux interfaces goes down after changing the MTU of the underlying et interface. [PR1424770](#)

- Mixed link-speed aggregated Ethernet bundle are not able to add new sub-interface successfully. [PR1437929](#)
- Targeted-distribution for static demux interface over aggregate Ethernet interface does not take correct LACP link status into consideration when choosing primary and backup links. [PR1439257](#)
- Mgd processes increases because the mgd processes are not closed properly. [PR1439440](#)
- The cfmd process might crash after a restart on Junos OS Release 17.1R1 and later. [PR1443353](#)
- Unrelated aggregated Ethernet interfaces might go down if changes in the configuration are committed. [PR1409535](#)
- Need enhancement to add or delete a single VLAN in vlan-id-list under interface family bridge. [PR1443536](#)
- ISSU might fail when you upgrade a device that has an aggregated Ethernet bundle with more than 64 logical interfaces. [PR1445040](#)
- The OAM CCM messages are sent with single-tagged VLAN even when configuring with two VLANs. [PR1445926](#)
- Not able to connect to newly installed Routing Engine from other Routing Engines in Routing Engines in MX Series Virtual Chassis. [PR1446418](#)
- Initiating a Routing Engine switchover on VRRP backup router through a CLI command (even protocols vrrp delegate-processing ae-irb) might cause VRRP state for aggregated Ethernet bundle interfaces transitions to the master state, then very shortly afterward to backup again. [PR1447028](#)
- The Layer 2 ald might fail to update composite next hop. [PR1447693](#)
- The ifinfo daemon might crash on the execution of the **show interface extensive** command. [PR1448090](#)
- Dual VRRP mastership might be seen after ungraceful Routing Engine switchover. [PR1450652](#)
- LACP daemon crashed continuously. [PR1450978](#)
- The severity level log might be flooded when the QSFP-100GE-DWDM2 is inserted. [PR1453919](#)
- In the CFM UP MEP over Layer 2 VPN or Layer 2 circuit service, the CFM UP MEP session might get stuck in the failed state. [PR1454187](#)
- The VRRP traffic loss is longer than 1 second for some backup groups after performing GRES. [PR1454895](#)
- Mismatched MTU value causes the RLT interface to flap. [PR1457460](#)

J-Web

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

Layer 2 Features

- LSI interface might not be created, causing remote MACs not to be learned and display of the following error log: **RPD_KRT_Q_RETRIES: ifl iff add: Device busy**. [PR1295664](#)
- VPLS neighbors might stay in the down state after configuration changes in vlan-id. [PR1428862](#)

- Traffic drop might be seen when one MX Series Virtual Chassis member reboots and rejoins the Virtual Chassis. [PR1453430](#)

Layer 2 Ethernet Services

- DHCP request might get dropped in DHCP relay. [PR1435039](#)
- The `jdhcpd` process might go into infinite loop and cause CPU full utilization. [PR1442222](#)
- On MX10008 or MX10016 platforms, the `dhcp-relay` command might not work. [PR1447323](#)
- Some additional information can be provided in DHCPv6 option 17. This option can be in SOLICIT or REQUEST messages. BNG should relay the information from this option to RADIUS servers in ACCESS REQUEST message in the attribute 26-207. Before the fix from the PR the information was not relayed. [PR1448100](#)
- PPPoE holding DHCPv6 prefix causes DHCPv6 binding failure due to duplicate prefix. [PR1453464](#)
- DHCP subscriber might not come online after the router reboots. [PR1458150](#)
- DHCP packet might not be processed correctly if DHCP option 82 is configured. [PR1459925](#)
- The ISSU might fail during subscriber in-flight login. [PR1465964](#)

MPLS

- The FPC might be stuck in the Ready state after making a change in the configuration that removes RSVP and triggers FPC restart. [PR1359087](#)
- Static MPLS LSP label might not get installed in MPLS.0 after the link flaps. [PR1457432](#)
- Traffic is silently discarded after the LSP protection link on Huawei transit router goes down. [PR1439251](#)
- Continuous `rpdc` core files are generated at `l2ckt_alloc_label`, `l2ckt_standby_assign_label`, and `l2ckt_intf_change_process` in new backup during GRES in MX2010 box. [PR1427539](#)
- The LDP might withdraw a label for an FEC once the IGP route is inactive in `inet.0`. [PR1428843](#)
- Dynamic SPRING-TE tunnel creation to LDP (non SR) speaking nodes are supported. [PR1432791](#)
- Root XML tag in the output is changed from **rsvp-pop-and-fwd-info** to **rsvp-pop-and-fwd-information** to be consistent with the XML tag convention. [PR1365940](#)
- SRLG entry shows unknown after removing it from configuration in **show mpls lsp extensive** or **show mpls srlg** output. [PR1433287](#)
- The P2MP LSP branch traffic might be dropped for a while when the sender PE device performs switchover. [PR1435014](#)
- The flow label is not pushed when **chained-composite-next-hop ingress l2ckt/l2vpn** is enabled. [PR1439453](#)
- LSI interface Layer 2 Virtual Chassis goes down on one router in VPLS domain through the MPLS path is still available in `inet.3`. Reason shows as **mpls label out of range**. [PR1442495](#)

- The backup LSP path messages are rejected if the bypass tunnel path is an inter-area LSP. [PR1442789](#)
- RSVP path message with long refresh interval is dropped between devices running Junos OS releases earlier than Release 16.1 and devices running versions later than Release 16.1. [PR1443811](#)
- TRUE POC: rpd core files are generated with SNMP polling. [PR1457681](#)
- P2MP LSP might get stuck in the down state after link flaps. [PR1444111](#)
- The rpd memory leak might be seen when the inter domain RSVP LSP is in the down state. [PR1445024](#)
- Traffic might be silently dropped or discarded if two consecutive PLRs along the LSP performs local repair simultaneously under certain misconfigured conditions. [PR1445994](#)
- The transit packets might be dropped if an LSP is added or changed on an MX Series or PTX Series device. [PR1447170](#)
- Traffic drop might be seen after traceoption configuration is committed in RSVP P2MP. [PR1447480](#)
- The rpd generated a core file at **ted_delete_abstract_hop (instance=0x75d33c0, hop_name=< optimized out>)** during abstract-hop testing. [PR1448769](#)
- The LDP route timer resets when committing unrelated configuration changes. [PR1451157](#)
- All LDP adjacencies flap after changing LDP preference. [PR1459301](#)
- The previously configured credibility preference is not considered by CSPF even though the configuration has been deleted or changed to prefer another protocol in the traffic engineering database. [PR1460283](#)
- High CPU usage and rpd core file might be observed if **ldp track-igp-metric** is configured and IGP metric is changed. [PR1460292](#)
- MPLS trace route does not trace the SRUDP tunnel ingress router. [PR1460516](#)

Network Address Translation (NAT)

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

Network Management and Monitoring

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

Platform and Infrastructure

- On all the EX9200 line of switches, MX Series routers, and T4000, LACP DDoS policer is incorrectly triggered by other protocols traffic. [PR1409626](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- Packet drops, replication failure, or ksyncd crash might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- The RPM udp-ping probe does not work in multiple routing instance scenario. [PR1442157](#)

- With CNH for 6PE, MPLS EXP rewrite rule for non-VPN IPv4 over MPLS traffic might not work. [PR1430878](#)
- Traffic from the same physical interface cannot be forwarded. [PR1434933](#)
- The BGP session might flap after Routing Engine switchover is done simultaneously on both boxes of BGP peer in scaled BGP session setup. [PR1437257](#)
- GRE traffic might get dropped if the terminating routing-instance name contains dots. [PR1437872](#)
- ARP resolution might fail after ARP HOLD net hops are added and deleted continuously. [PR1442815](#)
- Some duplicate flowtap filters are programmed after the restart of dynamic-flow-capture. [PR1442868](#)
- When host-bound packet is received in MAP-E BR router, service interface statistics counter shows incorrect number of bytes. [PR1443204](#)
- Packets drop due to missing destination MAC address in the Packet Forwarding Engine. [PR1445191](#)
- Python op scripts are executed as user nobody if started from NETCONF session, not as logged in user, resulting in failing PyEZ connection to the device. [PR1445917](#)
- On certain MPC line cards cm errors need to be reclassified. [PR1449427](#)
- Some hosts behind unnumbered interface are unreachable after the router or FPC restarts. [PR1449615](#)
- FPC might reboot with vmcore due to memory leak. [PR1449664](#)
- The DF flag BGP packets are dropped over MPLS LSP path. [PR1449929](#)
- REST API process becomes non-responsive when a number of requests come in at a high rate. [PR1449987](#)
- In EVPN-VXLAN scenario, sometimes host-generated packets gets dropped as hitting reject route in Packet Forwarding Engine. [PR1451559](#)
- The Routing Engine originated IPv6 packets might be dropped when interface-group rule is configured under IPv6 filter [PR1453649](#)
- Multicast traffic loss occurs in rare case in a seamless MPLS with MVPN configuration is observed. [PR1456905](#)
- Port mirroring does not occur with VPLS. [PR1458856](#)
- DDoS-protection does not stop logging when remote tracing is enabled. [PR1459605](#)
- Traceroute initiated from PE device does not show the tunnel endpoint hop in the output. [PR1461441](#)
- CLI configuration flag **version-03** must be optional. [PR1462186](#)

Routing Policy and Firewall Filters

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

Routing Protocols

- The rpd crashes in Junos OS Release 16.1 or later during BGP convergence. [PR1351639](#)
- The rpd process might crash with BGP multipath and damping configured. [PR1472671](#)
- Need to install all possible next hops for OSPF network LSAs. [PR1463535](#)
- The **other querier present interval** timer cannot be changed in a IGMP or MLD snooping scenario. [PR1461590](#)
- BGP peers might flap if the parameter of hold-time is set as small. [PR1466709](#)
- The rpd might crash under a rare condition if GR helper mode is triggered. [PR1382892](#)
- The rpd crashes in Junos OS 16.1 or later during BGP convergence. [PR1351639](#)
- BFD link failure detection of the broken path gets delayed when IGP link-state update is received from the same peer through an alternative path. [PR1410021](#)
- BGP might become stuck in the Idle state when the peer triggers a GR restart event. [PR1412538](#)
- BGP might get stuck in the Idle state when the peer triggers a GR restart event. [PR1412538](#)
- TI-LFA cannot find backup path when IS-IS overLoad bit is set on computing. node [PR1412923](#)
- Per-prefix LFA might not work as expected where the last hop needs to be protected on the penultimate node. [PR1432615](#)
- Unsupported configuration (EPE with dynamic-next-hop GRE tunnels) continuously rpd to generat core files. [PR1431536](#)
- The **show isis adjacency extensive** output does not contain the state transition details. [PR1432398](#)
- The next-hop of IPv6 route remains empty when a new IS-IS link comes up. [PR1430581](#)
- With SR enabled, 6PE next hop is not installed. [PR1435298](#)
- Clearing BGP neighbors takes longer time to delete routes. [PR1435466](#)
- Wrong next hop might be seen when BGP PIC edge is enabled. [PR1437108](#)
- The rpd might crash in case multipath is enabled, as BGP multipath teardown is called for secondary route even though secondary routes are considered for multipath. [PR1437837](#)
- The backup Routing Engine might go out of synchronization if you clear BGP sessions on the master Routing Engine. [PR1439620](#)
- Removing SSH Protocol version 1 from configuration. [PR1440476](#)
- RIP routes might be discarded by Juniper device over a /31 subnet interface. [PR1441452](#)
- The rpd might crash with SR-TE configuration change. [PR1442952](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- The rpd crash might be seen after configuring OSPF nssa area-range and summaries. [PR1444728](#)
- The rpd might crash in OSPF scenario due to invalid memory access. [PR1445078](#)

- The SSH login might fail if a user account exists in both the local database and RADIUS/TACACS+. [PR1454177](#)
- MoFRR with MLDP inband signaling is not working. [PR1454199](#)
- BRP: RPC call is not available for **show bgp output-scheduler**. [PR1445854](#)
- The BGP route prefixes are not being advertised to the peer. [PR1446383](#)
- The as-external route might not work in OSPF overload scenario for VRF instance. [PR1446437](#)
- The rpd uses full CPU utilization due to incorrect path selection. [PR1446861](#)
- The multicast traffic might be dropped in PIM with BGP PIC setup. [PR1447187](#)
- The rpd crashes and commit fails when trying to commit configuration changes. [PR1447595](#)
- On the MX2000 and PTX10000 lines of devices , Layer 3 VPN PE-CE link protection exhibits unexpected behavior. [PR1447601](#)
- Junos OS BFD sessions with authentication flaps occurs after sometime. [PR1448649](#)
- The connection between ppmdd (Routing Engine) and ppmann (FPC) might get lost due to session time out. [PR1448670](#)
- The BGP routes might fail to be installed in routing instance if the **from next-hop** policy match condition is used in the VRF import policy. [PR1449458](#)
- SPRING-LDP interoperability issues are observed with colocated SRMS+SR-client+LDP-stitching. [PR1452956](#)
- The rpd scheduler slip for BGP GR might be up to 120 second after the peer goes down. [PR1454198](#)
- The rpd memory might leak in a certain MSDP scenario. [PR1454244](#)
- Permanent rpd core files are seen with BGP configuration option **optimal-route-reflection** set. [PR1454803](#)
- Rpd might crash when multipath is in use. [PR1454951](#)
- The rpd might crash continuously due to memory corruption in IS-IS setup. [PR1455432](#)
- Prefix SID conflict might be observed in IS-IS. [PR1455994](#)
- Packet drop and CPU spike on Routing Engine might be seen in certain conditions if **labeled-unicast protection** is enabled for a CsC-VRF peer. [PR1456260](#)
- Rpd core file is seen at **rt_nhn_tree_stop,rt_table_tree_free_family,bgp_sync_free_tsp** after deactivating protocols. [PR1457358](#)
- The rpd might crash when OSPF router-id gets changed for NSSA with area-range configured. [PR1459080](#)
- The rpd memory leak might be observed on backup Routing Engine due to BGP flap. [PR1459384](#)
- Rpd scheduler slips might be seen on RPKI route validation enabled BGP peering router in a scaled setup. [PR1461602](#)
- Rpd core file is seen with BMP configured and BGP peer flapping. [PR1462441](#)

- IS-IS IPv6 multi-topology routes might flap every time when there is an unrelated commit under protocol stanza. [PR1463650](#)
- The rpd might crash if both BGP add-path and BGP multipath are enabled. [PR1463673](#)
- MX80 EVPN-VXLAN RT5 does not work properly and ip-prefix-routes are not reachable. [PR1466602](#)

Services Applications

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

Subscriber Access Management

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

User Interface and Configuration

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

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the Layer 2 circuit flaps repeatedly. [PR1282875](#)
- The rpd core file is seen at `rtbit_reset`, `rte_tgtexport_rth`. [PR1379621](#)
- The rpd crash might be seen if Layer 2 circuit or local switching connections flap continuously. [PR1418870](#)
- P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. [PR1432434](#)
- The resumed multicast traffic for certain groups might be stopped in overlapping MVPN scenario. [PR1441099](#)
- Result of the **show task replication** command shows MVPN as InProgress when the active master Routing Engine is forcibly removed and NSR are enabled. [PR1441292](#)
- Memory leak might happen if PIM messages are received over an MDT (mt- interface) in Draft-Rosen MVPN scenario. [PR1442054](#)
- The rpd process might crash due to memory leak in **MVPN RPF Src PE** block. [PR1460625](#)
- The Layer 2 circuit displays MM status which might cause traffic loss. [PR1462583](#)

SEE ALSO

[What's New | 85](#)

[What's Changed | 106](#)

[Known Limitations | 111](#)

[Open Issues | 114](#)

[Documentation Updates | 152](#)

[Migration, Upgrade, and Downgrade Instructions | 153](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 153](#)

This section lists the errata and changes in Junos OS Release 19.4R1 documentation for MX Series.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

[What's New | 85](#)

[What's Changed | 106](#)

[Known Limitations | 111](#)

[Open Issues | 114](#)

[Resolved Issues | 131](#)

[Migration, Upgrade, and Downgrade Instructions | 153](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.4 | 154](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 154](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 157](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 159](#)
- [Upgrading a Router with Redundant Routing Engines | 159](#)
- [Downgrading from Release 19.4 | 160](#)

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

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

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

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

Basic Procedure for Upgrading to Release 19.4

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

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

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

For more information about the installation process, see [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).

Procedure to Upgrade to FreeBSD 11.x based Junos OS

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install FreeBSD 11.x based Junos OS:

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

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

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-32-19.4R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-19.4R1.9-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-x86-32-19.4R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.4R1.9-limited.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname**

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.x) to Junos OS (FreeBSD 11.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 11.x, and Junos OS (FreeBSD 6.x) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

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

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 19.4R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following MX Series routers:
 - MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008

[See <https://kb.juniper.net/TSB17603>.]

NOTE: After you install a Junos OS Release 19.4 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. See the VM Host Software Administrative Commands in the [Installation and Upgrade Guide](#).

Procedure to Upgrade to FreeBSD 6.x based Junos OS

Products impacted: MX5, MX10, MX40, MX80, MX104.

To download and install FreeBSD 6.x based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.

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

- All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/jinstall-ppc-19.4R1.9-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-19.4R1.9-limited-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname**

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

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

Rebooting occurs only if the upgrade is successful.

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform the following Junos OS installation on each Routing Engine separately to avoid disrupting network operation:

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

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

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

Downgrading from Release 19.4

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

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

What's New 85
What's Changed 106
Known Limitations 111
Resolved Issues 131
Open Issues 114
Documentation Updates 152

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 161](#)
- [What's Changed | 163](#)
- [Known Limitations | 164](#)
- [Open Issues | 165](#)
- [Resolved Issues | 168](#)

- Documentation Updates | 172
- Migration, Upgrade, and Downgrade Instructions | 173

These release notes accompany Junos OS Release 19.4R1 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- General routing | 162
- Hardware | 162
- Architecture | 162

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

General routing

- **Support for MAP-E customer edge encapsulation and decapsulation (NFX Series)**—Starting in Junos OS release 19.4R1, Mapping of Address and Port with Encapsulation (MAP-E) customer edge (CE) encapsulation and decapsulation are supported on NFX Series devices. MAP-E is an IPV6 transition technique that encapsulates an IPv4 packet in an IPv6 and carries the packet over IPv4-over-IPv6 tunnel from MAP-E CE devices to the MAP-E provider edge (PE) devices (also called as border relay [BR] devices) through an IPv6 routing topology, where the packet is de-tunneled for further processing.

MAP-E uses network address port translation (NAPT) features for restricting transport protocol ports, Internet Control Message Protocol (ICMP) identifiers, and fragment identifiers to the configured port sets. Existing NAPT feature is enhanced to add this capability.

[See [How to Configure the NFX150.](#)]

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

Hardware

- **NFX350 Platform**— With Junos OS Release 19.4R1, the NFX portfolio introduces the NFX350 Network Services Platform, which is a secure, automated, software-driven customer premises equipment (CPE) platform that delivers virtualized network and security services on demand. The NFX350 is part of the Juniper Cloud CPE solution, which leverages Network Functions Virtualization (NFV). The NFX350 platform completes the uCPE portfolio to provide end-to-end platforms for medium, large, and extra-large deployments. In addition to IPsec, Layer 2 features, and SD-WAN functionality, the NFX350 provides features such as LAN or WAN isolation, software and hardware resiliency, redundant power supply, and serial over LAN. The NFX350 device supports two external SSD and LTE expansion module.

The NFX350 devices are available in the following variants:

- **NFX350-S1**—Rack-mount model with 8-core Intel Skylake D-2146NT CPU, 100-GB SSD, 32-GB RAM, eight 1-Gigabit Ethernet RJ-45 LAN ports, and eight 10-Gigabit Ethernet SFP+ WAN ports.
- **NFX350-S2**—Rack-mount model with 12-core Intel Skylake D-2166NT CPU, 100-GB SSD, 64-GB RAM, eight 1-Gigabit Ethernet RJ-45 LAN ports, and eight 10-Gigabit Ethernet SFP+ WAN ports.
- **NFX350-S3**—Rack-mount model with 16-core Intel Skylake D-2187NT CPU, 100-GB SSD, 128-GB RAM, eight 1-Gigabit Ethernet RJ-45 LAN ports, and eight 10-Gigabit Ethernet SFP+ WAN ports.

[See [NFX350 Hardware Guide.](#)]

[See [How to Configure the NFX350.](#)]

Architecture

- **NFX350 Architecture**—The NFX350 architecture enables unified management of its components through the Junos Control Plane (JCP). It supports the following modes to effectively manage system resources:

- Throughput mode—Provides maximum resources (CPU and memory) for Junos software. The default mode is throughput mode.
- Hybrid mode—Provides a balanced distribution of resources between the Junos software and third-party VNFs.
- Compute mode—Provides minimal resources for Junos software and maximum resources for third-party VNFs

[See [NFX350 Hardware Guide](#).]

[See [How to Configure the NFX350](#).]

SEE ALSO

[What's Changed | 163](#)

[Known Limitations | 164](#)

[Open Issues | 165](#)

[Resolved Issues | 168](#)

[Documentation Updates | 172](#)

[Migration, Upgrade, and Downgrade Instructions | 173](#)

What's Changed

IN THIS SECTION

- [System Logging | 164](#)

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

System Logging

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

SEE ALSO

[What's New | 161](#)

[Known Limitations | 164](#)

[Open Issues | 165](#)

[Resolved Issues | 168](#)

[Documentation Updates | 172](#)

[Migration, Upgrade, and Downgrade Instructions | 173](#)

Known Limitations

IN THIS SECTION

• [Interfaces | 165](#)

• [Platform and Infrastructure | 165](#)

Learn about known limitations in Junos OS Release 19.4R1 for NFX Series devices. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Interfaces

- On an NFX150 device running Junos software release 19.4R1 or later, whenever there is no vmhost mapping configuration for a specific heth interface, then the **show interfaces extensive heth-<x>-<y>** command output displays the **Mapped to** status as **None** for **VF Number 0**. [PR1484388](#)
- On NFX150 devices, the link does not come up if a 1-Gigabit SFP transceiver is connected from heth-0-4 and heth-0-5 to a peer device. As a workaround, disable the auto-negotiation for the interface connected to the NFX150 on the remote device. [PR1428020](#)
- On NFX350 devices, when there is any change in vmhost interface mapping for an FPC, that FPC is restarted to activate the changed mapping. [PR1471694](#)

Platform and Infrastructure

- On NFX150 devices, random RPM probe losses are noticed if the probe packets are fragmented because the data-size more than the inet MTU. [PR1447082](#)

SEE ALSO

What's New 161
What's Changed 163
Open Issues 165
Resolved Issues 168
Documentation Updates 172
Migration, Upgrade, and Downgrade Instructions 173

Open Issues

IN THIS SECTION

- [Mapping of Address and Port with Encapsulation \(MAP-E\) | 166](#)
- [Interfaces | 166](#)
- [Platform and Infrastructure | 166](#)
- [Virtual Network Functions \(VNFs\) | 167](#)

Learn about open issues in Junos OS Release 19.4R1 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.

Mapping of Address and Port with Encapsulation (MAP-E)

- On NFX Series devices, IP identification (IP ID) is not changed after MAP-E NAT44 is performed on fragment packets when the packets reach the customer edge (CE) device. To avoid this issue, you can configure the Border Relay (BR) device as follows:

```
user@host# set services software software-concentrator map-e mape-domain-1 v4-reassembly
user@host# set services software software-concentrator map-e mape-domain-1 v6-reassembly
```

[PR1478037](#)

Interfaces

- When you issue a **show interface** command on NFX150 devices to check the interface details, the system will not check whether the interface name provided is valid or invalid. The system will not generate an error message if the interface name is invalid. [PR1306191](#)
- When a DHCP server assigns a conflicting IP address to the NFX device interfaces, the NFX device will not send a **DHCP DECLINE** message in response. [PR1398935](#)
- On NFX150 and NFX250 NextGen devices, when you add, modify, or delete a VNF interface that is mapped to an L2 or L3 data plane, kernel traces might be observed on the NFX Series device console. [PR1435361](#)
- Only ge-1/0/1 is mapped to OVS by default and 8 logical interfaces are created on this. [PR1452743](#)
- On NFX Series devices, the static MAC address is replaced by random MAC address. [PR1458554](#)
- On NFX350 devices, if you delete and add SXE interfaces, the SXE interface moves to Spanning Tree Protocol blocking (STP BLK) state, and the traffic drops on that interface. [PR1475854](#)
- On NFX350 devices, sxe interfaces take more time to come online after fpc0 restart via CLI or after device reboot. [PR1465436](#)

Platform and Infrastructure

- If you plug an unsupported SFP-T transceiver into an NFX150 device and reboot the device, the FPC1 WAN port does not come online. [PR1411851](#)
- Jumbo frames are not supported through OVS on an NFX250 device. [PR1420630](#)
- On NFX250 devices, Virtual Port Peer (VPP) is not running on dual CPE and occasionally on single CPE. [PR1461238](#)

- After upgrading the NFX Series devices from Junos OS Release 15.1X53-D47.4 to Junos OS Release 18.4R1, upgrades revert the file to default and JDM subsystem becomes unavailable. [PR1456900](#)
- On NFX350 devices, if you execute the **show vmhost mode** command multiple times, JDM may crash and cause the show commands to stop working. [PR1474220](#)
- On NFX350 devices, secure boot is not enabled after BIOS recovery using SSD and USB. [PR1480165](#)
- On NFX350 devices, an srxpfe core file is generated when VF mapping to srxpfe changes. When mapping the backplane's NIC changes for FPC1 to a VF, the srxpfe restarts. In NFX350 devices, the internal NICs are Intel NICs and the DPDK library in srxpfe is unable to handle the PF reset event generated during the remapping. This causes the srxpfe to crash just before the restart. There is no impact on functionality as a result of this issue; however, graceful restart doesn't happen, and instead the srxpfe generates a core file. [PR1469201](#)
- With SRX1500 device used as HUB and NFX350 device as SPOKE, **IPSEC replay-errors** are seen with HTTP traffic when the APPQOE passive probing is enabled. As a workaround, use SRX4200 as HUB. [PR1461068](#)
- When the AppQoe passive probing is enabled on NFX350 devices, packet drops are seen with the **First path drop: Policy check failed** message. [PR1473810](#)
- On NFX350 devices, the LCM peer connection un-stable alarm is raised during the device reboot, and the alarm clears after 2-3 minutes. [PR1473531](#)

Virtual Network Functions (VNFs)

- On NFX150 and NFX250 NextGen devices, when two srxpfe interfaces are mapped to the same physical interface and if you delete the interface mapping to VF0, the traffic flow is disrupted. Even though the mapping is moved to VF0, the MAC address is not cleared in VF1, which disrupts the traffic. As a workaround, reboot the device, which resets the MAC address to the default value. [PR1448595](#)
- On NFX350 devices, management daemon (MGD) is broken when cross-connect is committed before VNF is created. As a workaround, create VNF before creating cross-connect. [PR1480740](#)

SEE ALSO

[What's New | 161](#)

[What's Changed | 163](#)

[Known Limitations | 164](#)

[Resolved Issues | 168](#)

[Documentation Updates | 172](#)

[Migration, Upgrade, and Downgrade Instructions | 173](#)

Resolved Issues

IN THIS SECTION

- Class of Service | 169
- High Availability | 169
- Interfaces | 169
- Layer 2 Ethernet Services | 169
- Platform and Infrastructure | 169
- Routing Protocols | 171
- SNMP | 171
- Virtual Network Functions (VNFs) | 171

This section lists the issues fixed in the Junos OS main release and the maintenance releases for NFX Series devices.

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

Class of Service

- On NFX Series devices, when CoS rewrite rule is configured for st0 interface, the CoS value will not take effect on corresponding forwarding class. It causes the CoS not to work as expected. This issue has traffic impact. [PR1439401](#)

High Availability

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

Interfaces

- When you transition NFX150 devices from PPPoE configuration to non-PPPoE configuration in a non-promiscuous mode, the interface hangs without any traffic flow. [PR1409475](#)
- The limit on maximum OVS interfaces is restored to the originally defined limit of 25 for backward compatibility. As a workaround, reduce the number of OVS interfaces in the configuration to 20 or fewer. [PR1439950](#)
- On NFX150 and NFX250 NextGen devices, cross-connect stays down even if all linked interfaces are up. [PR1443465](#)
- On NFX Series devices, ping is not working between the cross-connected interfaces with interface deny-forwarding configuration. [PR1442173](#)
- When traffic goes through vSRX3.0 platforms, core-dump files are generated and traffic is dropped. This issue might result in Packet Forwarding Engine being inactive and all interfaces being down. [PR1465132](#)

Layer 2 Ethernet Services

- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case the subscriber may need more time to get IP address assigned. The subscriber may remain in this state until it's lease expires if it had previously bound with the address in the option 50. [PR1435039](#)

Platform and Infrastructure

- On NFX Series devices, the HTTP traffic flow is created with a different routing instance when an APBR profile is configured with category and application in the same profile. [PR1447757](#)
- When applying firewall filters on lo0.0 on an NFX250 NextGen device, FPC0 disappears. [PR1448246](#)

- On NFX150 devices, the **show security dynamic-address** command does not work for port 3. [PR1448594](#)
- Half duplex configuration on 1G ports is not supported when autonegotiation is disabled. [PR1453911](#)
- Informational log message, **LIBCOS_COS_RETRIEVE_FROM_PVIDB: feature cos_fc_defaults num elems 4 rc 0** is displayed on the console when you commit after you configure AppQoS rule set. [PR1457328](#)
- REST API process will get non responsive when a number of requests start coming at a high rate. [PR1449987](#)
- On NFX Series devices, if there are any conditional groups, the l2cpd process might crash and generate a core dump when interfaces are flapping and the lldp neighbors are available. It might cause the dot1x process to fail and all the ports have a short interruption at the time of process crash. As a workaround, delete the conditional group in the device. [PR1431355](#)
- Packet drops, replication failure or ksyncd crashes might be seen on the logical system of a Junos OS device after Routing Engine switchover. [PR1427842](#)
- After upgrading the NFX Series devices to Junos OS Release 19.2R2-S1.4, the following commit warning is seen even though there is no configuration change under the **forwarding-options vxlan-overlay-load-balance** option:

```
# commit and-quit
re0:
[edit]
  'forwarding-options'
    warning: vxlan-overlay-load-balance configuration for forwarding options has
    been changed. A system reboot is mandatory. Please reboot *ALL* routing engines
    NOW. Continuing without a reboot might result in unexpected system behavior.
configuration check succeeds
re1:
configuration check succeeds
commit complete
re0:
commit complete
Exiting configuration mode
```

[PR1459833](#)

Routing Protocols

- On NFX Series devices, changing the **other querier present interval** timer is not working on IGMP or MLD snooping device in the existing Bridge Domain (BD) or Listener Domain (LD). As a workaround, deactivate or activate the IGMP snooping via configuration or run the **restart multicast-snooping** command. [PR1461590](#)

SNMP

- On NFX150 devices, SNMP does not work for the following commands:
 - `show snmp mib walk jnxlpSecTunMonOutEncryptedBytes`
 - `show snmp mib walk jnxlpSecTunMonOutEncryptedPkts`
 - `show snmp mib walk jnxlpSecTunMonInDecryptedBytes`
 - `show snmp mib walk jnxlpSecTunMonInDecryptedPkts`
 - `show snmp mib walk jnxlpSecTunMonLocalGwAddr`
 - `show snmp mib walk jnxlpSecTunMonLocalGwAddrType`[PR1386894](#)
- Version compare in phc may fail causing the phc to download the same image. [PR1453535](#)

Virtual Network Functions (VNFs)

- On NFX150 devices with VNFs configured, when the VNF interfaces are moved from default OVS bridge to custom OVS bridge, there will be duplicate VNF host entries in the `/etc/hosts` file on JDM. [PR1434679](#)
- On NFX150 devices, when you need to change the vmhost mappings of a particular NIC or NICs, you must delete the existing vmhost mapping and commit the configuration. Now you can configure the new mappings for the respective NICs. You cannot change the NIC vmhost mappings in the same commit to delete and add a new mapping to the heth NICs. [PR1450147](#)
- NFX250 devices do not allow *jdm* (case-insensitive) as a VNF name. You can use *jdm* as a part of the name. For example, *jdm123*, *abcJDM*, *abcJDM123* are valid VNF names, whereas, *jdm*, *JDM*, *Jdm*, *JDm* are not valid VNF names. [PR1463963](#)

SEE ALSO

[What's New | 161](#)

[What's Changed | 163](#)

[Known Limitations | 164](#)[Open Issues | 165](#)[Documentation Updates | 172](#)[Migration, Upgrade, and Downgrade Instructions | 173](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 172](#)

This section lists the errata and changes in Junos OS Release 19.4R1 documentation for the NFX Series.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

[What's New | 161](#)[What's Changed | 163](#)[Known Limitations | 164](#)[Open Issues | 165](#)[Resolved Issues | 168](#)[Migration, Upgrade, and Downgrade Instructions | 173](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 173
- Basic Procedure for Upgrading to Release 19.4 | 173

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

Basic Procedure for Upgrading to Release 19.4

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

NOTE: NFX150, NFX250 NextGen, and NFX350 devices run VMhost supported routing engine, and should follow the [VMhost Support on Routing Engines](#) upgrade procedure.

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

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

To download and install Junos OS Release 19.4R1:

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 | 161](#)[What's Changed | 163](#)[Known Limitations | 164](#)[Open Issues | 165](#)[Resolved Issues | 168](#)[Documentation Updates | 172](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- [What's New | 176](#)
- [What's Changed | 184](#)
- [Known Limitations | 187](#)
- [Open Issues | 188](#)
- [Resolved Issues | 191](#)
- [Documentation Updates | 195](#)
- [Migration, Upgrade, and Downgrade Instructions | 196](#)

These release notes accompany Junos OS Release 19.4R1 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

- [General Routing | 177](#)
- [Hardware | 177](#)
- [High Availability \(HA\) and Resiliency | 177](#)
- [Junos OS, XML, API, and Scripting | 177](#)
- [Junos Telemetry Interface | 178](#)
- [MPLS | 180](#)
- [Routing Protocols | 181](#)
- [Services Applications | 182](#)
- [Software Defined Networking | 183](#)
- [System Logging | 183](#)

Learn about new features introduced in this release for PTX Series routers.

General Routing

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

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

Hardware

- **Support for 40-Gbps ports to operate at 10-Gbps speed (PTX10002-60C)**—You can use the Mellanox 10-Gbps pluggable adapter (model number: MAM1Q00A-QSA) to convert quad-lane-based ports to a single-lane-based SFP+ port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ module. Use the QSA adapter to convert a 40-Gbps port to a 10-Gbps port. You can plug a 10-Gbps SFP+ transceiver into the QSA adapter, which is then inserted into the QSFP or QSFP+ port of the PTX10002-60C router.

See [Supported Transceivers on PTX10002-60C](#).

High Availability (HA) and Resiliency

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

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

Junos OS, XML, API, and Scripting

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

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

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

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

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

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

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

Junos Telemetry Interface

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

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

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

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

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

- `/components/components/transceiver/`

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

- **Support for Segment Routing telemetry statistics and binding SIDs routes for uncolored Segment Routing Traffic Engineering policies (PTX1000, PTX3000, and PTX5000)**—Starting in Junos OS Release 19.4R1, Junos OS supports collection of traffic statistics for both ingress IP traffic and transit mpls traffic that take non-colored SR-TE paths on PTX Series routers. Binding SIDs for SRTE paths that have labels as first-hops in their segment lists are also now supported on PTX Series routers.

The `show spring-traffic-engineering lsp` command now has a `tunnel-source` filter, to display only the tunnels created from the specified sources by which the SRTE policy was provisioned. Also, the `show`

spring-traffic-engineering lsp detail command now displays information on the source of the tunnel configuration and statistics. By default, traffic sensors and statistic collection are disabled for static SRTE routes. To enable provisioning of JVISION traffic sensors in Junos OS data plane to stream out traffic statistics on SR policies and their Binding-SID routes, enable **statistics** under **telemetry** at the [**edit source-packet-routing telemetry**] hierarchy level, and sensors will be created for both the SRTE policy nexthop and Binding SID that are installed in the forwarding plane.

[See [source-packet-routing](#)]

MPLS

- **update-threshold statement modified to generate IGP update for lower bandwidth reservation (PTX Series)**—Starting in Junos OS Release 19.4R1, you can configure the threshold value of the **update-threshold** statement to accept:

- an integer or floating point values up to 3 significant digits after decimal point using the **threshold-percent** option
- an absolute value of bandwidth threshold which generates an IGP update using the **threshold-value** option

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

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

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

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

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

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

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

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

- **Support for segment routing features (PTX10002)**—Starting with Junos OS Release 19.4R1, PTX10002 router support the following segment routing features:

- BGP link-state distribution with SPRING extensions
- SRGB for SPRING in IS-IS domain
- Anycast and prefix segments in SPRING for IS-IS protocols
- IS-IS SPRING and RSVP coexistence
- Segment routing policy for traffic engineering on BGP
- Static adjacency segment identifier for ISIS and OSPF
- Static adjacency segment identifier for aggregate Ethernet member links
- Interoperability of segment routing with LDP
- RSVP-TE pop-and-forward LSP tunnels
- BGP Labeled Unicast traffic statistics collection
- Static segment routing label switched path
- Interoperability of segment routing with LDP
- Topology Independent Loop-Free Alternate for IS-IS and OSPF
- MPLS ping and traceroute for segment routing
- Anycast and prefix segments in SPRING for OSPF protocols
- Configurable SRGBs used by SPRING in OSPF protocols

[See [Link-State Distribution Using BGP Overview](#), [Understanding Adjacency Segments](#), [Anycast Segments](#), [and Configurable SRGB in SPRING](#), [BGP Egress Traffic Engineering](#), [Static Adjacency Segment Identifier for ISIS](#), [Static Adjacency Segment Identifier for OSPF](#), [IS-IS User Guide](#), [OSPF User Guide](#).]

Routing Protocols

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

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

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

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

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

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

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

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

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

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

Services Applications

- **Inline J-Flow scale enhancement (PTX10002)**—Starting in Junos OS Release 19.4R1, 100,000 flows per Packet Forwarding Engine are supported.

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

- **Support for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP inline flow monitoring (PTX10002-60C)**—Starting in Junos OS Release 19.4R1, you can perform inline flow monitoring for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP traffic. Both IPFIX and version 9 templates are supported.

[See [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers](#).]

- **MPLS-over-UDP inner payload flow monitoring with IPFIX and version 9 formats (PTX10002-60C)**—Starting in Junos OS Release 19.4R1, on the PTX10002-60C router, you can perform flow monitoring for MPLS-over-UDP traffic to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. This feature supports MPLS IPv4 and IPv6 payloads and both IPFIX and version 9 templates. Only ingress sampling is supported.

[See [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers](#).]

Software Defined Networking

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

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

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

System Logging

- **Improved intermodule communication between FFP and MGD (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 19.4R1, intermodule communication is improved to enhance software debugging. To enhance error messages with more context, the exit conditions from libraries have been updated as follows:

- Additional information is now logged for MGD-FFP intermodule communication.
- Commit errors that previously were only shown onscreen are now logged.

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

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

SEE ALSO

[What's Changed | 184](#)

[Known Limitations | 187](#)

[Open Issues | 188](#)

[Resolved Issues | 191](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 196](#)

What's Changed

IN THIS SECTION

- General Routing | 184
- Interfaces and Chassis | 184
- Junos Telemetry Interface | 185
- Routing Protocols | 186
- Software-Defined Networking | 186
- System Logging | 186

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

General Routing

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

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

- **IPv6 address in the prefix TIEs displayed correctly (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.

Interfaces and Chassis

- **Updates to the show interfaces and show policer commands (PTX Series)**—Starting in Junos OS release 19.4R1, on PTX Series routers, when you issue the **show interfaces** command or the **show policer** command, the output does not display the default arp policer (**_default_arp_policer_**). In earlier releases, when you issue the **show interfaces** command or the **show policer** command, the output displays the default arp policer (**_default_arp_policer_**) though the default arp policer is not supported on PTX series routers.
- **Change in Fabric Error Handling Behavior (PTX10008, PTX10016, PTX5000 routers (with FPC3-PTX-U2, FPC3-PTX-U3 FPCs), QFX10008, QFX10016, and QFX10002 switches)**—Starting in Junos OS release 19.4R1, when the PFE encounters ECC errors or parity errors related to fabric which are fatal, major, or correctable minor errors, the interfaces on the PFE are disabled. You must reboot the FPC manually to recover from the error. If you still face an issue after rebooting the FPC, contact our Customer Service.

In earlier releases, when the PFE encounters any error (fatal, major, minor_correctable, minor_transient, and info), the errors were incorrectly classified as info and as a result, ignored.

Junos Telemetry Interface

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

[See [Guidelines for gRPC and gNMI Sensors.](#)]

Routing Protocols

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

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

Software-Defined Networking

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

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

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

System Logging

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

SEE ALSO

[What's New | 176](#)

[Known Limitations | 187](#)

[Open Issues | 188](#)

[Resolved Issues | 191](#)

Known Limitations

IN THIS SECTION

- General Routing | 187

Learn about known limitations in Junos OS Release 19.4R1 for PTX Series routers. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

General Routing

- On the PTX Platform with FPC Model FPC-PTX-P1-A or FPC2-PTX-P1A, you might encounter single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message gets reported: **Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002** The Junos OS Chassis Management Error handling detects such a condition, raises an Alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper support representative if the issue persists even after the FPC restarts. [PR1254415](#)
- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error.** [PR1268678](#)

- Traffic loss for more than 15 seconds is seen when 50 percent of the aggregated Ethernet links are brought down by restarting multiple FPCs. [PR1412578](#)
- Because of an issue in the BIOS:QFXS_SFP_00.32_02.01 version, when the watchdog is disabled, the device does not reboot. [PR1441963](#)
- Call trace is observed during image upgrade from WRL6 to WRL9. [PR1442017](#)
- For scaled MACs as per the current design, the learn rate is expected. [PR1473334](#)

SEE ALSO

[What's New | 176](#)

[What's Changed | 184](#)

[Open Issues | 188](#)

[Resolved Issues | 191](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 196](#)

Open Issues

IN THIS SECTION

- [General Routing | 189](#)
- [Infrastructure | 190](#)
- [Layer 2 Ethernet Services | 190](#)
- [MPLS | 190](#)
- [Routing Protocols | 191](#)

Learn about open issues in Junos OS Release 19.4R1 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

- Power budget values for a PTX5000 chassis, FPC, and PICs have been revised. For routers operating on limited power, this can change the point where alarms for power-over-budget or insufficient power are raised or cleared. [PR1216404](#)
- On PTX1000, PTX5000, and PTX10000, when using the outbound firewall filter with a syslog option, adding a commit warning that the host interface might stop sending packets. [PR1462634](#)
- The firewall counter for lo0 interface might not increase. As a workaround, set the lo0 filter family inet and family inet6 counters instead of filter family any. [PR1420560](#)
- Aggregated Ethernet Interface Egress/Output statistics are not be in synchronization with its aggregated members or ports Egress/Output statistics. [PR1459633](#)
- Memory leaks are expected in this release. [PR1438358](#)
- On PTX5000 and PTX3000 router with 15x100G and 96x10G PIC, the interface **bcm8238** line side amplitude setting is incorrect and might cause optic reliability issues. [PR1453217](#)
- Layer 3 traffic fragmentation fails without DF bit. [PR1459738](#)
- Statistic comparison between CLI and JVISION for queue fails as the buffers shows incorrect values. [PR1460246](#)
- Traffic fails with **gcm-aes-xpn-128** cipher when you perform an event. [PR1460254](#)
- After transient SIB voltage spikes, SIB 0 is moved to fault and off-lined with major alarm, then FHP processes and tried to restart the SIB 0 only. However, the restart of SIB is not supported on PTX1000 and hence, the box might stay in the black hole states. [PR1460406](#)
- Some steps in test scripts causes image to go in the bad state. This is due to bad image of 19.3 release. [PR1461832](#)
- On a PTX Series router with a third-generation FPC, an error message is displayed when the FPC goes online or offline. [PR1322491](#)
- On FPC P2 line card, interface might stay down after maintenance. The issue is observed on links connected to another vendors equipment. [PR1412126](#)
- Alarm action does not work for minor errors after the threshold is changed to 1. [PR1345154](#)
- You might not be able to stop the ZTP bootstrap process when a PTX10016 or PTX10008 router with many line cards is powered on with the factory-default configuration. [PR1369959](#)
- The em2 interface configuration causes the FPC to crash during initialization and the FPC does not come online. After deleting the em2 configuration and restarting the router, the FPC comes online. [PR1429212](#)
- Interface statistics does not get updated with port-mirroring. [PR1431607](#)
- There is an increase in traffic loss after a unified ISSU with InterAS Layer 3 VPN OptionB configuration. [PR1435578](#)

- The full-resolve tunnel uses chain composite next-hop to program tunnel composite next-hop. Since the chain composite next-hop is created from the resolver, it has to delete logic to save the IPC call to the kernel or Packet Forwarding Engine. If the full-resolve dynamic tunnel (in this case, IPoIP tunnel) is deleted and created within 10 seconds, it reclaims the old tunnel. Consequently, you can see the old statistics of the tunnel. [PR1444081](#)
- The XML output for colored routes displays <c> instead of colored routes. [PR1447156](#)
- The **show route** command does not give all the next-hop information in the case of multipath routes. [PR1458000](#)
- After injecting the errors on the FPC, all interrupts are not recorded. [PR1459367](#)
- While loading and unloading of the firewall configuration, cda-zh core is observed. [PR1467741](#)
- When tunnel-services are configured on a PIC, the optics measurements that subscribed via gRPC might not be streamed. [PR1468435](#)
- On PTX10001 router with Junos OS Release 19.4R1 image, IPv6 does not work. [PR1475673](#)
- mib2d generates core files while deleting channelized interfaces. [PR1479642](#)

Infrastructure

- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1. [PR1462986](#)

Layer 2 Ethernet Services

- In EVPN multi-homed ACTIVE-ACTIVE scenario when LACP is enable on PE-CE child member links and after recovering from a core-isolation on PE device, the PE-CE child member links might get stuck in the detached state if LACP synchronization-reset feature is enabled on CE device. The child links on the CE device might show lacp state as Collecting Distributing. However, on the PE, the lacp state might be detached. [PR1463791](#)

MPLS

- In RSVP LSP with loose or undefined path, the LSP might stay in a down state due to loop detection after the link in the path flaps. [PR1384929](#)
- Kernel might crash and device might restart. [PR1478806](#)

Routing Protocols

- Post IGP convergence backup IPoIP tunnel remains up. As a workaround, you must deactivate or activate dynamic tunnel. [PR1447153](#)

SEE ALSO

[What's New | 176](#)

[What's Changed | 184](#)

[Known Limitations | 187](#)

[Resolved Issues | 191](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 196](#)

Resolved Issues

IN THIS SECTION

- [Forwarding and Sampling | 192](#)
- [General Routing | 192](#)
- [Infrastructure | 194](#)
- [Interfaces and Chassis | 194](#)
- [Layer 2 Ethernet Services | 194](#)
- [MPLS | 194](#)
- [Platform and Infrastructure | 194](#)
- [Routing Protocols | 194](#)
- [VPNs | 195](#)

This section lists the issues fixed in Junos OS Release 19.4R1 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.

Forwarding and Sampling

- PFED core files are seen and MIB2D is reported as slow peer due to a Packet Forwarding Engine accounting issue. [PR1452363](#)

General Routing

- The agentd sensor transmits multiple interface telemetry statistics per FPC slot. [PR1392880](#)
- On PTX10000, the FPC might restart during run time. [PR1464119](#)
- On PTX platforms, reclassification policy applied on the route prefixes might not work. [PR1430028](#)
- The Layer 2 cpd process might crash and generate a core dump when interfaces flaps. [PR1431355](#)
- On the PTX1000 or PTX10002 devices, the PIC and interfaces might not come up after FPC reboot. [PR1441256](#)
- On the PTX3000 devices, if the IPLC card is present in the device when you perform the GRES operation, the IPLC card crashes. [PR1415145](#)
- On PTX3000 and PTX5000, PIC might restart if the temperature of QSFP optics is overheated. [PR1462987](#)
- Incorrect counter values are observed for the arrival rate and peak rate for DDoS commands. [PR1470385](#)
- The aggregated Ethernet interface does not have LACP enabled over the circuit cross-connect between R0 and R3. [PR1424553](#)
- After you reboot the FPC, an interface comes up. [PR1428307](#)
- On the PTX10000 devices that use the LC1105 line card, you might observe traffic loss. [PR1433300](#)
- On the PTX10002 devices, chassis alarm is not raised when a PEM is removed or power to the PEM is lost. [PR1439198](#)
- On the PTX Series devices, the CPU or an interface might become unresponsive on a particular 100-Gigabit port. [PR1440526](#)
- Interfaces on the PTX Series devices might not come up after the FPC restarts or a port flaps. [PR1442159](#)
- BCM FW needs to be upgraded to DE2E. [PR1445473](#)
- Receipt of a malformed packet for J-Flow sampling might create a FPC core file. [PR1445585](#)
- The option to use wildcard <*> is not available at the group level of the Junos CLI. [PR1445651](#)
- The jdncpd process might crash after the **show access-security router-advertisement-guard** command is issued. [PR1446034](#)
- Upon steering of underlay dynamic tunnel PNHs to a different set of ECMP NHs, the tunnel that shared the same PNH might send traffic with wrong VLAN. [PR1446132](#)

- On the PTX Series devices, if sFlow is configured on more than eight interfaces, egress sampling might stop working. [PR1448778](#)
- Currently, ISIS sends system host name instead of system ID in the OC paths in lsdbs or adjacency xpaths in periodic streaming and on change notification. [PR1449837](#)
- Interfaces might flap after deleting the interface disable configuration. [PR1450263](#)
- JNP10K-LC2101 FPC generates "Voltage Tolerance Exceeded" major alarm for EACHIP 2V5 sensors. [PR1451011](#)
- Firewall filter applied at the interface level does not work when entropy level is present in certain scenarios. [PR1452716](#)
- The FPC might crash when the severity of error is modified. [PR1453871](#)
- GRPC updates on_change does not work when performing delete operations. [PR1459038](#)
- On the PTX1000 devices, scaling with 5000 tunnels adds JENCAP error messages in log and drops traffic. [PR1459484](#)
- Traffic is on hold when the interface flaps interface flap after DRD automatically recovers. [PR1459698](#)
- The forwarding option is not present in the routing instance type. [PR1460181](#)
- Hardware failure in CB2-PTX causes traffic interruption. [PR1460992](#)
- IPv6 ping does not work between CE to CE in the Layer 3 VPN network. [PR1466659](#)
- Traffic loops for pure Layer 2 packets coming over EVPN tunnel with destination MAC matching IRB MAC. [PR1470990](#)

Infrastructure

- On all Junos OS VM-based platforms, the FPC might reboot if jlock hog occurs. [PR1439906](#)

Interfaces and Chassis

- Due to the an issue in DWDM media, if any LAG member interface flaps, the LAG/ae stop receiving the LACP RX packets and fails to come UP. The LAG interface can be recovered by disabling/enabling the LAG interface. [PR1429279](#)

Layer 2 Ethernet Services

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

MPLS

- On a PTX Series router, the transit packets might be dropped if an LSP is added or changed. [PR1447170](#)

Platform and Infrastructure

- The REST service might become nonresponsive when the REST API receives several continuous HTTP requests. [PR1449987](#)
- Packet drops, replication failure or ksyncd crash might be seen on the logical system of a Junos device after Routing Engine switchover. [PR1427842](#)

Routing Protocols

- PTX Series devices cannot intercept PIM BSR message. [PR1419124](#)
- The rpd might crash with a change in SRTE configuration. [PR1442952](#)
- SSH login might fail if a user account exists in both local database and RADIUS/TACACS+. [PR1454177](#)
- On the PTX1000 devices, the Layer 3 VPN PE-CE link protection exhibits unexpected behavior. [PR1447601](#)
- The **other querier present interval** timer cannot be changed in a IGMP/MLD snooping scenario. [PR1461590](#)

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](#)
- Memory leak might happen if PIM messages are received over an MDT (mt interface) in Draft-Rosen MVPN scenario. [PR1442054](#)

SEE ALSO

[What's New | 176](#)

[What's Changed | 184](#)

[Known Limitations | 187](#)

[Open Issues | 188](#)

[Documentation Updates | 195](#)

[Migration, Upgrade, and Downgrade Instructions | 196](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed as User Guides | 196](#)

This section lists the errata and changes in Junos OS Release 19.4R1 documentation for the PTX Series.

Feature Guides Are Renamed as User Guides

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

SEE ALSO

[What's New | 176](#)

[What's Changed | 184](#)

[Known Limitations | 187](#)

[Open Issues | 188](#)

[Resolved Issues | 191](#)

[Migration, Upgrade, and Downgrade Instructions | 196](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.4 | 196](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 199](#)
- [Upgrading a Router with Redundant Routing Engines | 200](#)

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

Basic Procedure for Upgrading to Release 19.4

When upgrading or downgrading Junos OS, use the **jinstall** package. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#). Use

other packages, such as the **jbundle** package, only when so instructed by a Juniper Networks support representative.

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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

To download and install Junos OS Release 19.4R1:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Click the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the router.

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

All customers except the customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-19.4R1.9.tgz
```

Customers in the Eurasian Customs Union (currently comprised of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-19.4R1.9-limited.tgz
```

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

- **/pathname**— For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname**

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

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

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 19.4R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following PTX Series routers:
 - PTX3000, PTX5000, PTX10016, PTX10008, and PTX10002-XX

[See <https://kb.juniper.net/TSB17603>.]

NOTE: After you install a Junos OS Release 19.2 **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.

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 176
What's Changed 184
Known Limitations 187
Open Issues 188
Resolved Issues 191
Migration, Upgrade, and Downgrade Instructions 196

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- What's New | 201
- What's Changed | 212
- Known Limitations | 215
- Open Issues | 217
- Resolved Issues | 224
- Documentation Updates | 232
- Migration, Upgrade, and Downgrade Instructions | 232

These release notes accompany Junos OS Release 19.4R1 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

- EVPN | 202
- General Routing | 206
- Interfaces and Chassis | 206
- Junos OS XML API and Scripting | 207
- Junos Telemetry Interface | 208
- Layer 2 Features | 208
- MPLS | 209
- Routing Protocols | 209
- Software Defined Networking (SDN) | 209
- System Logging | 210

- System Management | 211
- VLAN Infrastructure | 212

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

NOTE: The following QFX Series platforms are supported in Release 19.4R1: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5200-32CD, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.

Junos on White Box runs on Accton Edgecore AS7816-64X switches in this release. The software is based on Junos OS running on QFX5210 switches, so release-note items that apply to QFX5210 switches also apply to Junos on White Box.

EVPN

- **EVPN pure type-5 route support (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, you can configure pure type-5 routing in an EVPN-VXLAN environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which carries the MAC address of the sending switch and provides next-hop reachability for the prefix. To configure pure type-5 routing, include the **ip-prefix-routes advertise direct-nexthop** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level. To enable two-level equal-cost multipath (ECMP) next hops in an EVPN-VXLAN overlay network, you must also include the **overlay-ecmp** statement at the **[edit forwarding-options vxlan-routing]** hierarchy level.

[See [ip-prefix-routes](#).]

- **EVPN control plane and VXLAN data plane support (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, QFX5120-32C switches support EVPN-VXLAN. By using a Layer 3 IP-based underlay network coupled with an EVPN-VXLAN overlay network, you can place endpoints anywhere in the network and remain connected to the same logical Layer 2 network.

EVPN-VXLAN is commonly deployed over the following physical underlay architectures:

- A two-layer IP fabric that includes spine devices (Layer 3 VXLAN gateways) and leaf devices (Layer 2 VXLAN gateways). You can deploy EX4650 and QFX5120 switches as spine or leaf devices in this fabric.

- A one-layer IP fabric that includes leaf devices that function as both Layer 2 and Layer 3 VXLAN gateways. You can deploy EX4650 and QFX5120 switches as leaf nodes in this fabric.

[See [Understanding EVPN with VXLAN Data Encapsulation](#).]

- **Dynamic load balancing in an EVPN-VXLAN overlay network (QFX5200 and QFX5210)**—In Junos OS Releases before Release 19.4R1, QFX5200 and QFX5210 switches support a static load-balancing scheme based on destination MAC addresses. This scheme distributes traffic on a round-robin basis among virtual tunnel endpoints (VTEPs) in an EVPN-VXLAN overlay network.

Starting in Junos OS Release 19.4R1, QFX5200 and QFX5210 switches that function as leaf or spine devices in an EVPN-VXLAN overlay network (centrally-routed and edge-routed bridging overlays) support dynamic load balancing among different equal-cost VTEPs. When enabled, the dynamic load-balancing feature supersedes the static load-balancing feature. With the dynamic feature, traffic is hashed among equal-cost paths based on packet fields. We support this feature in the following use cases:

- A leaf device is multihomed to multiple spine devices.
- A host is multihomed to multiple leaf devices.

In both use cases, each multihomed physical, aggregated Ethernet, or logical interface is configured with an Ethernet segment identifier (ESI). Dynamic load balancing supports a maximum of 255 ESIs. If you exceed this maximum (for example, you configure 256 ESIs), traffic destined for the 256th ESI is flooded to the VLAN associated with the ESI.

The hashing takes place before a packet undergoes VXLAN encapsulation. We use these fields to load-balance traffic:

- Packets with an IP header:
 - IP header fields:
 - Source IP address
 - Destination IP address
 - Protocol
 - VLAN ID
 - Layer 4 (TCP and UDP) source and destination ports
- Packets with an MPLS/IP header:
 - Up to three top labels
 - IP header fields:
 - Source IP address
 - Destination IP address

- Layer 4 (TCP and UDP) source and destination ports
- Packets with a Layer 2 header only:
 - Source MAC address
 - Destination MAC address
 - VLAN ID

To enable dynamic load balancing, include the **vxlان-overlay-load-balance** configuration statement at the **[edit forwarding-options]** hierarchy level and restart your switch.

To further control the hashing input used by this feature, include the [enhanced-hash-key](#) configuration statement at the **[edit forwarding-options]** hierarchy level.

- **Assisted replication in data centers with EVPN-VXLAN overlay networks (QFX Series switches)**—Starting in Junos OS Release 19.4R1, QFX Series switches support assisted replication (AR) in data centers with EVPN-VXLAN networks to optimize replication of BUM traffic forwarded into the EVPN core. Instead of flooding BUM traffic using ingress replication to multiple remote virtual tunnel endpoints (VTEPs) for a VLAN or virtual network identifier (VNI), devices configured as AR leaf devices (also called AR clients) forward the traffic to an AR replicator device that can better handle the replication load. The AR replicator then replicates and forwards the traffic to the VXLAN overlay tunnels. For further optimization, you can configure AR with IGMP snooping.

Switches in the QFX10000 line can be AR replicators. Any QFX Series switches that support EVPN-VXLAN can be AR leaf devices.

[See [Assisted Replication Multicast Optimization in EVPN Networks](#).]

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

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

- **Features supported on EX4650 and QFX5120 switches**—Starting with Junos OS Release 19.4R1, the following Junos OS features are supported on EX4650 and QFX5120 switches:
 - Automatically generated Ethernet segment identifiers (ESIs) in EVPN-VXLAN and EVPN-MPLS networks.
[See [Understanding Automatically Generated and Assigned ESIs in EVPN Networks](#).]
 - Firewall filtering and policing on EVPN-VXLAN traffic.
[See [Understanding VXLANs](#) and [Overview of Firewall Filters](#).]
 - Graceful restart on EVPN-VXLAN.
[See [Graceful Restart in EVPN](#).]
 - IGMPv2 snooping for EVPN-VXLAN in a multihomed environment.
[See [Overview of IGMP Snooping in an EVPN-VXLAN Environment](#).]
 - IPv6 data traffic support through an EVPN-VXLAN overlay network.
[See [Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay](#).]
 - Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface.
[See [Understanding Flexible Ethernet Services Support with EVPN-VXLAN](#).]
 - MAC limiting, storm control, and port mirroring support in EVPN-VXLAN overlay networks.
[See [MAC Limiting, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment](#).]
 - Multihomed proxy advertisement.
[See [EVPN Multihoming Implementation](#).]
 - Selective multicast forwarding and SMET route support in EVPN-VXLAN.
[See [Overview of Selective Multicast Forwarding](#).]
 - Standard class-of-service (CoS) features—classifiers, rewrite rules, and schedulers—are supported on VXLAN interfaces.
[See [Understanding CoS on OVSD-Managed VXLAN Interfaces](#).]
 - VMTO for ingress traffic.
[See [Ingress Virtual Machine Traffic Optimization](#).]

General Routing

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

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

Interfaces and Chassis

- **Support for dynamic load balancing (QFX5120-32C and QFX5120-48Y)**—Starting in Junos OS Release 19.4R1, QFX5120-32C and QFX5120-48Y switches support dynamic load balancing (DLB) for ECMP and LAG. DLB is an enhancement to static load balancing. DLB considers member bandwidth utilization along with packet content for member selection.

You can use the following DLB modes to load-balance traffic:

- Flowlet
- Assigned flow
- Per-packet

To configure DLB for ECMP, include the **ecmp-dlb** statement at the **[edit forwarding-options enhanced-hash-key]** hierarchy level.

To configure DLB for LAG, include the **dlb** statement at the **[edit interfaces aex aggregated-ether-options]** hierarchy level.

NOTE: You cannot configure both DLB and resilient hashing at the same time. Otherwise, commit error will be thrown.

[See [Understanding Dynamic Load Balancing](#) and [Configuring Dynamic Load Balancing](#).]

- **Support for 10-Gbps speed using JNP-SFP-25G-DAC (QFX5120-48Y)**—Starting in Junos OS Release 19.4R1, you can use any of the following JNP-SFP-25G-DAC cables to set 10-Gbps speed on the SFP28 ports of a QFX5120-48Y switch:
 - JNP-SFP-25G-DAC-1M

- JNP-SFP-25G-DAC-3M
- JNP-SFP-25G-DAC-5M

If you've plugged a JNP-SFP-25G-DAC cable into a QFX5120-48Y switch, then the SFP28 ports come up with 10-Gbps speed by default. To configure the SFP28 ports to operate at 25-Gbps speed, you must explicitly configure the speed of the first port in the port group using the **set chassis fpc 0 pic 0 port *port-num* speed 25g** command.

[See [Channelizing Interfaces on QFX5120-48Y Switches](#).]

- **Support for 10-Gbps speed on JNP-SFPP-10GE-T transceiver (QFX5100-48S)**—Starting in Junos OS Release 19.4R1, QFX5100-48S switches support JNP-SFPP-10GE-T transceiver. This transceiver supports 10-Gbps speed by default.

Junos OS XML API and Scripting

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

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

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

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

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

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

Junos Telemetry Interface

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

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

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

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

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

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

Layer 2 Features

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

ERPS version 1 comprises the following features:

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

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

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

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

MPLS

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

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

Routing Protocols

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

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

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

Software Defined Networking (SDN)

- **OVSDB support with VMware NSX for vSphere (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, the Open vSwitch Database (OVSDB) management protocol provides a control plane through which an NSX controller can provision QFX5120-32C switches. In an environment in which

NSX Release 6.4.5 or later is deployed, an NSX controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and the reverse.

The physical underlay network over which OVSDB-VXLAN is commonly deployed is a two-layer IP fabric that includes spine and leaf devices. The spine devices function as Layer 3 VXLAN gateways, and the leaf devices function as Layer 2 VXLAN gateways. You can deploy QFX5120 switches as leaf devices in this fabric.

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

- **Layer 2 and Layer 3 VXLAN gateways (QFX5120-32C switches)**—Starting with Junos OS Release 19.4R1, you can deploy QFX5120-32C switches as follows:
 - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
 - As a Layer 2 VXLAN gateway in an OVSDB overlay network

VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See [Understanding VXLANs.](#)]

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

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

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

System Logging

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

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

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

System Management

- **Precision Time Protocol (PTP) transparent clock (QFX5120 and QFX5210)**—Starting in Junos OS Release 19.4R1, you can use a transparent clock to update the PTP packets with the residence time as the packets pass through the switch. There is no master/slave designation. The switches support end-to-end transparent clocks, which include only the residence time. The transparent clock can update the residence time in a one-step process, which means it sends the timestamps in one packet.

To use a transparent clock, enable the **e2e-transparent** statement at the **[edit protocols ptp]**.

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

- **Additional support for Bidirectional Forwarding Detection (QFX5110, QFX5120, QFX5200, and QFX5210)**—Starting in Junos OS Release 19.4R1, Bidirectional Forwarding Detection (BFD) can support sessions of less than 1-second intervals. Performance might vary depending on the configuration load within the system.

NOTE: IPv4 and standalone BFD sessions, as well as inline single-hop sessions are supported. Micro BFD implementation and logical router support are not supported.

[See [Understanding Bidirectional Forwarding Detection \(BFD\)](#).]

VLAN Infrastructure

- **Support for multiple Q-in-Q tags (QFX10000 switches)**—Starting in Junos OS Release 19.4R1, the QFX10000 line of switches support the third and fourth Q-in-Q tags as payload (also known as pass-through tag) along with the existing two tags (for VLAN matching and operations). The QFX10000 switches support multiple Q-in-Q tags for both layer 2 bridging and EVPN-VXLAN cases. The Layer 2 access interfaces accept packets with three or four tags (all tags with the TPID value 0x8100). All the tags beyond the fourth tag (that is, from the fifth tag onward) are considered part of the Layer 3 payload and are forwarded transparently.

NOTE: In a one or two tagged packet, the tags (tag 1 and tag 2) can carry any TPID values (0x8100, 0x88a8, 0x9100, and 0x9200).

[See [Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation](#).]

SEE ALSO

What's Changed 212
Known Limitations 215
Open Issues 217
Resolved Issues 224
Documentation Updates 232
Migration, Upgrade, and Downgrade Instructions 232

What's Changed

IN THIS SECTION

- [General Routing | 213](#)
- [Interfaces and Chassis | 213](#)
- [Junos Telemetry Interface | 213](#)
- [Management | 214](#)
- [Routing Protocols | 214](#)

- Software Defined Networking (SDN) | 214
- System Logging | 214

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

General Routing

- **Automatic installation of YANG-based CLI for RIFT protocol (MX Series, QFX Series, and vMX with 64-bit and x86-based servers)**—In RIFT 1.2 Release, installation of the CLI for RIFT protocol occurs automatically along with the installation of the junos-rift package. In the pre-1.0 releases of the junos-rift package, the RIFT CLI had to be installed separately using **request system yang** command after installation of the junos-rift package.
- **IPv6 address in the prefix TIEs displayed correctly (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.

Interfaces and Chassis

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

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

Junos Telemetry Interface

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

[See [Guidelines for gRPC and gNMI Sensors.](#)]

Management

- **entPhysicalTable fetched on QFX10002**—In Junos OS Release 19.4R1, the MIB data for entPhysicalTable will be fetched on a QFX10002-72Q or QFX10002-36Q switch.

[See [SNMP Explorer.](#)]

Routing Protocols

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

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

Software Defined Networking (SDN)

- **Increase in the maximum value of delegation-cleanup-timeout (QFX Series)**—You can now configure a maximum of 2147483647 seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

With the increase in maximum value of **delegation-cleanup-timeout** from 600 to 2147483647 seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

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

System Logging

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

SEE ALSO

What's New		201
Known Limitations		215
Open Issues		217
Resolved Issues		224
Documentation Updates		232
Migration, Upgrade, and Downgrade Instructions		232

Known Limitations

IN THIS SECTION

- [Layer 2 Ethernet Services](#) | [216](#)
- [Layer 2 Features](#) | [216](#)
- [Network Management and Monitoring](#) | [216](#)
- [Platform and Infrastructure](#) | [216](#)
- [Routing Protocols](#) | [217](#)

Learn about known limitations in Junos OS Release 19.4R1 for QFX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Layer 2 Ethernet Services

- In EVPN multihomed active/active scenario when LACP is enabled on PE-CE child member links, LACP force-up feature should not be enabled in conjunction with EVPN core isolation feature (enabled by default) because it is currently not supported in this scenario as these two features are contradictory in terms of action they take. [PR1461581](#)

Layer 2 Features

- The **Targeted-broadcast forward-only** command does not broadcast the traffic. [PR1359031](#)

Network Management and Monitoring

- The number of possible output interfaces in remote port mirroring varies among the various switches in the QFX5000 line of switches:
 - QFX5110, QFX5120, and QFX5210—Support a maximum of 4 output interfaces.
 - QFX5100 and QFX5200—Support a maximum of 3 output interfaces.

Platform and Infrastructure

- The chip has VLAN-based logical interface statistics. Since for a given logical interfaces on both IPv4 and IPv6 use the same VLAN, stats will count both V4 and V6 together. There is no way to separately count them. Hence, "IPv6 transit statistics" is always 0. However, the total transit statistics (IPv4 + IPv6) will be displayed under "Transit statistics". [PR1327811](#)
- VLAN is not deleted in the hardware on IRB disable earlier, leading to ARP getting refreshed even though IRB is disabled. [PR1421382](#)
- On QFX5110-32Q running Junos OS Release 18.1R1 and earlier, due to a platform limitation, the channelization of the ports should follow the following design recommendations:
 - With 100-gigabit transceivers connected in the port range 28–31, only ports 0–19 can be channelized in default system-mode.
 - If a 40-gigabit transceiver is connected in any of the 100G supported ports, only ports in the range 1–18 can be channelized in default system-mode.

- If all 32 ports have 40-gigabit transceivers connected, only ports in the range 1–18 can be channelized in default system-mode.
- In non-oversubscribed mode, all the valid ports (that is, 0-23) can be channelized as expected. [PR1438319](#)
- There is a limitation regarding this behavior because 500 million is not sufficient for `/var/rundb` if there is a scaled configuration, which usually keeps the history of the configuration causes this issue and it is mostly seen during rollback and commit with scaled configurations. As a workaround, clean up `/var/rundb` when it is full and then proceed with commit. [PR1452154](#)

Routing Protocols

- Targeted broadcast functionality with VXLAN is not supported yet on QFX5000 platforms. In case of a non-VXLAN case, broadcast destination IP lookup results in next hop with destination MAC of all 0xffs and gives the class-id for IFP to match and action to redirect to IPMC with vlan membership check. VXLAN case, L3 egress interface, egress L3 next hop, and ingress L3 entry creations are failing. [PR1397086](#)

SEE ALSO

What's New 201
What's Changed 212
Open Issues 217
Resolved Issues 224
Documentation Updates 232
Migration, Upgrade, and Downgrade Instructions 232

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 218](#)
- [EVPN | 218](#)
- [High Availability \(HA\) and Resiliency | 219](#)
- [Interfaces and Chassis | 219](#)
- [Junos Fusion for Provider Edge | 219](#)

- [Layer 2 Features | 219](#)
- [Layer 2 Ethernet Services | 220](#)
- [MPLS | 220](#)
- [Platform and Infrastructure | 220](#)
- [Routing Protocols | 223](#)

Learn about open issues in Junos OS Release 19.4R1 for QFX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- PFC feature will not be supported with QFX5120/EX4650 2-member Virtual Chassis currently due to BCM limitation. [PR1431895](#)
- On QFX5110, QFX5100 and EX4600 platforms, if **shaping-rate** is configured, the shaping feature might not work after a reboot. The service might be impacted as the traffic cannot be rate limited. [PR1432078](#)
- In Junos Fusion scenario, when traffic from AD (aggregation device) to SD (satellite device) is exported with different dscp marking, it might be changed into network-control queue on extended port of SD. [PR1433252](#)

EVPN

- OVSDB-managed QFX5100 or QFX5110 is encapsulating VXLAN traffic and sending to incorrect destination MAC address when multiple remote VTEPs are in the same subnet and reachable via IRB interface in stretched VLAN. [PR1424698](#)
- In an EVPN environment, proxy ARP and ARP suppression is enabled on the Provider Edge device by default for reducing the flooding of ARP packets; however, in the case of ARP probe packets used in the process of Duplicate Address Detection (DAD), the client may treat the IP address that it is in use as duplicated address after receiving the proxied packets from PE device. [PR1427109](#)
- In Ethernet Virtual Private Network - Virtual Extensible LAN (EVPN-VXLAN) core Isolation scenario, the server is multihomed to the leaf devices through Link Aggregation Control Protocol (LACP) interfaces. If GR (Graceful Restart) is enabled, upon system reboot or restart routing on the leaf device, the core Isolation will not work. In the system reboot case, the issue results in the leaf device being dropped silently the traffic sent from the server during the time window between LACP coming up and Border Gateway Protocol (BGP) coming up. In the restart routing case, there might be no traffic drop because of the GR. [PR1461795](#)

High Availability (HA) and Resiliency

- During unified ISSU from previous releases to Junos OS Release 19.4R1 for QFX5000 platforms, dc-pfe will crash continuously and hence ISSU will not work. [PR1472183](#)

Interfaces and Chassis

- Customers might notice the flooding of ARP reply unicast packets as a result of an ARP request sent for the device's VRRP MAC address. This should not cause major issues. The ARP reply that is flooded in the VLAN by the device has the correct DMAC of the originator of the ARP request. In other words, the ARP reply is flooded but with the correct unicast DMAC. The ARP reply is not broadcast. For example:
15:15:58.378813 In -----original packet----- e4:5d:37:5e:e0:40 > ff:ff:ff:ff:ff:ff, ethertype 802.1Q (0x8100), length 64: vlan 1123, p 0, ethertype ARP, arp who-has 10.110.59.158 tell 10.110.59.146
15:15:58.378854 Out -----original packet----- 00:00:5e:00:01:3b > e4:5d:37:5e:e0:40, ethertype 802.1Q (0x8100), length 50: vlan 1123, p 0, ethertype ARP, arp reply 10.110.59.158 is-at 00:00:5e:00:01:3b
[PR1454764](#)
- When dynamic DHCP sessions are existing in the device, if multiple commits in parallel are performed, the commit might hang up. [PR1470622](#)
- Commit error was not thrown when member link was added to multiple aggregation group with different interface specific options. When member interface added to bundle with both ether and gig-ether interface specific options, gig-ether option takes precedence over ether options. [PR1475634](#)

Junos Fusion for Provider Edge

- IGMP membership is not getting learned by the AD fully even when the IGMP queries is being sent out. [PR1419265](#)

Layer 2 Features

- In case of QFX5000 Virtual Chassis/VCF setups, when IGMP snooping is enabled, multicast traffic is forwarded based on IGMP joins/reports. But when the IGMP report times out, traffic should be dropped; instead it will be flooded in the VLAN. This happens only in case of QFX5000 Virtual Chassis/VCF; this issue is not seen on stand-alone QFX5000. [PR1431893](#)
- If Packet Forwarding Engine process is restarted manually for QFX5110-32Q, some of MAC address might not be seen on software MAC table in case of EVPN-VXLAN case even though MAC address will be present in hardware table. This is a timing issue and not always seen. This issue is seen only on QFX5110-32Q and not seen on any other platforms. [PR1467466](#)

Layer 2 Ethernet Services

- In EVPN multihomed active/active scenario when LACP is enable on PE-CE child member links and after recovering from a core-isolation on PE device, the PE-CE child member links might be stuck in DETACHED state if LACP sync-reset feature is enabled on CE device. The child links on the CE device might show lacp state as **Collecting Distributing**, but on the PE the lacp state might be "DETACHED". [PR1463791](#)

MPLS

- The `show mpls static-lsp | display xml` command produces INVALID XML when more than 100 static LSPs are configured. [PR1469378](#)

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as `nh_ucast_change:291Referenced l2ifl not found`. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- On a QFX Series switch with a third-generation FPC, the error message is displayed when the FPC goes online or offline. [PR1322491](#)
- QFX10000 platform drops the Aruba wireless access point (AP) heartbeat packets. As a result, the Aruba wireless AP cannot work. [PR1352805](#)
- User might not be able to stop the ZTP bootstrap, when a QFX10016/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](#)
- With MLD-snooping enabled and when we have two receivers in the same VLAN interested in the same group address but from a different source, traffic will be received only on the receiver which sent the latest MLD report. This is because we do not install S, G routes in hardware when MLD snooping is enabled. [PR1386440](#)
- On Junos OS Release 18.4R1 branch, Intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface. [PR1388082](#)
- On EX Series and QFX Series platforms, when bringing up clients (most likely in DHCP/PPP subscriber scenario), the subscribers might fail to bind. The reason is that when installing new software images, it might cause shared memory (created by previously running image) not to be cleared out. The issue will persist until the previous values in shared memory are removed, and the daemons affected by the data in shared memory might continue to generate core files and to crash and thus not be able to function properly. [PR1396470](#)
- On QFX5000 platforms with scaled setup of the aggregated Ethernet (ae) bundles and VLANs, if Link Aggregation Control Protocol (LACP) is enabled, and there are scaled configuration changes, for example,

delete 4000 VLANs/VXLANs and reapply them again, some interfaces of ae bundle might go to the detached state. Due to this issue, the running routing protocols (for example, LACP and BGP) will go down over the affected ae bundles. [PR1406691](#)

- On QFX5110 and QFX5120 platforms, uRPF check in strict mode might not work properly. [PR1417546](#)
- On QFX Series Virtual Chassis during shutdown, if an interrupt is received, the system gets into this state and vmcore is observed. [PR1421250](#)
- LLDP frames received on a QFX5210 management em0 port might not show in show LLDP operational queries. Other non-em0 interfaces will display statistics. [PR1426753](#)
- Hardware is not getting programmed on rebooting QFX5000 node for CoS rewrite on aggregated Ethernet (LAG) interfaces. [PR1430173](#)
- When routing process is restarted, if system is configured with EVPN service, memory of I2 learning daemon is increasing by 4000 when you use **show system processes extensive | match I2ald**. [PR1435561](#)
- The time it takes to install or delete IPv4/Pv6 routes into the FIB is slowed down in Junos OS Release 19.3. Analysis shows that rpd learning rates are not degraded, but RIB to FIB download rate is degraded. [PR1441737](#)
- During unified ISSU upgrade process sometime if ISSU does not complete or state machine does not move forward, ISSU process is forced to time out and abort. during abort, ISSU states are cleared and ISSU is aborted . This is one of the corner cases where ISSU is not progressed and ISSU. After any such abort, it is good to have reboot of the system as recommended by ISSU console logs, which is also a workaround. [PR1443342](#)
- On QFX10000 platforms and EVPN-VXLAN (spine-leaf) scenario, the QFX10000 spine switches are configured with VXLAN Layer3 gateway (utilizing the virtual-gateway) on an IRB interface, if enabling and then subsequently remove the VXLAN L3 gateway on this IRB interface on one or some of these spine switches, traffic drop might be observed. If all virtual-gateways are configured with a unique v4 or v6 mac-address, this issue would not happen. This is also the workaround. [PR1446291](#)
- When all the members of QFX5120-32C/QFX5120-48T Virtual Chassis are rebooted while traffic is running, for a quite period of time (~15 minutes) backup will be disconnected from master as the master-backup socket connection will be down and will be re-established after 15 minutes and backup will join the Virtual Chassis. [PR1453399](#)
- QFX5200-32c-32q : vmcore occurred at
/amd/svl-engdata1vs1/occamdev/build/freebsd/stable_11/20190614.234225
 __ci_fbsd_builder_stable_11.0.269d466/src/sys/kern/kern_shutdown.c:313 after upgrade from 18.3Throttle image to Junos OS 19.3R1. [PR1455851](#)
- When traceoptions is enabled for OSPF, rpd might crash if the interface cost is changed. [PR1456054](#)
- In VXLAN setup containing very large number of child interfaces, significant link up delay was seen when one of FPCs in QFX5100 Virtual Chassis is rebooted. [PR1456336](#)
- Show dynamic-tunnels database does not show **v6 mapped** next-hop flag for the 6PE routes that have labels. It is just a display issue. [PR1458634](#)

- On QFX5100, when unified ISSU is performed with Layer3 protocols configured then traffic loss of 0.8 seconds is observed. [PR1459701](#)
- A libvirtMib_suba core file might be seen after an image upgrade from Junos OS Release 17.3R3-S5.2 to 17.3R3-S6.3 on a QFX5110 device. There is no functional impact due to this core file generation. [PR1462725](#)
- BGP route addition and deletion time increased in Junos OS Release 19.4. [PR1464572](#)
- A few of DHCPvX INFORM Messages, specific to particular VLAN are not receiving any ACK from server. [PR1467182](#)
- When **tunnel-services** are configured on a PIC, the optics measurements that subscribed via gRPC might not be streamed. [PR1468435](#)
- In Junos Fusion Data Center environment, when a VM is moved from one satellite port to another using VMotion, MAC address of VM might not move to new satellite port in Aggregate Device's switching table. [PR1468732](#)
- If system has 1000 BGP-V4 VRF (120000 routes) + 700 OSPF-V2 VRF (70000 routes) + 300 P2P ISIS V4 VRF (30000 routes) (nearly 220000 hardware routes), then deleting/reading VRF configurations might lead to all BGP sessions going down. [PR1469881](#)
- On QFX5100 and EX4300 mixed-mode Virtual Chassis, the speed 10m might not be configured on the GE interface. [PR1471216](#)
- In VXLAN scenario on QFX1000 series platforms, when VTEP source interface is configured in multiple routing instances, the traffic loss might occur if one of such routing instances is deleted. [PR1471465](#)
- Because of change in new SDK, egress filter's slice usage became double and there will be only 512 entries in Junos OS Release 19.4R1. [PR1472206](#)
- On QFX platforms, the shaping of CoS does not work after reboot when the shaping rate is configured with an absolute value. The line rate of traffic is sent out, no shaping occurs. This issue has traffic or service impact. [PR1472223](#)
- In EVPN-VXLAN scenario, when an SP style interface is configured both with **native-vlan-id** and LLDP on QFX5000 platforms, continuous log messages might be observed. [PR1474545](#)
- QFX5000 Leaf device might fail to forward the traffic to AR-replicator/Spines, on flapping bgp neighbors on AR-replicator/Spines. [PR1475430](#)
- Interfaces are not detected on some of the ports when we swap the 25g sfps and insert 10g sfp. [PR1475574](#)
- The commit synchronize command fails because the kernel socket gets stuck. [PR1177692](#)
- On the QFX10002-60C, filter operation with log action is not supported for protocols other than Layer 2, IPv4, and IPv6. The following message is seen in firewall logs: **Protocol 0 not recognized**. [PR1325437](#)

Routing Protocols

- QFX5110, used as VRRP peers, in some specific scenarios involving configuration of bpdu-block-on-edge might claim to be VRRP masters. [PR1367439](#)
- Value added in Hexa after Unknown Ext-Community is getting reset to 0. [PR1371448](#)
- On QFX-5100 VC/VCF, the following error is observed: **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:13 nh 6594 unintsall failed** in hardware with mini-PDT base configurations. There is no functionality impact because of this error message. [PR1407175](#)
- On QFX5000/EX4600 platforms, the received traffic will be dropped if the destination UDP port is 521. [PR1429543](#)
- In EVPN/VXLAN scenario with QFX5000 platform, all BGP sessions might go into down status if restarting overlay/underlay BGP session while traffic flow starts. [PR1431259](#)
- QFX5110 MCLAG: **L2_L3_INTF_OPS_ERROR** messages are seen after node reboot. [PR1435314](#)
- For Junos OS Release 19.3R1: On QFX5100 , when unified ISSU is performed, then traffic loss of 15-20 seconds is observed. [PR1449581](#)
- In an EVPN-VXLAN multihomed environment with QFX5110/QFX5120 acting as leaf devices, if IGMP snooping is used, it might override the local bias filters on designated forwarder (DF) and nondesignated forwarder (NDF) devices, and forwards the packets, causing multicast packets loops. [PR1457725](#)
- Example case: Send flow with interburst gap of 1000 micro seconds.
 - 1 Configure a flowlet timer value of 16 microseconds.
 - 2 The flows should split and the flows should be distributed among all the links in the ECMP.
 - 3 Now change the inactivity-timer to 10000 micro seconds.
 - 4 The flows should again fall back to single link because the inactivity timer is greater the interburst gap. The aforementioned steps list is the expected behavior in case of above timer value and interburst gap or IFG configured in the flows. But there may be cases when the above values are changed multiple times, the behavior can be unexpected, and flows will not move back to single link even though the inactivity-timer configured is more than the interburst gap or IFG. Trigger: Change the inactivity timer value from less (less than IFG or interburst gap) to more (more than IFG or interburst gap) multiple times. Expected behavior: Flows will move to single link when inactivity-timer is more than IFG. Current behavior: Flows will not move to single link when inactivity-timer is more than IFG. Recovery: Restart dc-pfe. [PR1471729](#)

SEE ALSO

[What's New | 201](#)

[What's Changed | 212](#)

[Known Limitations | 215](#)[Resolved Issues | 224](#)[Documentation Updates | 232](#)[Migration, Upgrade, and Downgrade Instructions | 232](#)

Resolved Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 224](#)
- [EVPN | 225](#)
- [Forwarding and Sampling | 225](#)
- [Interfaces and Chassis | 225](#)
- [Layer 2 Features | 225](#)
- [MPLS | 226](#)
- [Platform and Infrastructure | 226](#)
- [Routing Protocols | 230](#)
- [User Interface and Configuration | 231](#)

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

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

Class of Service (CoS)

- QFX10008: FPC0 generated core files after running the Packet Forwarding Engine command **show cos sched-usage**. [PR1449645](#)
- The **show cos scheds-per-pfe**, **show cos pfe-scheduler-ifds**, and **pfe** commands will restart forwarding planes on QFX10008 switches. [PR1452013](#)

EVPN

- Asynchronous result between ARP table and Ethernet switching table happens if EVPN ESI link flaps multiple times. [PR1435306](#)
- When using **no-arp-suppression** , an ARP request might not be sent out when an ARP entry aged out. [PR1441464](#)
- ARP and IPv6 neighbor entries cannot be cleared when they are learned from EVPN multihomed ESI. [PR1446957](#)
- EVPN-VXLAN NON-COLLAPSED: ARP will get resolved on QFX5100 for VXLAN having vlan-id of 2. [PR1453865](#)
- ARP request/NS might be sent back to the local segment by DF router. [PR1459830](#)

Forwarding and Sampling

- Commit error and dfwd core files might be observed when applying a firewall filter with action **then traffic-class** or **then dscp**. [PR1452435](#)

Interfaces and Chassis

- VRRPv6 state is flapping with init and idle states after configuring vlan-tagging. [PR1445370](#)
- On QFX10000 ARP entries might not be synchronized between MC-LAG devices. [PR1449806](#)
- The traffic might be forwarded to the incorrect interfaces in MC-LAG scenario. [PR1465077](#)
- Vrrpv3mibs are not working on QFX Series platform to poll VRRPv6 related objects. [PR1467649](#)

Layer 2 Features

- Storm control configuration might be disabled for the interface. [PR1354889](#)
- Packet loss might be seen when one of the spine switches fails or reboots. [PR1421672](#)
- Ethernet ring protection switching (ERPS) nodes might not converge to IDLE state after failure recovery or reboot. [PR1431262](#)
- EVPN-VXLAN NON-COLLAPSED: JTASK and multimove depth failed errors are seen after HALT. [PR1434687](#)
- The MAC/ARP learning might not work for copper base SFP-T on QFX5100/QFX5110/EX4600. [PR1437577](#)
- The traffic leaving QFX5000 and EX46000 switches might not be properly load-balanced over ae interfaces. [PR1448488](#)

- Unequal LAG hashing might happen on QFX devices. [PR1455161](#)
- The fxpc.core file might be seen when committing the configuration all together, for example, after the reboot. [PR1467763](#)

MPLS

- The l2circuit traffic might be silently dropped at EVPN SPINE/MPLS LSP TRANSIT device if VXLAN access interface flaps on remote PE node (QFX5110). [PR1435504](#)
- Packet loss might occur when ECMP resilient-hash is enabled on QFX5000 platforms. [PR1442033](#)

Platform and Infrastructure

- QFX5100-VC MacDrainTimeOut and bcm_port_update failed: Internal error. [PR1284590](#)
- On QFX5100 platforms, LR4 QSFP can take up to 15 minutes to come up after Virtual Chassis reboot. [PR1337340](#)
- When powering off an individual FPC, the other FPC Packet Forwarding Engine might go offline too. [PR1344395](#)
- Mib2d core file in mib2d_write_snmpidx at snmpidx_sync.c on both ADs while bringing up base traffic profile. [PR1354452](#)
- Need new CLI command to enable copying of Open vSwitch Database (OVSDb) to RAM on Virtual Chassis backup Routing Engine instead of SSD. [PR1382522](#)
- FEC error counts are not updating for QFX5110. [PR1382803](#)
- QSFP-100GBASE-SR4/LR4 might take a long time to come up after disabling interface or reboot. [PR1402127](#)
- Ping over loopback might not work over type 5 tunnel on QFX10000 platforms. [PR1405786](#)
- QFX5200/5100 might not be able to send out control plane traffic to the peering device. [PR1406242](#)
- No inner VLAN tag is added even with **input-vlan-map push** configured on QFX10000 platforms. [PR1407347](#)
- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. [PR1411015](#)
- QFX5120 : Route table full for IPv6 routes in some scenarios. [PR1412873](#)
- Intermittently chassis alarms might not be raised after power-cycle of the device. [PR1413981](#)
- IPv6 multicast traffic received on one Virtual Chassis member might be dropped when egressing on other Virtual Chassis member if MLD snooping is enabled. [PR1423310](#)

- Ports might get incorrectly channelized if they are 10-Gigabit Ethernet already and they are channelized to 10-Gigabit Ethernet again. [PR1423496](#)
- On QFX5000 or QFX10000 switches, packet drops might be seen for the traffic that has to go over type-5 overlay tunnel. [PR1423928](#)
- The dcpfe/Packet Forwarding Engine might not start on AS7816-64X and QFX5000 TVP platform devices. [PR1426737](#)
- QFX5210: Received LLDP frames on em0 not displaying in LLDP neighbor output. [PR1426753](#)
- QFX5100-VCF - rollback for uncommitted configuration takes 1 hour. [PR1427632](#)
- Packet drops, replication failure, or ksynd crashes might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- The dcpfe process might crash and restart in MC-LAG scenario when the ARP/NDP next hop is changed. [PR1427994](#)
- The **global-mac-limit** and **global-mac-ip-limit** might allow more entries than the configured values. [PR1428572](#)
- [QFX10008] After Routing Engine switchover, LED status is not set for missing fan tray. [PR1429309](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- The dcpfe might crash on all line cards on QFX10000 in a scaled setup. [PR1431735](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- Outer VLAN tag might not be pushed in the egress VXLAN traffic toward the host for Q-in-Q scenario. [PR1432703](#)
- Line card might crash due to plug in unsupported SFP-T module. [PR1432809](#)
- Traffic loss might be seen on QFX10000/PTX10000 platforms using line card LC1105. [PR1433300](#)
- Layer 3 filters applied to PVLAN IRB interface might not work after unified ISSU. [PR1434941](#)
- QFX5100-Virtual Chassis : NSSU: there might be approximate 1 minute traffic loss during NSSU with LACP link protection configuration. [PR1435519](#)
- The mc-ae interface might get stuck in waiting state in dual mc-ae scenario. [PR1435874](#)
- QFX5200 NSSU: dcpfe core file is seen after NSSU upgrade of backup followed by reboot. [PR1435963](#)
- DHCP discover packets sent to IP addresses in the same subnet as IRB interface cause the QFX5110 to send bogus traffic out of DHCP-snooping enabled interfaces. [PR1436436](#)
- Unknown SNMP traps (1.3.6.1.4.1.2636.3.69.1.0.0.1) are sent on QFX5110 restart. [PR1436968](#)
- The FPC might crash if both the ae bundle flapping on the local device and the configuration change on peer device occur at the same time. [PR1437295](#)
- BGP neighborship might not come up if the MACsec feature is configured. [PR1438143](#)
- The DHCP snooping table might be cleared for VLAN ID 1 after adding a new VLAN ID to it. [PR1438351](#)

- Port LED turns red when cable is connected on QFX5210. [PR1438359](#)
- Interfaces configured with **flexible-vlan-tagging** might loss connectivity. [PR1439073](#)
- The xSTP recognizes 1G SFP-T optic interface as LAN type resulting, in slow STP convergence. [PR1439095](#)
- LACP MUX state stuck in "Attached" after disabling peer active members when link protection is enabled on local along with force-up. [PR1439268](#)
- DHCPv6 relay binding is not up while verifying the DHCP Snooping along with DHCPv6 relay. [PR1439844](#)
- EX4600 Virtual Chassis does not comes up after replacing Virtual Chassis port from fiber connection to DAC cable. [PR1440062](#)
- MAC addresses learned on RTG might not be aged out after a Virtual Chassis member is rebooted. [PR1440574](#)
- QFX10002 MCLAG PDT: Layer 2, Layer 3 Traffic drop is seen on disabling/enabling MC-LAG. [PR1440732](#)
- The Layer 3 communication might break on an interface that is configured with **flexible-ethernet-services**. [PR1441690](#)
- The operational status of the interface in hardware and software might be out of synchronization in EVPN setup with **arp-proxy** feature enabled. [PR1442310](#)
- Flow control does not work as expected on 100-Gigabit Ethernet interface of QFX5110. [PR1442522](#)
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB. [PR1442587](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- When a line card is rebooted, the MC-LAG might not get programmed after the line card comes back online. [PR1444100](#)
- QFX5200: Observing **DCBCM[bcore_init]: ioctl call failed ret:0** failure message when changing UFT profile in FPC logs. [PR1445855](#)
- On QFX10008, traffic impact might be seen when the JSRV interface is used. [PR1445939](#)
- CoS classifier might not work as expected. [PR1445960](#)
- IPinIP: QFX - CoS rewrite happens to both inner and outer header. [PR1446128](#)
- IPinIP: ptx/qfx - Upon steering of underlay dynamic tunnel PNHs to a different set of ECMP next hops, unrelated IPv6 based tunnel traffic is tagged with the incorrect VLAN. [PR1446132](#)
- Traffic discarded for only specified VLAN in IPACL_VXLAN filters. [PR1446489](#)
- Long IPv6 address are not displayed fully on IPv6 neighbor table. [PR1447115](#)
- Unicast ARP requests are not replied with **no-arp-trap** option. [PR1448071](#)
- Rebooting QFX5120-48Y using **request system reboot** doesn't take physical links offline immediately. [PR1448102](#)

- QFX10000 -- QSFP28 100G AOC / 740-065632 & QSFP+ 40G / 740-043308 transceiver -- port LED remains lit green after disconnecting one end. [PR1448121](#)
- QFX5100-48t's in a mixed Virtual Chassis with QFX5110 switches are experiencing rx crc errors on vc-ports 53 and 52. [PR1449406](#)
- Except one AE member link, the other links do not send out sFlow sample packets for ingress traffic. [PR1449568](#)
- REST API process will get non-responsive when a number of request coming with a high rate. [PR1449987](#)
- RMPC core files are found after configuration changes are done on the network for PTP/Clock Synchronization. [PR1451950](#)
- Vgd core files might be generated when tunnel gets deleted twice. [PR1452149](#)
- DHCP offer packet with unicast flag set gets dropped by QFX10000 in a VXLAN multi-homed setup using anycast IP. [PR1452870](#)
- Configuration change in VLAN all option might affect the per-VLAN configuration. [PR1453505](#)
- The classifier configuration doesn't get applied to the interface in an EVPN/VXLAN environment. [PR1453512](#)
- The **show chassis led** shows incorrect status. [PR1453821](#)
- On QFX5100-VC VGD process hogs the CPU without **switch-options vtep-source-interface lo0.0** configuration. [PR1454014](#)
- Master FPC might come up in master state again after reboot instead of backup. [PR1454343](#)
- QFX10002-60c: EVPN-VXLAN: MAC+IP Count is shown as Zero. [PR1454603](#)
- QFX5120 : Untagged hosts ARP/NS connected on **encapsulation ethernet-bridge** interface are not being resolved. [PR1454804](#)
- The PFC feature doesn't work on QFX10000 platforms. [PR1455309](#)
- The laser from the 10G SFP+ interface is still on when the interface is disabled or the device is rebooted. [PR1456742](#)
- Over temperature SNMP trap messages are shown after update even though the temperatures are within the system thresholds. [PR1457456](#)
- Dual tag Q-in-Q is not working with EVPN-VXLAN. [PR1458206](#)
- QFX5210 : LED does not light on port 64 and 65 after upgraded to Junos OS Release 19.2R1. [PR1458514](#)
- The BPDU packet might be looped between leaf DF switch and non-DF switch and cause traffic blocking. [PR1458929](#)
- The dhcpv6 LDRA relay bounded count is not as expected after dhcp is configured. [PR1459499](#)
- The fxpc process might crash due to BGP IPV6 session flaps. [PR1459759](#)
- The **forwarding** option is missed in routing-instance type. [PR1460181](#)

- The 'entPhysicalTable' MIB is not fetching expected data on QFX10002-72Q / 36Q platforms. [PR1462582](#)
- The firewall filter does not get hit for traceroute packets when destination MAC address is VRRP virtual MAC. [PR1463425](#)
- On QFX5100 Virtual Chassis, the error **BRCM-VIRTUAL,brcm_vxlan_walk_svp(),6916:Failed to find L2-iff for ifl:** might appear during cleanup of EVPN-VXLAN configurations. These messages are harmless. [PR1463939](#)
- A few of the interfaces stay down and keep flapping for QFX ULC-3DWDM-MACsec line cards on reboot. [PR1464650](#)
- QFX5100-24Q: Not able to apply DSCP rewrite to firewall filter to a Layer 3 subinterface (for example, xe-0/0/0.100). [PR1464883](#)
- PEM is not present spontaneously on QFX5210. [PR1465183](#)
- The 10-Gigabit Ethernet port on QFX5100-48T negotiates with speed 1 GB with BRCM 10G/GbE 2+2P 57800-t rNDC. [PR1465196](#)
- The QSFP-100G-PSM4 could not be correctly identified on QFX5200 or QFX5110 platforms. [PR1465214](#)
- When BGP open messages with specific types of BGP optional capabilities are sent during BGP session establishment, incorrectly coded messages are later sent to the BMP Collector. [PR1466477](#)
- Slow packet drops might be seen on QFX5000 platforms. [PR1466770](#)
- Ingress drops to be included at CLI from interface statistics and added to InDiscards. [PR1468033](#)
- QFX5120 is looping the IP routed packet through IS-IS or MPLS. [PR1469998](#)
- l2ald core is seen (**l2ald_mem_free, l2ald_update_comp_vmenh**) after restarting dc-pfe in Virtual Chassis devices. [PR1473521](#)

Routing Protocols

- Host-destined packets with filter log action might not reach to Routing Engine if log/syslog is enabled. [PR1379718](#)
- The IRB transit traffic might not be counted for EVPN-VXLAN traffic. [PR1383680](#)
- QFX5100 : BGP IPv4 and IPv6 convergence and RIB installation and deletion time are degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The fxpc core file might be seen during the reboot of device on QFX5100/EX4600 switches. [PR1432023](#)
- The IPv4 fragmented packets might be broken if PTP transparent clock is configured. [PR1437943](#)
- Traffic might be dropped after the Q-in-Q enabled interface is flapped or a change is made to the vlan-id-list. [PR1441402](#)
- QFX5210: firewall Filter DSCP action modifier does not work when firewall filter is mapped to IRB. [PR1441444](#)

- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- PIM (S,G) joins can cause MSDP to incorrectly announce source active messages in some cases. [PR1443713](#)
- The QFX5120 might drop the tunnel encapsulated packets if it acts as a transit device. [PR1447128](#)
- Loopback address exported into other VRF instances might not work on ACX Series, EX Series, and QFX Series platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even the LDP neighbor's MAC changes. [PR1451217](#)
- A few seconds of traffic drop might be seen on the existing receivers when another receiver joins/leaves. [PR1457228](#)
- The egress interface in Packet Forwarding Engine for some end-hosts might not be correct on the Layer 3 gateway switch after it is rebooted. [PR1460688](#)
- The "other querier present interval" timer cannot be changed in IGMP/MLD snooping scenario. [PR1461590](#)
- When deleting IRB on the Layer 3 gateway, IRB does not get removed from Packet Forwarding Engine and will silently drop traffic to IRB MAC address. [PR1463092](#)

User Interface and Configuration

- EX4600 and QFX5100 were unable to commit baseline configuration after being returned to zero. [PR1426341](#)

SEE ALSO

[What's New | 201](#)

[What's Changed | 212](#)

[Known Limitations | 215](#)

[Open Issues | 217](#)

[Documentation Updates | 232](#)

[Migration, Upgrade, and Downgrade Instructions | 232](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 232](#)

This section lists the errata and changes in Junos OS Release 19.4R1 for the QFX Series switches documentation.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

[What's New | 201](#)

[What's Changed | 212](#)

[Known Limitations | 215](#)

[Open Issues | 217](#)

[Resolved Issues | 224](#)

[Migration, Upgrade, and Downgrade Instructions | 232](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 233](#)
- [Installing the Software on QFX10002-60C Switches | 235](#)

- Installing the Software on QFX10002 Switches | 235
- 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 | 236
- Installing the Software on QFX10008 and QFX10016 Switches | 238
- Performing a Unified ISSU | 242
- Preparing the Switch for Software Installation | 243
- Upgrading the Software Using Unified ISSU | 243
- Upgrade and Downgrade Support Policy for Junos OS Releases | 245

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

Upgrading Software on QFX Series Switches

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

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

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **19.4** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 19.4 release.

An Alert box appears.

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

A login screen appears.

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

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

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

```
user@host> request system software add  
source/jinstall-host-qfx-5-x86-64-19.4-R1.n-secure-signed.tgz reboot
```

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

- **/pathname**—For a software package that is installed from a local directory on the switch.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname** (available only for Canada and U.S. version)

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 19.4jinstall package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.

NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add <pathname><source>** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-19.4R1.9.tgz
```

If the Install Package resides remotely from the switch, execute the **request vmhost software add <pathname><source>** command.

For example:

```
user@switch> request vmhost software add
ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-19.4R1.9.tgz
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

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

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R1.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-f-x86-64-19.4R1.n-secure-signed.tgz reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-19.4R1.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

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

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in `/var/home`. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

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

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-19.4R1.n-secure-signed.tgz
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the **request system reboot** command:

```
user@switch> request system reboot
```

NOTE: You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete <package-name>** command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

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

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
Routing Engine status:
Slot 1:
  Current state           Master
  Election priority       Backup (default)
```

14. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-19.4R1.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 243](#)
- [Upgrading the Software Using Unified ISSU on page 243](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

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

To verify that nonstop active routing is enabled:

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

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

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

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

Upgrading the Software Using Unified ISSU

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

To upgrade the switch using unified ISSU:

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

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

where *package-name.tgz* is, for example, *jinstall-host-qfx-10-f-x86-64-19.4R1.n-secure-signed.tgz*.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

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

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
```

```

ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item                Status                Reason
  FPC 0                Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

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

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

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

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

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases

provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

SEE ALSO

What's New 201
What's Changed 212
Known Limitations 215
Open Issues 217
Resolved Issues 224
Documentation Updates 232

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 247](#)
- [What's Changed | 255](#)
- [Known Limitations | 262](#)
- [Open Issues | 263](#)
- [Resolved Issues | 266](#)
- [Documentation Updates | 273](#)
- [Migration, Upgrade, and Downgrade Instructions | 274](#)

These release notes accompany Junos OS Release 19.4R1 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

- [Application Security | 247](#)
- [Chassis Clustering | 248](#)
- [Flow-Based and Packet-Based Processing | 248](#)
- [General Packet Radio Switching \(GPRS\) | 249](#)
- [Hardware | 249](#)
- [Interfaces and Chassis | 250](#)
- [Intrusion Detection and Prevention \(IDP\) | 250](#)
- [Junos OS XML API and Scripting | 251](#)
- [J-Web | 252](#)
- [Logical Systems and Tenant Systems | 252](#)
- [Network Management and Monitoring | 253](#)
- [System Logging | 254](#)
- [Unified Threat Management \(UTM\) | 254](#)
- [VPNs | 254](#)

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

Application Security

- **Selectively disable midstream APBR (SRX Series and vSRX)**—Starting in Junos OS Release 19.4R1, you can selectively turn-off midstream routing for a specific APBR rule while retaining the global APBR functionality for the remaining sessions.

When you disable the midstream routing for a specific APBR rule, the system does not apply midstream APBR for corresponding application traffic, and routes the traffic through a non-APBR route

[See [Advanced Policy-Based Routing](#).]

- **DSCP support for AppQoS (SRX Series and vSRX)**—Starting in Junos OS Release 19.4R1, AppQoS supports SLA-based path selection for an incoming traffic based on Differentiated Services Code Point (DSCP) value.

AppQoS depends on AppID and APBR to select the best possible link for the application traffic to meet the performance requirements specified in SLA. Junos OS Release 19.3R1 introduced APBR functionality for DSCP-tagged traffic. Using this enhancement, AppQoS selects the best possible link for the application traffic based on the application signature, or DSCP value, or a combination of both application signature and DSCP value.

With this enhancement, now you can apply AppQoS for the encrypted traffic based on the DSCP value.

[See [Application Quality of Experience](#).]

- **Support for server certificates with key size 4096 bits (SRX300 and SRX320)**—Starting in Junos OS Release 19.4R1, SRX300 and SRX320 devices support RSA certificates with key size 4096 bits. You must explicitly configure the SSL proxy profile on these devices to use the server certificate with key size 4096 bits.

The RSA certificates with key size 4096 bits support is available only when the SRX300 and SRX320 devices are operating in standalone mode.

[See [Managing Certificates and Keys for SSL Proxy](#).]

Chassis Clustering

- **Increase in the maximum number of child links (SRX4600)**—Starting in Junos OS Release 19.4R1, you can configure up to eight child links in a redundant Ethernet bundle on each node of the chassis cluster.

See [[Configuring Chassis Cluster Redundant Ethernet Interfaces on SRX4600](#).]

Flow-Based and Packet-Based Processing

- **Express Path for Flow Processing (SRX4600)**—Starting from Junos OS 19.4R1, Express Path is enabled by default on SRX4600 devices. You must configure Express Path only in policies. There is no need to configure Express Path on Flexible PIC Concentrator (FPC) or on Physical Interface Cards (PIC).

See [[Express Path](#).]

- **Support of IPFIX formatting for SRX J-Flow functionality (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vSRX, and vSRX3.0)** —Starting with Junos OS Release 19.4R1, you can use IPFIX flow templates to define a flow record for IPv4 traffic or IPv6 traffic. IPFIX is an enhanced version of J-flow version 9 template. Using IPFIX, you can collect a set of sampled flows and send the record to a specified host.

See [[Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches.](#)]

- **Symmetric Fat Tunnel (SRX5400, SRX5600, and SRX5800 devices with SPC3 card, and vSRX)**—Starting from Junos OS 19.4R1, fat tunnel technology is introduced to improve the single IPsec tunnel throughput value to 10 times of current value.

To enable this feature, a new CLI command **fat-core** is introduced at the **set security distribution-profile** hierarchy level.

See [[Understanding Symmetric Fat IPsec Tunnel.](#)]

General Packet Radio Switching (GPRS)

- **Increase in GTP scale for IoT and roaming firewall applications (SRX5400, SRX5600, SRX5800, and SRX4600)**—Starting in Junos OS Release 19.4R1, to enable Internet of Things (IoT) and roaming firewall use cases, the GTP tunnel scale per SPU is increased for the following SRX devices:
 - SRX5000 (SRX5400, SRX5600, SRX5800) SPC3: 1.2M to 12M
 - SRX5000 (SRX5400, SRX5600, SRX5800) SPC2: 600K to 3M
 - SRX4600: 400K to 4M

[See [Understanding Policy-Based GTP.](#)]

Hardware

- **Wi-Fi mini-physical interface module (SRX320, SRX340, SRX345, and SRX550M)**—The Wi-Fi mini-physical interface module (mini-PIM) provides an integrated wireless LAN access point solution for branch SRX Series Services Gateways. The Mini-PIM supports the 802.11ac Wave 2 wireless standards and is backward-compatible with 802.11a, 802.11b, 802.11g, and 802.11n.

The Mini-PIM supports the following key features:

- 2x2 MU-MIMO
- Dual radios, which provide concurrent dual bands of 2.4 GHz and 5 GHz
- Eight virtual access points (VAPs) per radio
- Configurable transmit power
- 128 concurrent users

The Wi-Fi Mini-PIM is available in three models based on the regional wireless standards:

- SRX-MP-WLAN-US (United States)
- SRX-MP-WLAN-IL (Israel)
- SRX-MP-WLAN-WW (other countries)

[See [Wi-Fi Mini-Physical Interface Module](#)].

- **SRX5K-SPC3 LTC firmware version check and upgrade**—Starting in Junos OS Release 19.4R1, you can check the current LTC Firmware version on an SRX5K-SPC3 card and upgrade the firmware version manually.

The LEDs on the front panel of the services gateway chassis indicate a major alarm when the chassis detects that a newer version of LTC firmware is available and the firmware on the SRX5K-SPC3 card is outdated. The CLI commands:

- **show chassis alarm**—displays the alarm description
- **show system firmware**—displays the current version, available version, and the Status of the LTC firmware
- **request system firmware upgrade pic fpc-slot 0 pic-slot 0 tag 0**—updates the LTC firmware version.

[See [Chassis Component Alarm Conditions on SRX5400, SRX5600, and SRX5800 Services Gateways](#).]

Interfaces and Chassis

- **Wi-Fi Mini-Physical Interface Module (SRX320, SRX340, SRX345, and SRX550M)**—In Junos OS Release 19.4R1, we introduce the Wi-Fi Mini-Physical Interface Module (Mini-PIM). For retail and small offices, the Wi-Fi Mini-PIM provides secure wireless LAN connectivity to endpoint devices. The Wi-Fi Mini-PIM supports 802.11ac wave 2 wireless standards.

[See [Wi-Fi Mini-Physical Interface Module Overview](#).]

- **LTE Support in HA deployments (SRX300, SRX320, SRX340, SRX345, and SRX550HM)**—Starting in Junos OS Release 19.4R1, you can use one LTE interface on an LTE High Availability (HA) node separately and enable failover and backup between interfaces.

[See [Configuring the LTE Mini-PIM as a Backup Interface](#).]

Intrusion Detection and Prevention (IDP)

- **IDP utility to read packet capture and generate protocol contexts (SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM)**—Starting from Junos OS Release 19.4R1, on SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM devices, to improve the IDP validation process, a CLI command is introduced to display and clear the contexts and the associated data only for the packet capture (PCAP) traffic. You can run the packet capture utility in either inet mode or transparent mode to generate attack contexts.

See [[IPD Utility for PCAP](#).]

- **Signature Language Constructs (SRX Series)**—Starting from Junos OS 19.4R1, signature language constructs are supported in the IDP engine code to write more efficient signatures that helps in reducing false positives.

The following constructs are supported:

- Depth
- Offset
- Within
- Distance
- Ipopts

See [[IDP Signature Language Enhancements](#).]

Junos OS XML API and Scripting

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

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

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

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

- **pyasn1** (0.4.5)
- **yaml** (PyYAML package) (5.1)

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

J-Web

- **Threat Assessment report enhancement (SRX Series)**—Starting in Junos OS Release 19.4R1, the Threat Assessment report displays a new *Filename* column in the *Malware downloaded by User* table. This column helps you to identify the malware filename.

[See [About Reports Page](#).]

- **UTM enhancement (SRX Series)**—Starting in Junos OS Release 19.4R1, the following UTM pages (Configure > Security Services > UTM) are refreshed for a seamless experience:
 - Antivirus
 - Content Filtering
 - Policy

[See [About the Antivirus Page](#), [About the Content Filtering Page](#), and [About the Policy Page](#).]

- **Support for Wi-Fi Mini-PIM (SRX320, SRX340, SRX345, and SRX550M devices)**—Starting in Junos OS Release 19.4R1, J-Web supports the Wi-Fi Mini-Physical Interface Module (Mini-PIM). The physical interface for the Wi-Fi Mini-PIM uses the name **wl-x/0/0**, where x identifies the slot on the services gateway where the Mini-PIM is installed.

You can monitor and configure the wireless LAN settings using the J-Web interface.

[See [Dashboard Overview](#), [Monitor Ports](#), [About the Ports Page](#), [Monitor Wireless LAN](#), and [About the Settings Page](#).]

Logical Systems and Tenant Systems

- **Flow trace support at logical system and tenant system level (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.4R1, you can trace the packet flow at the logical system level and tenant system level. Traceoptions enables you to monitor traffic flow into and out of an SRX Series device.

When you trace traffic flow, you can generate and save the trace logs to the respective logical system and tenant system log files.

Flow trace at the level of logical system and tenant system helps you avoid generating large log files from the root level.

[See [Flow Trace Support for Logical Systems](#) and [Flow Trace Support for Tenant Systems](#).]

- **AppID statistics at tenant system level (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.4R1, you can view or clear the application identification (AppID) statistics, counters, and application system cache at the tenant system level.

[See [Application Security for Tenant Systems](#).]

Network Management and Monitoring

- **SNMP support for Wi-Fi Mini-Physical Interface Module (Mini-PIM) monitoring (SRX320, SRX340, SRX345, and SRX550M)**—Starting in Junos OS Release 19.4R1, you can monitor the Wi-Fi Mini-PIM status from remote network using SNMP. Use the **show snmp mib walk ascii jnxWlanWAPStatusTable** and **show snmp mib walk jnxWlanWAPClientTable** commands to monitor the Wi-Fi Mini-PIM status and client information.

[See [SNMP MIB Explorer](#) and [show snmp mib](#).]

- **SNMP support for IPsec VPN flow monitoring (SRX5000 Series devices with SRX5K-SPC3 card)**—Starting in Junos OS Release 19.4R1, we have enhanced the existing IPsec VPN flow monitor MIB **jnxIpSecFlowMonMIB** to support the global IKE statistics for tunnels using IKEv2. Use the **show security ike stats** command to display the global statistics of tunnels such as in-progress, established, and expired negotiations using IKEv2.

[See [Enterprise-Specific SNMP MIBs Supported by Junos OS](#) and [show security ike stats](#).]

- **Improved query performance in on-box reporting (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 19.4R1, we've upgraded the on-box logging database to improve query performance. For example, if you expect fewer traffic logs, you can use the default configuration with a start time and a stop time. If you expect a large number of traffic logs and greater time intervals for which the logs will be generated, we recommend you enable table dense mode.

[See [Understanding On-Box Logging and Reporting](#).]

- **Enhanced support for the non-default management instance (SRX Series)**—Starting in Junos OS 19.4R1, you can access information related to all routing instances and logical system networks and not specific to ingress routing instance by configuring the SNMPv3 management interface in a required management instance. Configuring the SNMPv3 management interface in a required management instance enables all the SNMPv3 requests coming from non-default routing instance is treated as if the requests are coming from default routing instance. You can configure the management instance configuration statement at the **[edit SNMP v3]** hierarchy level.

[See [SNMPv3 Management Routing Instance](#).]

System Logging

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

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

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

Unified Threat Management (UTM)

- **UTM support for active/active chassis cluster (SRX Series devices)**—Starting in Junos OS Release 19.4R1, you can configure all Unified Threat Management (UTM) features when the device is in active/active chassis cluster mode. The UTM features supported on an active/active chassis cluster include, Antispam Filtering, Content Filtering, Sophos Antivirus Scanning, URL (Web) Filtering, Enhanced Web Filtering, Local Web Filtering, and Websense Redirect Web Filtering, and On-box/Avira AV. Enhanced Web Filtering and Sophos Antivirus Scanning remain active on both the primary node and the secondary node.

[See [Understanding UTM Support for Active/Active Chassis Cluster](#).]

- **UTM support for SMTPS, IMAPS, POP3S, and FTPS (SRX Series devices)** —Starting in Junos OS Release 19.4R1, UTM supports implicit and explicit SMTPS, IMAPS, and POP3S protocol and explicit passive-mode FTPS. SMTPS, IMAPS, POP3S, and FTPS are methods for securing SMTP, IMAP, POP3, FTP protocols using Transport Layer Security (TLS). Antivirus and content filtering feature supports SMTPS, IMAPS, POP3S and FTPS protocol. Antispam feature only supports SMTPS protocol.

[See [Antispam Filtering Overview](#) and [Understanding Content Filtering Protocol Support](#).]

VPNs

- **Extended Sequence Number (SRX5400, SRX5600, and SRX5800 devices using SPC3)**—Starting from Junos OS Release 19.4R1, Extended Sequence Number (ESN) is introduced in IPsec VPN using IKE version 2 (IKEv2).
IPSec uses a 32-bit sequence number by default for the sequence number. When all sequence numbers are consumed, a rekey must be issued. By enabling ESN this 32-bit sequence numbering is increased to 64-bit.

You can enable ESN using the **set extended-sequence-number** command at the **edit security ipsec proposal proposal-name** level.

[See [Understanding Extended Sequence Number \(ESN\)](#).]

- **VPN support for inserting Services Processing Cards in Chassis Cluster (SRX5400, SRX5600, SRX5800)**—Starting in Junos OS Release 19.4R1, on all SRX5000 Series devices chassis cluster, you can insert a new SRX5K-SPC3 (SPC3) on the device without affecting or disrupting the traffic on the existing IKE or IPsec VPN tunnels. When you insert the new SPC3 in each chassis of the cluster, the existing tunnels are not affected and traffic continues to flow without disruption. You must reboot the node after you insert the SPC3 to activate the card. After the node reboot is complete, IPsec tunnels are distributed to the cards. After you reboot the secondary node where the new spc3 card is inserted, the IPsec sessions remain active on the other active node, without disruption to those sessions except during the failover time.

[See [Understanding VPN Support for Inserting Services Processing Cards](#).]

- **IPsec Encapsulating Security Payload authentication-only mode in PowerMode IPsec (SRX5000 Series devices with SRX5K-SPC3 card, and vSRX)**—Starting in Junos OS Release 19.4R1, you can enable the IPsec Encapsulating Security Payload (ESP) authentication-only mode in the PowerMode IPsec (PMI). The ESP authentication-only mode provides authentication, integrity checking, and replay protection in the PMI.

[See [Improving IPsec Performance with PowerMode IPsec](#).]

SEE ALSO

What's Changed	 255
Known Limitations	 262
Open Issues	 263
Resolved Issues	 266
Documentation Updates	 273
Migration, Upgrade, and Downgrade Instructions	 274

What's Changed

IN THIS SECTION

- [Application Security](#) | 256
- [Authentication and Access Control](#) | 259
- [Class of Service](#) | 259
- [General Routing](#) | 259

- J-Web | 259
- Network Management and Monitoring | 259
- Port Security | 261
- Routing Protocols | 261
- System Logging | 261
- VPNs | 261

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

Application Security

- Starting in Junos OS Release 19.4R1, you have the flexibility to limit the application identification inspection as follows:

- **Inspection Limit for TCP and UDP Sessions**

You can set the byte limit and the packet limit for application identification (AppID) in a UDP or in a TCP session. AppID concludes the classification based on the configured inspection limit. On exceeding the limit, AppID terminates the application classification.

If AppID does not conclude the final classification within the configured limits, and a pre-matched application is available, AppID concludes the application as the pre-matched application. Otherwise, the application is concluded as `junos:UNKNOWN` provided the global AppID cache is enabled. The global AppID cache is enabled by default.

To configure the byte limit and the packet limit, use the following configuration statements from the **[edit]** hierarchy:

- ```
user@host# set services application-identification inspection-limit tcp byte-limit byte-limit-number
packet-limit packet-limit-number
```
- ```
user@host# set services application-identification inspection-limit udp byte-limit byte-limit-number
packet-limit packet-limit-number
```

[Table 3 on page 257](#) provides the range and default value for configuring the byte limit and the packet limit for TCP and UDP sessions.

Table 3: Maximum Byte Limit and Packet Byte Limit for TCP and UDP Sessions

Session	Limit	Range	Default Value
TCP	Byte limit	0 through 4294967295	<ul style="list-style-type: none"> 6000 For Junos OS Release 15.1X49-D200, the default value is 10000.
	Packet limit	0 through 4294967295	Zero
UDP	Byte limit	0 through 4294967295	Zero
	Packet limit	0 through 4294967295	<ul style="list-style-type: none"> 10 For Junos OS Release 15.1X49-D200, the default value is 20.

The byte limit excludes the IP header and the TCP/UDP header lengths.

If you set 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 19.4R1, the maximum packet threshold for DPI performance mode option **set services application-identification enable-performance-mode max-packet-threshold *value*** is deprecated—rather than immediately removed—to provide backward compatibility and an opportunity

to bring your configuration into compliance with the new configuration. This option was used for setting the maximum packet threshold for the DPI performance mode.

If your configuration includes enabled performance mode option with **max-packet-threshold** in Junos OS releases 15.1X49-D200 and 19.4R1, AppID concludes the application classification on reaching the lowest value configured in the TCP or UDP inspection limit or in the global offload byte limit, or in the maximum packet threshold for DPI performance mode option.

[See [Application Identification Inspection Limit](#) and [application-identification](#)]

- Starting in Junos OS Release 19.4R1, the **apbr-rule-type** field in the system log message displays the value as **none** if no rule is applied when you have disabled midstream for the application. Updated syslog message sample is as following:

```
<14>1 2019-07-11T03:06:27.276-07:00 pavna RT_FLOW - APPTRACK_SESSION_CLOSE
[junos@2636.1.1.1.2.140 reason="TCP FIN" source-address="4.0.0.1"
source-port="33810" destination-address="5.0.0.1" destination-port="80"
service-name="junos-http" application="HTTP" nested-application="YAHOO"
nat-source-address="4.0.0.1" nat-source-port="33810"
nat-destination-address="5.0.0.1" nat-destination-port="80" src-nat-rule-name="N/A"
dst-nat-rule-name="N/A" protocol-id="6" policy-name="1" source-zone-name="trust"
destination-zone-name="untrust" session-id-32="370" packets-from-client="969"
bytes-from-client="50518" packets-from-server="1107" bytes-from-server="1057897"
elapsed-time="253" username="N/A" roles="N/A" encrypted="No"
profile-name="profile1" rule-name="N/A" routing-instance="default"
destination-interface-name="xe-1/1/4.0" uplink-incoming-interface-name="xe-1/1/4.0"
uplink-tx-bytes="50518" uplink-rx-bytes="1057897" category="Web"
sub-category="miscellaneous" apbr-policy-name="sla1" multipath-rule-name="N/A"
src-vrf-grp="N/A" dst-vrf-grp="N/A" dscp-value="N/A" apbr-rule-type="none"]
```

- Starting in Junos OS Releases 19.4R1, security policy does not support using following applications as dynamic-applications match criteria:
 - junos:HTTPS
 - junos:POP3S
 - junos:IMAPS
 - junos:SMTPS

Software upgrade to the Junos OS Releases 19.4R1 fails during the validation if any of the security policies are configured with junos:HTTPS, junos:POP3S, junos:IMAPS, junos:SMTPS as dynamic-applications as match criteria. We recommend you to remove any configuration that includes these dynamic-applications as match criteria in security policies.

We recommend you to use the **request system software validate package-name** option before upgrading to the above mentioned releases.

Authentication and Access Control

- **Enabling and disabling SSH login password or challenge-response authentication (SRX Series)**—Starting in Junos OS Release 19.4R1, you can disable either the SSH login password or the challenge-response authentication at the `[edit system services ssh]` hierarchy level.

In Junos OS releases earlier than Release 19.4R1, you can enable and disable both SSH login password and the challenge-response authentication simultaneously at the `[edit system services ssh]` hierarchy level.

[See [Configuring SSH Service for Remote Access to the Router or Switch](#).]

Class of Service

- We've corrected the output of the `show class-of-service interface | display xml` command. Output of the following sort:

`<container> <leaf-1>data <leaf-2> data<leaf-3> data<leaf-1> data<leaf-2> data <leaf-3> data` will now appear correctly as:`<container> <leaf-1> data <leaf-2> data <leaf-3> data <container> <leaf-1> data <leaf-2> data <leaf-3> data`.

General Routing

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

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

J-Web

- Deactivated policy rules are not visible in the J-Web UI (SRX Series)—J-Web does not support disabling or enabling the security firewall or global policy rules from Junos OS Release 19.4R1. The policy rules that are deactivated through CLI are also not visible in the J-Web UI. As a workaround, use CLI to disable or enable the policy rules on the device.

Network Management and Monitoring

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

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

- **Change in On-box reporting factory-default configuration (SRX1500, SRX4100, SRX4200, SRX4600 and vSRX)**—Starting in Junos OS Release 19.4R1, the factory-default configuration does not include on-box reporting configuration to increase the solid-state drive (SSD) lifetime. You can enable the on-box reporting by configuring the **set security log report** CLI command at **[edit security log]** hierarchy.

[See [Understanding On-Box Logging and Reporting](#).]

- **Change in jnxJsFlowMIB statistics display (SRX Series)**—Starting in Junos OS Release 19.4R1, in a chassis cluster, you can see the statistics on all SPUs of both nodes using the **show snmp mib walk jnxJsFlowMIB** command. In the earlier releases, you can see the statistics only on local SPUs.

[See [SNMP MIB Explorer](#).]

Port Security

- **Configuring source mac filters (SRX300 and SRX550 Services Gateway)**—In this release of Junos OS, fixed an issue that prevented source mac filters from being configured on an interface. The error effected both the **accept-source-mac** and **source-address-filter** statements and resulted in one of the following error messages: **accept-source-mac not allowed in switching mode** and **source mac filters not allowed in switching mode**.

Routing Protocols

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

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

System Logging

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

VPNs

- **IKE gateway dynamic distinguished name attributes (SRX Series devices)**—Starting in Junos OS Release 19.4R1, you can now configure only one dynamic distinguished name (DN) attribute among **container-string** and **wildcard-string** at **[edit security ike gateway gateway_name dynamic distinguished-name]** hierarchy. If you try configuring the second attribute after you configure the first attribute, the first attribute is replaced with the second attribute. Before you upgrade your device, you must remove one of the attributes if you have configured both the attributes.

[See [distinguished-name \(Security\)](#) and [Understanding IKE Identity Configuration](#).]

- **CoS Forward Class name (SRX Series devices)**—Starting in Junos OS Release 19.4R1, we have deprecated the CLI option **fc-name** (CoS Forward Class name) in the new **iked** process that displays security associations (SAs) under show command **show security ipsec sa**.

[See [show security ipsec security-associations](#).]

SEE ALSO

[What's New | 247](#)[Known Limitations | 262](#)[Open Issues | 263](#)[Resolved Issues | 266](#)[Documentation Updates | 273](#)[Migration, Upgrade, and Downgrade Instructions | 274](#)

Known Limitations

Learn about known limitations in Junos OS Release 19.4R1 for SRX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Application Layer Gateways (ALGs)

- Due to an SSL-FP limitation, the active mode of TLS-based FTP is not supported in Junos OS Release 19.4R1. [PR1450924](#)

Ethernet 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](#)

Flow-Based and Packet-Based Processing

- For any WiFi configuration change, the access point restarts to make the configuration active. [PR1436587](#)
- The SSID in different WLANs uses the same IP address as the source IP address of the radius packet. [PR1445276](#)
- TKIP is not supported in acn mode. [PR1459160](#)

J-Web

- The CA profile group imported using J-Web is not populated in the Certificate Authority Group initial landing page grid, but all the CA profiles of a group are populated on the Trusted Certificate Authorities landing page. [PR1426682](#)
- J-Web does not provide an option to enable or disable security objects. Security objects that are deactivated or disabled through the CLI are not displayed in the J-Web UI.

SEE ALSO

[What's New | 247](#)

[What's Changed | 255](#)

[Open Issues | 263](#)

[Resolved Issues | 266](#)

[Documentation Updates | 273](#)

[Migration, Upgrade, and Downgrade Instructions | 274](#)

Open Issues

IN THIS SECTION

- [ALG | 264](#)
- [Flow-Based and Packet-Based Processing | 264](#)
- [IDP | 264](#)
- [J-Web | 264](#)
- [Platform and Infrastructure | 264](#)
- [Routing Policy and Firewall Filters | 264](#)
- [VPNs | 265](#)

Learn about open issues in Junos OS Release 19.4R1 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.

ALG

- On SRX5000 line of devices, the H323 call with NAT64 could not be established. [PR1462984](#)

Flow-Based and Packet-Based Processing

- Use 512 anti-replay window size for IPv6 in **fat-tunnel**. The esp sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). So that there is no out-of-order with 512 anti-replay window size. [PR1470637](#)

IDP

- Rogue .gz files in **/var/tmp/sec-download/** might cause offline secpack update to fail. [PR1466283](#)

J-Web

- While adding a policy rule, the creation of inline scheduler or UTM or redirect profile objects automatically refreshes the policy grid, and the changes made to the policy rule are lost unless the changes are updated. [PR1451274](#)
- The **Interconnect Ports** page cannot be used from J-Web because the Type list does not contain any values. [PR1478333](#)

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](#)
- On SRX4600 devices, the Packet Forwarding Engine stops due to a segmentation problem. [PR1422466](#)
- If security datapath configuration is applied on tunnel transit traffic, ESP traffic is not captured. [PR1442132](#)
- On the SRX300 line of devices with Mini-PIM installed, tail-drop might happen on all ports when the serial egress port gets congested. [PR1468430](#)
- The **request chassis fpc restart** command does not work in Layer 2 mode on Wi-Fi Mini-PIM. [PR1479396](#)

Routing Policy and Firewall Filters

- SSL reverse proxy feature must be used instead of SSL inspection feature. SSL inspection on IDP level is being deprecated in favor of SSL reverse proxy. [PR1450900](#)

- On SRX5400, SRX5600, and SRX5800 devices, in some scenarios, DHCPv4 client might return to INIT state after chassis reboot if DHCPv4 retries are insufficient. [PR1458490](#)
- Whenever a high CPS traffic being passed through the SRX device making use of SSL proxy feature with SSL session resumption enabled (by default SSL session-resumption is enabled) the device might run into low memory and the sessions being bypassed by SSL proxy. [PR1472077](#)

VPNs

- On SRX Series devices, if 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 rekey. New negotiation of those traffic selectors might trigger through other mechanisms such as traffic or by peer. [PR1287168](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, when SRX Series device is configured in IKEv1 and NAT traversal is active, after a successful IPsec rekey, the IPsec tunnel index might change. In such a scenario, there might be some traffic loss for a few seconds. [PR1409855](#)
- On SRX5400, SRX5600, and SRX5800 devices 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](#)
- On SRX5400, SRX5600, and SRX5800 devices, during in-service software upgrade (ISSU), the IPsec tunnels flap, disrupting traffic. The IPsec tunnels recover automatically after the ISSU process is completed. [PR1416334](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with IKEv1 enabled IKE daemon might core, when IKESA is expired and IPsec tunnel associated with the expired IKESA exists in case of an RGO failover. Daemon recovers eventually. [PR1463501](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, when RGO failover happens while a IPsec and/or IKE rekey in progress, those rekeying tunnels might go down and possible traffic loss seen till tunnel is reestablished by the peer. [PR1471499](#)

SEE ALSO

[What's New | 247](#)

[What's Changed | 255](#)

[Known Limitations | 262](#)

[Resolved Issues | 266](#)

[Documentation Updates | 273](#)

[Migration, Upgrade, and Downgrade Instructions | 274](#)

Resolved Issues

This section lists the issues fixed in hardware and software in Junos OS Release 19.4R1 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.

Application Layer Gateways

- Unexpected forwarding sessions might appear for tenant SIP ALG traffic in the case of cross-tenants. [PR1409748](#)
- The flowd or srpxfe process might stop in SRX Series devices with chassis cluster when SIP ALG is used. [PR1445766](#)
- Packet loss happens during cold synchronization from the secondary node after rebooting. [PR1448252](#)
- After Layer 3 HA is enabled, ALG H.323 group or resource cannot be synchronized to the peer node correctly. [PR1456709](#)

Application Security

- The AAMW diagnostic script generates incorrect error: **Error: Platform does not support SkyATP: srx300.** [PR1423378](#)
- If automatic application-identification download is configured with a start-time specified, the automatic download stops when the time has progressed to the next year and a reboot is done before the start-time is reached that year. [PR1436265](#)
- SSL-based AppID simplification effort (removal of HTTPS, POP3S, IMAPS, SMTPS). [PR1444767](#)
- The flowd process core files might be generated when traffic hits the AppQoS policy. [PR1446080](#)
- The AAMW diagnostic script generates incorrect error when there is Internet latency: **Error: server unreachable is detected, please make sure port 443 is reachable.** [PR1468114](#)

Authentication and Access Control

- Same-source IP sessions are cleared when the IP entry is removed from the UAC table. [PR1457570](#)

Chassis Clustering

- Hardware failure is seen on both nodes in the output of the **show chassis cluster status** command. [PR1452137](#)

- On SRX Series devices with chassis cluster, the control link remains up even though the control link is actually down. [PR1452488](#)

Class of Service

- Frequent issuance of the **show class-of-service spu statistics** command causes the rtlogd process to be busy. [PR1438747](#)

Flow-Based and Packet-Based Processing

- Throughput or latency performance of TCP traffic is dropped when TCP traffic passes from one logical system to another logical system. [PR1403727](#)
- Packet loss is caused by FPGA back pressure on the SPC3 card. [PR1429899](#)
- VPN traffic fails after the primary node is rebooted or powered off. [PR1433336](#)
- Currently, PMI doesn't support the mirror-filter functionality. If mirror filters are configured, PMI flaps all of the traffic to the regular flow path. [PR1434583](#)
- Intermittent packet drop might be observed if IPsec is configured. [PR1434757](#)
- On an SRX4600 device, core file might be generated and SPM might be in present state. [PR1436421](#)
- Security logs cannot be sent to the external syslog server through TCP. [PR1438834](#)
- Decryption traffic doesn't take PMI path after IPsec rekey (initiated by peer) when the loopback interface is configured as an external interface. [PR1438847](#)
- The IKE pass-through packet might be dropped after a NAT operation on the source. [PR1440605](#)
- New CLI option to show only useful group information for an Active Directory user. [PR1442567](#)
- While checking the flow session XML for source NAT under tenant, there is no value identifier for **tenant-name**. [PR1440652](#)
- The flowd or srpxfe process might stop when processing fragmented packets. [PR1443868](#)
- Junos OS: SRX5000 Series: flowd process crash due to receipt of specific TCP packet (CVE-2019-0064). [PR1445480](#)
- J-Flow version 5 stops working after changing the input rate value. [PR1446996](#)
- Packet loss happens during cold synchronization from secondary node after rebooting. [PR1447122](#)
- On the SRX1500 device, automatic installation is removed from CLI. [PR1447796](#)
- SPC3 talus FPGA stuck on 0x3D or 0x69 golden version. [PR1448722](#)
- Host inbound or host outbound traffic on VR does not work when the SRX5000 line of devices works in SPC3 mixed mode. [PR1449059](#)

- SPU priority does not work when PMI is enabled on the SRX5000 line of devices with an SPC3 card. [PR1449587](#)
- All ingress packets are dropped if the traffic transit network is also the same network for LTE mPIM internal management. [PR1450046](#)
- The flowd or srxpfe process might stop when SSL proxy service is used. [PR1450829](#)
- The AAWM policy rules for IMAP traffic sometimes might not get applied when passed through SRX Series devices. [PR1450904](#)
- FTP data cannot pass through SRX320 4G wireless from FTP server to client. [PR1451122](#)
- Traffic forwarding on Q-in-Q port and VLAN tagging are not observed properly on R0. [PR1451474](#)
- The rpd process might stop and restart with the generation of an rpd core file when committing the configuration. [PR1451860](#)
- The SRX Series devices stop and several core files are generated. [PR1455169](#)
- Added some JP APN settings to default list in LTE mPIM. [PR1457838](#)
- Changing the **RESET** configuration button behavior on the SRX1500 does not work. [PR1458323](#)
- The **security flow traceoptions** fills in with RTSP ALG related information. [PR1458578](#)
- Optimizations were made to improve the connections-per-second performance of SPC3. [PR1458727](#)
- The **security-intelligence** CC feed does not block HTTPS traffic based on SNI. [PR1460384](#)
- The AAMWD process exceeds 85 percent RLIMIT_DATA limitation due to memory leak. [PR1460619](#)
- Added command to clear specified associated client. [PR1461577](#)
- The tunnel packets might be dropped because the gr0.0 or st0.0 interface is wrongly calculated after a GRE or VPN route change. [PR1462825](#)
- Fragmented traffic might get looped between the fab interface in a rare case. [PR1465100](#)
- Track Jbuf double free issue. [PR1465286](#)
- HTTP block message stops working after SNI check for HTTPS session. [PR1465626](#)
- The jbuf process usage might increase up to 99 percent after Junos OS upgrade. [PR1467351](#)
- The rpd process might stop after several changes to the **flow-spec** routes [PR1467838](#)
- FTP data connection might be dropped on dl interface. [PR1468570](#)

Interfaces and Chassis

- SCB4 or SCB3 ZF or XF2 fabric plane retraining is needed after switching the fabric redundancy mode. [PR1427119](#)
- MTU change after a CFM session is up can impact Layer 2 Ethernet ping (loopback messages). If the new change is less than the value in the initial incarnation, then Layer 2 Ethernet ping fails. [PR1427589](#)
- The LACP interface might flap while performing a failover. [PR1429712](#)
- LFM remote loopback is not working as expected. [PR1428780](#)
- The number of mgd processes increases as the mgd processes are not closed properly. [PR1439440](#)
- The fxp0 interface might redirect packets not destined to itself. [PR1453154](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srpxfe process crashes and generates a core file. [PR1437569](#)
- CLI helper text was added to the IDP's attack chain expressions. [PR1438620](#)

J-Web

- The default log query time in J-Web monitoring functionality has been reduced. This increases the responsiveness of the landing pages. [PR1423864](#)
- Phone home UI portal to be removed from SRX Series devices. [PR1428717](#)
- Some error messages might be seen when using J-Web. [PR1446081](#)
- The idle-timeout for J-Web access does not work properly. [PR1446990](#)
- J-Web fails to display the traffic log in event mode when stream mode host is configured. [PR1448541](#)
- Editing destination NAT rule in J-Web introduces a non-configured routing-instance field. [PR1461599](#)
- The **Go** button within the J-Web Monitor->Events view now correctly refreshes the logs even when using a blank search query. [PR1464593](#)
- J-Web security resources dashboard widget was not being populated correctly. [PR1464769](#)

Layer 2 Ethernet Services

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

Network Address Translation

- The nsd process might stop when SNMP queries deterministic NAT pool information. [PR1436775](#)
- Flowd process core files are generated in the device while testing NAT PBA in AA mode. [PR1443148](#)
- RTSP resource session is not found during NAT64 static mapping. [PR1443222](#)
- A port endian issue in SPU messages between SPC3 and SPC2 results in one redundant NAT binding being created in central point when one binding is allocated in SPC2 SPC. [PR1450929](#)
- Packet loss is observed when multiple source NAT pools and rules are configured. [PR1457904](#)

Network Management and Monitoring

- MIB OID `dot3StatsDuplexStatus` shows wrong status. [PR1409979](#)
- Snmpd process might generate core files after restarting NSD process by using the **restart network-security gracefully** command. [PR1443675](#)
- Control links are logically down on SRX Series devices with chassis cluster running Junos OS Release 12.3X48. [PR1458314](#)

Platform and Infrastructure

- On SRX4600 platform, when manual RGO failover is performed, sometimes node0 (the original primary node) stays in secondary-hold status for a long time and cannot change back to secondary status. [PR1421242](#)
- Packet drops, replication failure, or ksyncd stops might be seen on the logical system of a Junos OS device after Routing Engine switchover. [PR1427842](#)
- The PICs might go offline and split brain might be seen when interrupt storm happens on internal Ethernet interface em0 or em1. [PR1429181](#)
- REST API does not work properly. [PR1430187](#)
- Unable to launch J-Web when the device is upgraded through USB image. [PR1430941](#)
- Packet Forwarding Engine crashes might be seen on SRX1500 platform. [PR1431380](#)
- The ksyncd process might stop and restart. [PR1440576](#)
- The configured RPM probe server hardware timestamp does not respond with the correct timestamp to the RPM client. [PR1441743](#)

- ARP resolution might fail after **ARP HOLD NHs** are added and deleted continuously [PR1442815](#)
- The SRX300 line of device does not have MIB that can retrieve the fan status. [PR1443649](#)
- IS-IS adjacencies between the GE link are not up. [PR1446533](#)
- The flowd process might stop on SRX Series devices when chassis cluster and IRB interface are configured. [PR1446833](#)
- The **show security flow session** command fails with error messages when SRX4100 or SRX4200 has around 1 million routing entries in FIB. [PR1445791](#)
- LACP cannot work with the encapsulation **flexible-ethernet-services** configuration. [PR1448161](#)
- On certain MPC line cards, cm errors need to be reclassified. [PR1449427](#)
- The REST service might become nonresponsive when the REST API receives several continuous HTTP requests. [PR1449987](#)
- VM core files might be generated if the configured sampling rate is more than 65,535. [PR1461487](#)
- Loading CA certificate causes pkid core file to be generated. [PR1465966](#)

Routing Policy and Firewall Filters

- The NSD process might stop due to a memory corruption issue. [PR1419983](#)
- Two ipfd processes appear in **ps** command and the process pauses. [PR1444472](#)
- During commit, the **nsd_vrf_group_config_ls** log messages are displayed. [PR1446303](#)
- Traffic log shows wrong custom-application name when the **alg ignore** option is used in application configuration. [PR1457029](#)
- The NSD process might get stuck and cause problems. [PR1458639](#)
- The policy detail does not print out policy statistics counter, even when policy count is enabled. [PR1471621](#)

Services Applications

- The flowd process stops when the SRX5000 line of devices works in SPC3 mixed mode with one SPC3 card or seven SPC2 cards. [PR1448395](#)
- The srxpfe lcore-slave core files are generated. [PR1460035](#)

Unified Threat Management

- The **show security utm web-filtering status** command now provides additional context when the status of EWF is **down**. [PR1426748](#)

- Memory issue due to SSL proxy whitelist or whitelist URL category. [PR1430277](#)
- Adjust core allocation ratio for on-box antivirus. [PR1431780](#)

VLAN Infrastructure

- ISSU failed from Junos OS Release 18.4R2.7 to Junos OS Release 19.4, with secondary node PICs in present state after upgrading to Junos OS Release 19.4. [PR1468609](#)

VPNs

- IPsec SA inconsistent on SPCs of node0 and node1 in SRX Series devices with chassis cluster. [PR1351646](#)
- After RG1 failover, IKE phase 1 SA is getting cleared. [PR1352457](#)
- With a large number of IPsec tunnels established, a few tunnels might fail during rekey negotiation if the SRX Series device initiates the rekey. [PR1389607](#)
- Displaying incorrect port number when scale is 1,000 on IKEv1 AutoVPN tunnels. [PR1399147](#)
- The IKE and IPsec configuration under groups is not supported in this release. [PR1405840](#)
- The IKED process stops due to a misconfiguration. [PR1416081](#)
- The VPN tunnel might flap when IKE and IPsec rekey happen simultaneously. [PR1421905](#)
- Old tunnel entries are also seen when new tunnel negotiation happens from peer device after change in IKE gateway configuration at peer side. [PR1423821](#)
- IPsec packet throughput might be impacted if NAT-T is configured and the fragmentation operation of post fragment happens. [PR1424937](#)
- Tunnel does not come up after changing configurations from IPv4 to IPv6 tunnels in the script with **gateway lookup failed** error. [PR1431265](#)
- P1 configuration delete message is not sent on loading baseline configuration if there has been a prior change in VPN configuration. [PR1432434](#)
- IPsec rekey triggers for when sequence number in AH and ESP packet is about to exhaust. [PR1433343](#)
- P1 or P2 SAs are deleted after RG0 failover. [PR1433355](#)
- IPsec SA in and out key sequence number update missing after cold synchronization. [PR1433424](#)
- Sequence number reset to zero while recovering SA after SPC3 or flowd stops or reboots. [PR1433568](#)
- The kmd log shows resource temporarily unavailable repeatedly and VPNs might be down. [PR1434137](#)
- The IKED process stops on SRX5000 line of devices with SPC3 when IPsec VPN or IKE is configured. [PR1443560](#)
- The IPsec VPN traffic drop might be seen on SRX Series devices with NAT-T scenario. [PR1444730](#)

- Sometimes old SAs are not deleted after rekey and the number of IPsec tunnels shows up more than the configured tunnels. [PR1449296](#)
- Some IPsec tunnels flap after RGs failover on SRX5000 line of devices. [PR1450217](#)
- The VPN flaps on the primary node after a reboot of the secondary node. [PR1455389](#)
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. [PR1455951](#)
- IPsec VPN tunnels are losing routes for traffic selector randomly while tunnel is still up, causing traffic loss of these IPsec VPN tunnels. [PR1456301](#)
- On all SRX Series devices, the **no-anti-replay** option does not take effect immediately. Traffic is not sent out through IPsec VPN after upgrading to Junos OS Release 18.2 or later. [PR1461793](#)
- The IPsec VPN tunnels cannot be established if overlapped subnets are configured in traffic selectors. [PR1463880](#)

SEE ALSO

[What's New | 247](#)

[What's Changed | 255](#)

[Known Limitations | 262](#)

[Open Issues | 263](#)

[Documentation Updates | 273](#)

[Migration, Upgrade, and Downgrade Instructions | 274](#)

Documentation Updates

IN THIS SECTION

- [Feature Guides Are Renamed As User Guides | 274](#)

This section lists the errata and changes in Junos OS Release 19.4R1 for the SRX Series documentation.

Feature Guides Are Renamed As User Guides

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

SEE ALSO

[What's New | 247](#)

[What's Changed | 255](#)

[Known Limitations | 262](#)

[Open Issues | 263](#)

[Resolved Issues | 266](#)

[Migration, Upgrade, and Downgrade Instructions | 274](#)

Migration, Upgrade, and Downgrade Instructions

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

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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 15.1X49, 17.3, 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from one Junos OS Release to the next release or one release after the next release. For example you can upgrade from Junos OS Release 15.1X49 to Release 17.3 or 17.4, Junos OS Release 17.4 to Release 18.1 or 18.2, and from Junos OS Release 18.1 to Release 18.2 or 18.3 and so on.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

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

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

SEE ALSO

What's New 247
What's Changed 255
Known Limitations 262
Open Issues 263
Resolved Issues 266
Documentation Updates 273

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 is a framework, set of policies, and tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that have been developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information on 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**—Determine the features supported on MX Series, PTX Series, QFX Series devices. The Juniper Networks Feature Explorer is a Web-based app that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. <https://pathfinder.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home


NOTE: To obtain information about the components that are supported on the devices, and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. apps.juniper.net/compliance/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

21 April 2023—Revision 29, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 December 2021—Revision 28, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 October 2021—Revision 27, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 July 2021—Revision 26, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 April 2021—Revision 25, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 March 2021—Revision 24, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 23, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 November 2020—Revision 22, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 October 2020—Revision 21, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 20, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 July 2020—Revision 19, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 June 2020—Revision 18, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 May 2020—Revision 17, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 May 2020—Revision 16, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

30 April 2020—Revision 15, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 April 2020—Revision 14, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 April 2020—Revision 13, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 March 2020—Revision 12, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 March 2020—Revision 11, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 March 2020—Revision 10, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

5 March 2020—Revision 9, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 February 2020—Revision 8, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 January 2020—Revision 7, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 January 2020—Revision 6, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 January 2020—Revision 5, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 January 2020—Revision 4, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 January 2020—Revision 3, Junos OS Release 19.4R1— ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 December 2019—Revision 2, Junos OS Release 19.4R1—ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 December 2019—Revision 1, Junos OS Release 19.4R1—ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.