

Release Notes: Junos[®] OS Release 19.2R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

22 April 2021

Contents	Introduction 10
	Junos OS Release Notes for ACX Series 10
	What's New 11
	What's New in Release 19.2R3 11
	What's New in Release 19.2R2 11
	What's New in Release 19.2R1-S1 12
	What's New in Release 19.2R1 12
	What's Changed 16
	What's Changed in 19.2R3-S1 17
	What's Changed in 19.2R3 17
	What's Changed in 19.2R2 17
	What's Changed in 19.2R1 18
	Known Limitations 19
	General Routing 19
	Open Issues 22
	General Routing 22
	MPLS 25

Resolved Issues | 25

Resolved Issues: 19.2R3 | 26

Resolved Issues: 19.2R2 | 26

Resolved Issues: 19.2R1-S1 | 29

Resolved Issues: 19.2R1 | 30

Documentation Updates | 31

Installation and Upgrade Guide | 31

Migration, Upgrade, and Downgrade Instructions | 32

Upgrade and Downgrade Support Policy for Junos OS Releases | 32

Junos OS Release Notes for EX Series Switches | 34

What's New | 34

What's New in Release 19.2R3 | 35

What's New in Release 19.2R2 | 35

What's New in Release 19.2R1-S1 | 36

What's New in Release 19.2R1 | 36

What's Changed | 41

What's Changed in Release 19.2R3 | 41

What's Changed in Release 19.2R2 | 41

What's Changed in Release 19.2R1 | 43

Known Limitations | 44

EVPN | 45

Platform and Infrastructure | 45

Open Issues | 45

Authentication and Access Control | 46

Infrastructure | 46

Interfaces and Chassis | 47

Junos Fusion for Enterprise | 47

Layer 2 Features | 47

Network Management and Monitoring | 47

Platform and Infrastructure | 47

Routing Protocols | 48

Resolved Issues | 49

Resolved Issues: 19.2R3 | 49

Resolved Issues: 19.2R2 | 51

Resolved Issues: 19.2R1	59
Documentation Updates	63
Installation and Upgrade	63
Migration, Upgrade, and Downgrade Instructions	63
Upgrade and Downgrade Support Policy for Junos OS Releases	64
Junos OS Release Notes for Junos Fusion Enterprise	65
What's New	65
What's New in Release 19.2R3	66
What's New in Release 19.2R2	66
What's New in Release 19.2R1	66
What's Changed	67
What's Changed in 19.2R3	67
What's Changed in 19.2R2	67
What's Changed in 19.2R1	67
Known Limitations	68
Open Issues	68
Junos fusion for enterprise	68
Resolved Issues	69
Resolved Issues: 19.2R3	70
Resolved Issues: 19.2R2	70
Resolved Issues: 19.2R1	70
Documentation Updates	71
Migration, Upgrade, and Downgrade Instructions	71
Basic Procedure for Upgrading Junos OS on an Aggregation Device	71
Upgrading an Aggregation Device with Redundant Routing Engines	73
Preparing the Switch for Satellite Device Conversion	74
Converting a Satellite Device to a Standalone Switch	75
Upgrade and Downgrade Support Policy for Junos OS Releases	75
Downgrading from Junos OS	76
Junos OS Release Notes for Junos Fusion Provider Edge	77
What's New	77
What's New in Release 19.2R3	78
What's New in Release 19.2R2	78

What's New in Release 19.2R1 | 78

What's Changed | 78

Known Limitations | 79

Open Issues | 79

Junos Fusion Provider Edge | 80

Resolved Issues | 80

Resolved Issues: 19.2R3 | 81

Resolved Issues: 19.2R2 | 81

Resolved Issues: 19.2R1 | 81

Documentation Updates | 82

Migration, Upgrade, and Downgrade Instructions | 82

Basic Procedure for Upgrading an Aggregation Device | 83

Upgrading an Aggregation Device with Redundant Routing Engines | 85

Preparing the Switch for Satellite Device Conversion | 86

Converting a Satellite Device to a Standalone Device | 87

Upgrading an Aggregation Device | 90

Upgrade and Downgrade Support Policy for Junos OS Releases | 90

Downgrading from Junos OS Release 19.2 | 90

Junos OS Release Notes for MX Series 5G Universal Routing Platform | 91

What's New | 92

What's New in 19.2R3 | 92

What's New in 19.2R2 | 93

What's New in 19.2R1-S4 | 93

What's New in 19.2R1-S1 | 94

What's New in 19.2R1 | 95

What's Changed | 113

What's Changed in Release 19.2R3-S1 | 114

What's Changed in Release 19.2R3 | 114

What's Changed in Release 19.2R2 | 115

What's Changed in Release 19.2R1 | 120

Known Limitations | 123

General Routing | 124

Interfaces and Chassis | 126

Platform and Infrastructure | 127

[Routing Protocols | 127](#)

[Open Issues | 127](#)

[EVPN | 128](#)

[Forwarding and Sampling | 128](#)

[General Routing | 129](#)

[Infrastructure | 134](#)

[Interfaces and Chassis | 134](#)

[Junos Fusion for Provider Edge | 135](#)

[Layer 2 Features | 135](#)

[Layer 2 Ethernet Services | 135](#)

[MPLS | 135](#)

[Network Management and Monitoring | 136](#)

[Platform and Infrastructure | 136](#)

[Routing Protocols | 137](#)

[User Interface and Configuration | 138](#)

[VPNs | 138](#)

[Resolved Issues | 139](#)

[Resolved Issues: 19.2R3 | 139](#)

[Resolved Issues: 19.2R2 | 146](#)

[Resolved Issues: 19.2R1 | 175](#)

[Documentation Updates | 191](#)

[Installation and Upgrade Guide | 191](#)

[Subscriber Management Provisioning Guide | 191](#)

[Migration, Upgrade, and Downgrade Instructions | 192](#)

[Basic Procedure for Upgrading to Release 19.2 | 193](#)

[Procedure to Upgrade to FreeBSD 11.x based Junos OS | 193](#)

[Procedure to Upgrade to FreeBSD 6.x based Junos OS | 195](#)

[Upgrade and Downgrade Support Policy for Junos OS Releases | 197](#)

[Upgrading a Router with Redundant Routing Engines | 198](#)

[Downgrading from Release 19.2 | 198](#)

[Junos OS Release Notes for NFX Series | 199](#)

[What's New | 199](#)

[What's New in Release 19.2R3 | 200](#)

[What's New in Release 19.2R2 | 200](#)

What's New in Release 19.2R1	200
What's Changed	201
What's Changed in Release 19.2R3	202
What's Changed in Release 19.2R2	202
What's Changed in Release 19.2R1	202
Known Limitations	202
Interfaces	203
Platform and Infrastructure	203
Virtual Network Functions (VNFs)	204
Open Issues	204
Interfaces	205
Platform and Infrastructure	205
Routing Protocols	205
Virtual Network Functions (VNFs)	206
Resolved Issues	206
Resolved Issues: 19.2R3	207
Resolved Issues: 19.2R2	207
Resolved Issues: 19.2R1	209
Documentation Updates	210
Migration, Upgrade, and Downgrade Instructions	210
Upgrade and Downgrade Support Policy for Junos OS Releases	211
Basic Procedure for Upgrading to Release 19.2	211
Junos OS Release Notes for PTX Series Packet Transport Routers	212
What's New	213
What's New in 19.2R3	213
What's New in 19.2R2	214
New and Changed Features: 19.2R1-S4	214
New and Changed Features: 19.2R1-S1	215
New and Changed Features: 19.2R1	215
What's Changed	221
What's Changed in Release 19.2R3	222
What's Changed in Release 19.2R2	222
What's Changed in Release 19.2R1-S5	224
What's Changed in Release 19.2R1	224

Known Limitations | 227**General Routing | 227****Interfaces and Chassis | 228****Open Issues | 228****General Routing | 229****Interfaces and Chassis | 230****Layer 2 Ethernet Services | 230****Routing Protocols | 230****Resolved Issues | 230****Resolved Issues: 19.2R3 | 231****Resolved Issues: 19.2R2 | 231****Resolved Issues: 19.2R1 | 235****Documentation Updates | 238****Installation and Upgrade Guide | 238****Migration, Upgrade, and Downgrade Instructions | 238****Basic Procedure for Upgrading to Release 19.2 | 239****Upgrade and Downgrade Support Policy for Junos OS Releases | 241****Upgrading a Router with Redundant Routing Engines | 242****Junos OS Release Notes for the QFX Series | 243****What's New | 243****What's New in Release 19.2R3 | 244****What's New in Release 19.2R2 | 244****What's New in Release 19.2R1-S1 | 244****What's New in Release 19.2R1 | 245****What's Changed | 253****What's Changed in Release 19.2R3 | 254****What's Changed in Release 19.2R2 | 254****What's Changed in Release 19.2R1 | 257****Known Limitations | 259****EVPN | 259****Layer 2 Features | 259****Platform and Infrastructure | 259****Routing Protocols | 260**

Open Issues | 261

EVPN | 261

Infrastructure | 262

Interfaces and Chassis | 262

Layer 2 Features | 262

Platform and Infrastructure | 262

Routing Protocols | 264

Virtual Chassis | 264

Resolved Issues | 265

Resolved Issues: 19.2R3 | 265

Resolved Issues: 19.2R2 | 267

Resolved Issues: 19.2R1 | 281

Documentation Updates | 287

Installation and Upgrade guide | 287

Migration, Upgrade, and Downgrade Instructions | 288

Upgrading Software on QFX Series Switches | 288

Installing the Software on QFX10002-60C Switches | 291

Installing the Software on QFX10002 Switches | 291

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 | 292

Installing the Software on QFX10008 and QFX10016 Switches | 294

Performing a Unified ISSU | 298

Preparing the Switch for Software Installation | 299

Upgrading the Software Using Unified ISSU | 299

Upgrade and Downgrade Support Policy for Junos OS Releases | 301

Junos OS Release Notes for SRX Series | 302

What's New | 303

New and Changed Features: 19.2R3 | 303

New and Changed Features: 19.2R2 | 303

New and Changed Features: 19.2R1-S1 | 304

New and Changed Features: 19.2R1 | 304

What's Changed | 313

Release 19.2R3 Changes in Behavior and Syntax | 314

Release 19.2R2 Changes in Behavior and Syntax | 314

Release 19.2R1 Changes in Behavior and Syntax	316
Known Limitations	317
J-Web	318
VPNs	318
Open Issues	319
Chassis Clustering	319
Flow-Based and Packet-Based Processing	319
Intrusion Detection and Prevention (IDP)	320
J-Web	320
Platform and Infrastructure	320
Routing Policy and Firewall Filters	320
VPNs	320
Resolved Issues	321
Resolved Issues: 19.2R3	321
Resolved Issues: 19.2R2	323
Resolved Issues: 19.2R1	333
Documentation Updates	340
Migration, Upgrade, and Downgrade Instructions	341
Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases	341
Upgrading Using ISSU	342
Licensing	342
Compliance Advisor	342
Finding More Information	343
Documentation Feedback	343
Requesting Technical Support	344
Self-Help Online Tools and Resources	344
Opening a Case with JTAC	345
Revision History	345

Introduction

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

These release notes accompany Junos OS Release 19.2R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 11](#)
- [What's Changed | 16](#)
- [Known Limitations | 19](#)
- [Open Issues | 22](#)
- [Resolved Issues | 25](#)
- [Documentation Updates | 31](#)
- [Migration, Upgrade, and Downgrade Instructions | 32](#)

These release notes accompany Junos OS Release 19.2R3 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [What's New in Release 19.2R3 | 11](#)
- [What's New in Release 19.2R2 | 11](#)
- [What's New in Release 19.2R1-S1 | 12](#)
- [What's New in Release 19.2R1 | 12](#)

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

What's New in Release 19.2R3

There are no new features or enhancements to existing features for ACX Series routers in Junos OS Release 19.2R3.

What's New in Release 19.2R2

There are no new features on ACX Series in Release 19.2R2.

What's New in Release 19.2R1-S1

Routing Protocols

- **Decouple RSVP for IGP-TE (MX Series, PTX Series, ACX Series, QFX Series, SRX Series, and EX Series)**—Starting in Junos OS Release 19.2R1-S1, 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.

What's New in Release 19.2R1

Class of Service (CoS)

- **Support for class of service (CoS)(ACX6360 routers)**—Starting in Junos OS Release 19.2R1, ACX6360 routers support class of service (CoS) functionality.

CoS is the assignment of traffic flows to different service levels. Service providers can use router-based CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows.

[See [CoS on ACX Series Universal Metro Routers Features Overview](#).]

EVPN

- **EVPN support of VLAN ID ranges and lists in service provider style interface configurations (EX9200 switches, ACX5448 and MX Series routers, and vMX virtual routers)**—Starting in Junos OS Release 19.2R1, EX9200 switches, ACX5448 and MX Series routers, and vMX virtual routers support the use of VLAN ID ranges and lists in a service provider style interface configuration, which must be referenced in an EVPN routing instance. This configuration is supported with the following EVPN environments, services, and features:
 - Environments:
 - EVPN with VXLAN encapsulation
 - EVPN with MPLS encapsulation
 - VLAN bundle service:
 - E-LAN
 - E-Tree
 - E-Line
 - Feature:
 - EVPN multihoming:
 - All-active
 - Single-active

- Singlehoming

[See [VLAN ID Ranges and Lists in an EVPN Environment](#).]

Interfaces and Chassis

- **Support for 100-Mbps and 1-Gbps speeds on Tri-Rate Copper SFP (ACX5448 routers)**—Starting in Junos OS Release 19.2R1, ACX5448 routers support 100-Mbps and 1-Gbps speeds on Tri-Rate Copper SFP optics (part number 740-013111).

NOTE: 100-Mbps speed is supported only on ports xe-0/0/24 through xe-0/0/47.

10-Mbps speed is not supported on Tri-Rate Copper SFP due to hardware limitations.

- To set the speed for the optics, issue the **set interfaces *interface-name* speed auto** command. [See [Speed](#) for more details.]
- To enable autonegotiation, issue the **set interfaces *interface-name* gigether-options auto-negotiation** command. [See [auto-negotiation](#).]

Junos Telemetry Interface

- **Support for LSP statistics on JTI (ACX6360)**—Starting with Junos OS Release 19.2R1, you can provision the LSP statistics sensor using the resource path **/junos/services/label-switched-path/usage/** to monitor per-MPLS LSP statistics on the ACX6360 router and export telemetry data through Junos telemetry interface (JTI) to external collectors. You can stream data at configurable intervals through gRPC without involving polling.

JTI support is only for RSVP LSPs.

Statistics that are streamed are similar to the output displayed by the operational mode command **show mpls lsp bypass statistics**.

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

To enable statistics for export from the Junos OS, include the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level.

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

- **Specify Routing Instance for JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R1, you can specify the routing instance to use for remote procedure call (gRPC) services. Include the **routing-instance *instance-name*** at the **[edit system services extension-service request-response grpc]** hierarchy level. The routing instance name specified should match the name of

the existing routing instance, such as a name configured under the **[routing-instances]** hierarchy level or **mgmt_junos** if **system management-instance** is configured (the dedicated management routing instance).

Configuring the routing instance lets you choose the VRF for gRPC services. When the routing instance is not configured, the default behavior is that all gRPC-related services are available through the management **fxp0/em0** interface.

Layer 3 Features

- **Support for Layer 3 unicast features (ACX 6360)**—Starting in Junos OS Release 19.2R1, ACX routers support the following Layer 3 forwarding features for unicast IPv4 and IPv6 traffic:
 - Basic IPv6 forwarding
 - Virtual router (VRF-lite) for both IPv4 and IPv6
 - Layer 3 subinterfaces support for both IPv4 and IPv6
 - VRF-lite, subinterfaces, and IPv6 forwarding support on link aggregation groups (LAGs)
 - Statistics support for Layer 3 subinterfaces
 - 32-way equal-cost multipath (ECMP)
 - Centralized Bidirectional Forwarding Detection (BFD)
 - IPv4 Layer 3 protocols:
 - OSPF
 - IS-IS
 - BGP
 - IPv6 Layer 3 protocols:
 - OSPFv3

- RIPng

Network Management and Monitoring

- **Support for displaying valid user input in the CLI for command options and configuration statements in custom YANG data models (ACX Series)**—Starting in Junos OS Release 19.2R1, the CLI displays the set of possible values for a given command option or configuration statement in a custom YANG data model when you include the **action-expand** extension statement in the option or statement definition and reference a script that handles the logic. The **action-expand** statement must include the **script** child statement, which defines the Python action script that is invoked when a user requests context-sensitive help in the CLI for the value of that option or statement.

[See [Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules](#).]

Software Installation and Upgrade

- **Zero Touch Provisioning (ACX5448)**—Starting in Junos OS Release 19.2R1, Zero Touch Provisioning (ZTP) automates the provisioning of the device configuration and software image with minimal manual intervention on management interface **em0**.

When you physically connect a router to the network and boot it with a factory configuration, the router upgrades the Junos OS software image automatically and automatically installs a configuration file from the network through the management interface.

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

System Management

- **Support for transferring accounting statistics files and router configuration archives using HTTP URL (ACX Series)**—Starting in Junos OS Release 19.2R1, you can transfer accounting statistics files and router configuration archives to remote servers by using an HTTP URL. In addition to SCP and FTP, the following HTTP URL will be supported under the **archive-sites** statement:

`http://username@host:url-path password password`

- To transfer accounting statistics files, configure **archive-sites** under **[edit accounting-options file <filename>]** hierarchy.
- To transfer router configuration archival, configure **archive-sites** under **edit system archival configuration** hierarchy.
- To view the statistics of transfer attempted, succeeded, and failed, use the **show accounting server statistics archival-transfer** command.
- To clear the statistics of transfer attempted, succeeded, and failed, use the **clear accounting server statistics archival-transfer** command.

[See [archive-sites](#), [Backing Up Configurations to an Archive Site](#), [show accounting server statistics archival-transfer](#), and [clear accounting server statistics archival-transfer](#)].

- **Precision Time Protocol (PTP) Transparent Clock with IPv6 Transport (PTX10001-20C and ACX6360-OR devices)**—Starting with Junos OS Release 19.2R1, PTP uses IPv6 transport to synchronize clocks throughout a packet-switched network. With a transparent clock, the PTP packets are updated with the residence time as the packets pass through the switch. There is no master/slaved designation. End-to-end transparent clocks are supported. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that the timestamps are sent in one packet.

You can configure the transparent clock at the **[edit protocols ptp]** Junos OS CLI hierarchy.

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

SEE ALSO

[What's Changed | 16](#)

[Known Limitations | 19](#)

[Open Issues | 22](#)

[Resolved Issues | 25](#)

[Documentation Updates | 31](#)

[Migration, Upgrade, and Downgrade Instructions | 32](#)

What's Changed

IN THIS SECTION

- [What's Changed in 19.2R3-S1 | 17](#)
- [What's Changed in 19.2R3 | 17](#)
- [What's Changed in 19.2R2 | 17](#)
- [What's Changed in 19.2R1 | 18](#)

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

What's Changed in 19.2R3-S1

General Routing

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

[See [gether-options](#).]

What's Changed in 19.2R3

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the `error` option at the `edit system services extension-service traceoptions level` hierarchy.

See [traceoptions \(Services\)](#).

What's Changed in 19.2R2

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Release 19.2R2 and later, ACX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

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

- **Support for disabling RS-FEC (ACX6360-OX)**—By default, Junos OS software enables or disables forward error correction based on plugged-in optics. Starting with Junos OS Release 19.2R2, on ACX6360-OX routers functioning as transponders, you can now disable Ethernet FEC, also known as RS-FEC or FEC91. Previously, RS-FEC was enabled by default and could not be disabled.

[See [fec](#).]

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when a change in number of supported history records is detected (ACX Series and MX Series)**—In Junos OS Release 19.2R2, when Ethernet connectivity fault management starts, it detects the number of history records supported by the existing Performance Monitoring history database and if there is any change from the number of history records supported (that is, 12) in Release 19.2R2, then the existing performance monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

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 19.2R2, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.2R2, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

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

What's Changed in 19.2R1

Interfaces and Chassis

- **Monitoring information available only in trace log (ACX Series)**—In Junos OS Release 19.2R1 and later, the Ethernet link fault management daemon (lfmd) in the peer router stops monitoring the locally occurred errors until ISSU completes. You can view the monitoring-related details only through the trace log file.

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (ACX Series)**—Starting in Junos OS Release 19.2R1, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.
- **Custom YANG RPC support for input parameters of type empty (ACX Series)**—Starting in Junos OS Release 19.2R1, custom YANG RPCs support input parameters of type **empty** when executing the RPC's command in the Junos OS CLI, and the value passed to the action script is the parameter name. In earlier releases, input parameters of type **empty** are only supported when executing the RPC in a NETCONF or Junos XML protocol session, and the value passed to the action script is the string '**none**'.

[See [Creating Action Scripts for YANG RPCs on Devices Running Junos OS.](#)]

VLAN Infrastructure

- **Specifying a descending VLAN ID range (ACX5448 routers)**—In Junos OS releases prior to Junos OS Release 19.2R1, the system accepts a descending range—for example, 102-100, with the **vlan-id-range** configuration statement in the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy.

Starting with Junos OS Release 19.2R1, the system considers a descending range specified with **vlan-id-range** to be invalid and raises an error if you try to commit this configuration.

SEE ALSO

What's New 11
Known Limitations 19
Open Issues 22
Resolved Issues 25
Documentation Updates 31
Migration, Upgrade, and Downgrade Instructions 32

Known Limitations

IN THIS SECTION

- [General Routing | 19](#)

Learn about known limitations in this release for ACX Series routers.

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

General Routing

- All PTP packets go to the best-effort queue instead of the network control queue. This is because of the limitation on ACX5448 where DSCP values are not preserved. [PR1361315](#)
- ACX6360 Junos telemetry interface or telemetry infrastructure does not support the interface-filtering capability. Therefore, after you enable a particular sensor for telemetry, it is turned on for all the interfaces. [PR1371996](#)

- For an Ethernet (et) interface, only the PRE_FEC_SD defect is raised and no OTN alarm is raised. [PR1371997](#)
- The CLI **static-cak** command encryption does not work between two ACX-OX transponder nodes. [PR1389802](#)
- The ACX6360 TIC has only 8 CFP2-DCO ports, so chassis beacon show/requests to port numbers larger than 7 do not work (as the ports don't exist) but also do not report an error. **user@router> request chassis beacon fpc 0 pic-slot 1 port 15 on FPC 0 PIC 1 PORT 15 ON user@router> show chassis beacon fpc 0 pic-slot 1 port-range lower-limit 0 upper-limit 15 FPC 0 PIC 1 PORT 0 ON FPC 0 PIC 1 PORT 1 ON FPC 0 PIC 1 PORT 2 ON FPC 0 PIC 1 PORT 3 ON FPC 0 PIC 1 PORT 4 ON FPC 0 PIC 1 PORT 5 ON FPC 0 PIC 1 PORT 6 ON FPC 0 PIC 1 PORT 7 ON FPC 0 PIC 1 PORT 8 ON FPC 0 PIC 1 PORT 9 ON FPC 0 PIC 1 PORT 10 OFF FPC 0 PIC 1 PORT 11 OFF FPC 0 PIC 1 PORT 12 OFF FPC 0 PIC 1 PORT 13 OFF FPC 0 PIC 1 PORT 14 OFF FPC 0 PIC 1 PORT 15 ON.** [PR1399335](#)
- When a timing configuration and the corresponding interface configuration are flapped for multiple times in iteration, PTP is stuck in "INITIALIZE" state where the ARP for the neighbor is not resolved. In issue state, BCM hardware block gets into inconsistency state, where the lookup is failing. [PR1410746](#)
- The input packet count given under the traffic statistics includes all packets that are coming in. The statistics are not segregated as IPv4, IPv6, MPLS, and so on. This is the same behavior across all the ACX Series platforms. [PR1419143](#)
- Hardware-based fragmentation or reassembly is not supported. Software-based fragmentation rates are going to be extremely slow depending on CPU load. [PR1419371](#)
- In the output of show SNMP mib walk jnxBoxAnatomy, the chassis CLEI code and contents model is reading data from the I2C bus and EEPROM. Because the fan is not present on the I2C bus and does not have EEPROM, fan data can't be displayed for chassis CLEI code and contents model. [PR1420639](#)
- There is no support on separate counters for egress-dropped packets. Counters are reflected as part of RED-dropped packets. [PR1427148](#)
- Problem description: When end device (fan tray CPLD) I2C line is grounded or pulled low, other device write/reads are failing. From our testing, we have concluded the following: We have verified isolation functionality when the device is in the problem state, it's working fine (verified through software simulation). But in this case where we grounded one of the fan tray CPLD i2c lines to verify the failed test case, then entire circuit will get stalled and it leads to write fail for other devices also such as the PEM, and temperature. This is a hardware limitation. [PR1427222](#)
- These error messages can be seen sometimes if the optics module is being unplugged in between the EEPROM read. This is expected and does not impact any functionality. [PR1429016](#)
- Packet rates are not seen for aggregated Ethernet logical interface. [PR1429590](#)
- Multicast packets are flooded in a BD if snooping is not enabled. If interfaces x and y belong to a BD, then all multicast packets will be flooded to both x and y interface. If packets are received from interface x, packets will be flooded to x and y in ingress but discarded in the egress path for interface x because packet is received from the same interface. But these packets are also counted in the VOQ and hence

we are seeing more queue statistics. It is a known hardware limitation. **monitor interface xe-0/0/30**Input packets: 177958 (64 pps) [0]Output packets: 357306 (128 pps) [0] **monitor interface xe-0/0/12**Input packets: 361161 (128 pps) [642]Output packets: 179878 (63 pps) [320] **root@rioxd-p2a-a> show interfaces queue xe-0/0/30** Queue: 0, Forwarding classes: best-effortQueued:Packets : 544032 192 pps. => Sum of 64 + 128pps **root@rioxd-p2a-a> show interfaces queue xe-0/0/12** Queue: 0, Forwarding classes: best-effortQueued:Packets : 550929 192 pps. => Sum of 64 + 128pps. [PR1429628](#)

- Any packet greater than the MTU size are accounted for oversized packets. Packets exceeding MTU sizes are not considered for Jabber check. [PR1429923](#)
- The statistics are accessed through Broadcom API, which is the same for both tagged and untagged packets. This cannot be changed in accordance with MX Series, because it is direct access from Broadcom without any statistics changes specific to tagging from ACX5448 side. It will impact other statistics if changes are made. [PR1430108](#)
- The port LEDs glowing during system/vmhost halt state is the expected behavior across all ACX Series platforms. Even the system LED glows during halt state. [PR1430129](#)
- These are initial transient messages seen and do not have any functional impact. [PR1430355](#)
- 1-Gigabit Ethernet interfaces are shown as "xe". Therefore, the cosmetic issue is observed with respect to autonegotiation parameters although there is no impact on functionality. [PR1430835](#)
- If Layer 2 VPN sessions have OAM control-channel option set to router-alert-label, the no-control-word option in L2VPN shouldn't be used for BFD sessions to come up. [PR1432854](#)
- BCM SDK currently does support statistics, and we see routes are getting reinstalled on a periodic basis. SDK does not support statistics unless we move to Flex mode in KBP. This is a product limitation today. [PR1435579](#)
- New rate of 1.8 MB/s and which takes 16-17 min to copy the ACX5448 image (1.9G image size) - ACX5448 rates are less than Misha because the rate limiter is in bps and does not support pps-based (hardware limitation from DNX)*Avg size is 512. Hence rate is approximately one third of Misha rate. In file copy cases, normal pkt size seen are 1500 pkt sizes. [PR1439960](#)
- The hold timer expiry is common across all platforms. It is not specific to ACX5448-D. [PR1439980](#)
- Remote loopback is not supported on ACX5448-D. [PR1443517](#)
- The PEM entries for jnxFruName SNMP index are shown twice. [PR1446215](#)
- ACX Series routers support only 900 joins of IGMPv3 users per second. [PR1448146](#)
- The 2000 EVPN IRB scale is not hitting due to hardware limitation of filter entries that can be installed for EVPN instances. We can support only 1000 scale for Junos OS release 19.2. This is Broadcom limitation and cannot be changed. [PR1461309](#)
- Counters for filtering based on DA MAC and SA MAC are not supported, because QMX doesn't have any separate counter to count matched or dropped packets with interface MAC address. [PR1463981](#)

SEE ALSO

[What's New | 11](#)[What's Changed | 16](#)[Open Issues | 22](#)[Resolved Issues | 25](#)[Documentation Updates | 31](#)[Migration, Upgrade, and Downgrade Instructions | 32](#)

Open Issues

IN THIS SECTION

[General Routing | 22](#)[MPLS | 25](#)

Learn about open issues in this release for ACX Series routers.

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

General Routing

- Forwarding when using nonexistent SSM map source address in IGMPv3 instead of pruning. This is a day-1 design issue and needs a design solution. [PR1126699](#)
- When Layer 3 packets are classified, DiffServ code points are not preserved but are getting lost at the egress interface because of a chipset limitation. [PR1322142](#)
- ACX5448: When a 1-Gigabit SFP transceiver is plugged into the router, autonegotiation is enabled by default. There is no functional impact. Only the output of the **show interfaces <intf-name> extensive** CLI command shows the autonegotiation field as disabled. [PR1343679](#)
- If the **set interfaces ae<x>xaggregated-ether-options link-speed <x>** configuration is used in the router, after reboot AE interface remains down. Error seen in logs: **/kernel: kernel did not add link ae1, link speeds differ 1000000000 10000000000 /kernel: bundle ae1.0: link xe-0/3/0 not added due to speed mismatch.** [PR1357012](#)
- Loopback status is not shown for OT interfaces on CLI (available from vty only). [PR1358017](#)

- The SD (Signal Degrade) threshold is normally lower than the SF threshold (that is, so that as errors increase, SD condition is encountered first). For the ACX6360 optical links there is no guard code to prevent the user from setting the SD threshold, above the SF threshold which would cause increasing errors to trigger the SF alarm before the SD alarm. This will not cause any issues on systems with correctly provisioned SD/SF thresholds. [PR1376869](#)
- ACX6360-OR: Enhancement is needed for FRR BER threshold SNMP support. [PR1383303](#)
- On ACX6360/PTX10001 router, Tx power cannot be configured using + sign. [PR1383980](#)
- More than expected traffic loss is observed during link failure (FRR convergence) and link restoration test on L3VPN traffic over LDP-OSPF MPLS LSP.

Steps to re-create:

1. Layer 3 VPN traffic was flowing end to end on the active path in steady state.
2. Link down state is induced on primary path and traffic is shifted to backup path. Observation: More than expected traffic loss is observed (around 1.5 seconds).
3. Link is restored and the traffic is moved to the active path. Observation: More than expected traffic loss is observed (around 21 seconds).

The traffic loss percentage is not consistent and varies across each run, and this is the worst-case traffic loss percentage observed. [PR1387834](#)

- The switchover time observed was more than 50 minutes under certain soak test conditions with an increased scale with a multiprotocol multirouter topology. [PR1387858](#)
- IGMP packets over L2 Circuit with Control-Word are dropped in ACX5048. [PR1394301](#)
- The ccc logs are not compressed after rotation. [PR1398511](#)
- A jnxIfOtnOperState trap notification is sent for all ot interfaces. [PR1406758](#)
- Layer 2 rewrite is happening on regular bridge domain and VLAN interfaces, although there are some service dependencies (VPLS in this case) due to which, the egress interface map table is not updated properly with the Layer 2 rewrite map ID. As a result, the rewrite does not work. [PR1414414](#)
- Clock class value is incorrect in default data (show ptp clock) when the slave interface is down in the PTP-OC device. [PR1416421](#)
- On ACX5448 devices, the zero-touch provisioning (ZTP) process will proceed with image upgrade even in situations when there is a mismatch between the platform name of the software image stored on FTP or ZTP servers and the actual platform where the ZTP process is being run. [PR1418313](#)
- The issue is not fixed in Junos OS Release 19.1 where inner VLAN tag classification will not work. It is not fixed in 18.3R2 release. Also not fixed in 19.2R1 Release. There was a behavior change introduced as part of [PR1307666](#), where the inner VLAN tag is popped out on the ingress side when an IP packet with double-tagged VLAN is received, and this change is needed for IP packets to work on proper transmit on the egress interface. [PR1422515](#)

- Under scale conditions, if multiple bridge-domains are deleted together, then sometimes Broadcom SDK returns error **Entry not found**. This indicates that the resource is already freed up. [PR1423308](#)
- On an ACX5448, the **request system reboot** command triggers a reboot on the host (Linux) instead of just being limited to Junos OS. [PR1426486](#)
- Traffic loss is seen if the configuration has /128 prefix routes and its limited to /128 only. This is due to a known issue tracked in [PR1445231](#). [PR1429833](#)
- Any packets greater than the MTU size are accounted for as oversized packets. Packets exceeding MTU sizes are not considered for Jabber check. [PR1429923](#)
- The port LEDs glowing during system/vmhost halt state is the expected behavior across all ACX Series platforms. Even the system LED glows during halt state. [PR1430129](#)
- Packets dropped due to MTU checks in the output interface are not accounted for MTU errors. All packets above MTU size are accounted for Oversized-packets in the input interface. [PR1430446](#)
- Protocols get forwarded when using a nonexisting SSM map source address in IGMPv3 instead of pruning. This is a day 1 design issue, and needs a design solution. [PR1435648](#)
- On ACX5448 box, after issuing deactivate/activate "class-of-service", traffic drop might be seen. [PR1436494](#)
- In certain test conditions, it was observed that L2VPN at a scale of 16,000 had issues when all VPNs were brought down and up. [PR1439471](#)
- Recovery of Junos volume is not possible from the OAM menu. [PR1446512](#)
- Drop profile maximum threshold might not be reached when the packet size is other than 1000 bytes. This is due to the current design limitation. [PR1448418](#)
- When an xe interface working in 1-Gigabit mode in ACX5448-D is added to a member link of an aggregated Ethernet interface, the speed of the aggregated Ethernet interface is incorrectly shown as 10 Gbps. There is no functional impact. This is a display issue. [PR1449887](#)
- It is not possible to form 125,000 IGMP groups with the ACX5448 router receiving 125,000 IGMP v2 reports per second. This is a product limitation from BCM and CPU host path queuing model. [PR1454465](#)
- The issue is seen during unified ISSU to Junos OS Releases 20.1, 20.2, and 19.4. The upgrade is successful but the forwarding plane (PFE) does not function. As a result, forwarding is affected. [PR1483959](#)
- On ACX Series platforms with shaping configured, after deactivating and activating CoS the shaping might not work and traffic drop would appear. [PR1488935](#)
- In PTP environment some vendor devices acting as slave expecting announce messages at an interval of -3 (8pps) from upstream master device. As of today announce message are configurable in range of 0 to 3. To support the above requirement engineering provided a hidden cli knob "set protocol ptp master announce-interval -3". In the networks/design where we have this requirement. we can configure the hidden cli otherwise regular cli which is in the range (0 to 3) can be configurable. Both the cli knobs are mutually exclusive, commit error is expected if both are configured. This new change is applicable to ACX platforms only excluding ACX5000. [PR1507782](#)

MPLS

- The default behavior of local reversion has changed from Junos OS Release 16.1 and that impacts the LSPs for which the ingress does not perform make-before-break. Junos OS does not perform make-before-break for no-cspf LSPs. [PR1401800](#)

SEE ALSO

What's New	 11
What's Changed	 16
Known Limitations	 19
Resolved Issues	 25
Documentation Updates	 31
Migration, Upgrade, and Downgrade Instructions	 32

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3](#) | [26](#)
- [Resolved Issues: 19.2R2](#) | [26](#)
- [Resolved Issues: 19.2R1-S1](#) | [29](#)
- [Resolved Issues: 19.2R1](#) | [30](#)

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

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

Resolved Issues: 19.2R3

General Routing

- Policer discarded count is also shown incorrectly to the enq count of the interface-queue but traffic behavior is fine as expected. [PR1414887](#)
- ACX5448-D interfaces support: The input bytes value in the **show interfaces extensive** command output is not at par with older ACX Series or MX Series devices. [PR1430108](#)
- CoS: Egress queue statistics are not applicable to ae interfaces on ACX5048. [PR1472467](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- Link does not come up when 100-Gigabit Ethernet is channelized to 4x25-Gigabit Ethernet interfaces. [PR1479733](#)
- During speed mismatch, QSFP28/QSFP+ the optics/cables might or might not work. [PR1494600](#)
- Outbound SSH connection flap or memory leak issue might be observed during pushing configuration to ephemeral DB with high rate. [PR1497575](#)
- SFW sessions are not getting updated on ms interfaces in ACX500-i. [PR1505089](#)
- PIC slot may be shut down less than 240 seconds because the over-temperature start time is handled incorrectly. [PR1506938](#)
- L2ALD crash is seen during stability test with traffic on scaled setup. [PR1517074](#)

MPLS

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

Routing Protocols

- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

VPNs

- The Layer 2 circuit neighbor might be stuck in RD state at one end of MG-LAG peer. [PR1498040](#)

Resolved Issues: 19.2R2

Class of Service (CoS)

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

General Routing

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

- On an ACX Series router, the LED on a Gigabit Ethernet interface goes down when the 10-Mbps speed is added. [PR1385855](#)
- Link fault signaling (LFS) is not working on ACX5448 10/40/100GbE interfaces. [PR1401718](#)
- Kernel memory leak in virtual-memory due to interface flaps (CVE-2020-1625). [PR1407000](#)
- High CPU consumption for fxpc processes with class-of-service changes on AE interfaces. [PR1407098](#)
- 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](#)
- ACX5448: 40G FEC on ACX5448 is default FEC is enabled need to align with our platforms MX/QFX where FEC is NONE. [PR1414649](#)
- ACX5448: BFD Timer values are not as per the configured 900ms with multiplier 3. The values are showing 6.000 with multiplier 3 instead for most of the sessions. [PR1418680](#)
- [ARP] ACX5448-D: 96000 ARPs are getting populated but only 47,000 next-hop entries are present. Therefore, around 50% packet drop is observed. [PR1426734](#)
- Drift messages in ACX2200, which is a PTP hybrid (PTP + Synchronous Ethernet) device. [PR1426910](#)
- The chassisd process might crash with unsupported HCoS configuration when MX104 is used as the fusion aggregation device. [PR1430076](#)
- On ACX5448, upon reboot of an MC_LAG peer, when the peer comes up (but before hardware comes up), there is a 10-20 second traffic hit on node1. [PR1430910](#)
- ACX5448-D interface support: After chassis control restart, load balancing on the child interfaces of an ae interface stops. [PR1431206](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- ACX5448 might malfunction in encapsulating small packets if egress link is 40G or 100G. [PR1434900](#)
- In ACX Series platforms, **no-vrf-propagate-ttl** might not work after the CoS configuration is deactivated and then activated. [PR1435791](#)
- LACP state might get stuck in 'Attached' state after disabling peer active members. [PR1439268](#)
- Packet drop might be seen on an ACX Series platform when chained composite next hop is enabled for L3VPN. [PR1439317](#)
- Interface on ACX1100 remains down when using SFP-1FE-FX (740-021487). [PR1439384](#)
- On ACX5448, DHCP packets are not transparent over Layer 2 circuit. [PR1439518](#)
- When the interface is flapped between channelized configurations (25GbE to 100GbE), the AE interface configuration is not cleaned up properly. [PR1441374](#)
- ACX5448: Packet buffer error from Packet Forwarding Engine leading to memory leak when IGMP is sent from NNI AC in Layer 2 circuit and VPLS. [PR1442901](#)
- RED drops might be seen after link flaps or CoS configuration changes. [PR1443466](#)

- ACX5448: The encapsulation flexible-vlan-tagging is not supported with the MPLS family; need to provide commit error. [PR1445046](#)
- ACX5448/18.3R1-S4.1 not performing proper dot1p CoS rewrite on interfaces configured with l2circuit/local-switching/family ccc. [PR1445979](#)
- In ACX Series routers, auto-exported routes between VRFs might not reply for ICMP echo requests. [PR1446043](#)
- Fans on an ACX5448-M might not be running at the correct speed. [PR1448884](#)
- Layer 2 circuit 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](#)
- FPC core files might be seen after changing the configuration of PTP or Synchronous Ethernet. [PR1451950](#)
- Platforms: ACX5448-D interfaces support: After the 100-Gbps and 40-Gbps interfaces are disabled, the Laser output power in **show interfaces diagnostics optics** shows some values. [PR1452323](#)
- ACX5448 FPC crashed due to segmentation fault. [PR1453766](#)
- Incorrect operating state displayed in SNMP trap for fan removal. [PR1455577](#)
- ACX5048 SNMP polling will be stalled after a link flap or an SFP transceiver replacement, and **ACX_COS_HALP(acx_cos_gport_sched_set_strict_priority:987): Failed to detach** logs will be seen. [PR1455722](#)
- ACX6360-OX: Enable the **gigether** option to configure Ethernet FEC on client ports. [PR1456293](#)
- ACX5448-D and ACX5448-M devices do not display airflow information and temperature sensors as expected. [PR1456593](#)
- ACX5448 L2VPN with encapsulation-type **ethernet** stops passing traffic after a random port is added with VLAN configuration. [PR1456624](#)
- The rpd crash might be seen if a BGP route is resolved over the same prefix protocol next hop in an inet.3 table that has both RSVP and LDP routes. [PR1458595](#)
- Route resolution is not happening when the packet size is 10,000. [PR1458744](#)
- Traffic might be silently dropped 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-D interfaces and optics support: Sometimes, when the AE interfaces are brought up, there are ARP resolution issues. [PR1461485](#)
- On ACX Series platforms, the LLDP neighbor not up on a LAG after software upgrade to Junos OS Release 18.2R3-S1. [PR1461831](#)
- Memory leak on l2cpd process might lead to l2cpd crash. [PR1469635](#)

- RED drop on interface even without any congestion. [PR1470619](#)
- Egress queue statistics are not applicable to ae interface on model ACX5048. [PR1472467](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- dcpfe core files are seen when disabling/enabling MACsec using Toby scripts. [PR1479710](#)
- ACX5448 Layer 2 VPN with **interface ethernet-ccc input-vlan-map/output-vlan-map** can cause traffic to be dropped silently. [PR1485444](#)

Interfaces and Chassis

- Upgrade from releases before Junos OS Release 17.4R1 results in cfmd core files. [PR1425804](#)
- MC-AE interface might show as unknown status if you are adding the subinterface as part of the VLAN on the peer MC-AE node. [PR1479012](#)

Layer 2 Ethernet Services

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

Platform and Infrastructure

- The REST API process becomes nonresponsive when a number of requests come at a high rate. [PR1449987](#)

Routing Protocols

- Export of loopback address to other VRF instances might not work on ACX Series, EX Series and QFX Series platforms. [PR1449410](#)
- MPLS LDP might still use stale MAC addresses of the neighbor even if the LDP neighbor's MAC address changes. [PR1451217](#)
- The rpd might crash continuously due to memory corruption in IS-IS setup. [PR1455432](#)
- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

VPN

- The l2circuit neighbor might be stuck in RD state at one end of MG-LAG peer. [PR1498040](#)

Resolved Issues: 19.2R1-S1

General Routing

- Link Fault Signaling (LFS) do not work on ACX5448, ACX5410, ACX5440, and 100-Gigabit Ethernet interfaces. [PR1401718](#)

- In an ACX5448 platforms, when the Packet Forwarding Engine failed to allocate packet buffer, portion of packet memories might not be free. [PR1442901](#)

Resolved Issues: 19.2R1

Class of Service (CoS)

- The error message **STUCK_BUFF : port_sp not empty for port 35 sp 1 pkts:1** is seen when a lag bundle is configured with 64 lag links.. [PR1346452](#)

General Routing

- The 1G copper module interface shows "Link-mode: Half-duplex". [PR1286709](#)
- On an ACX ring topology, after link between ACX and MX flaps, VPLS RI on PE (MX) has no MAC of CE over I2circuit. [PR1360967](#)
- ACX5000: **fpc0 (acx_rt_ip_uc_lpm_install:LPM route add failed** error) Reason : Invalid parameter after configuring lpm-profile. [PR1365034](#)
- ACX5448: **LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** prints while committing on configuration prompt. [PR1376665](#)
- On ACX5448, channelized ET interface of 25-Gigabit interface will not come up after chassis-control restart. [PR1379288](#)
- ACX 5448:100 Gigabit link FEC enabled by default on 100G LR4. [PR1389518](#)
- On ACX Series platforms, the **forwarding-option dhcp-relay forward-only** command stops working and the DHCP packets are dropped. [PR1392261](#)
- On ACX5048, RPM RFC 2544 benchmarking test failed to start. [PR1395730](#)
- CFM adjacency is not going down with distinct intervals. [PR1397883](#)
- Dynamic tunnels are not supported on ACX Series routers. [PR1398729](#)
- VLAN tagged traffic arriving on VPLS interface might get dropped. [PR1402626](#)
- ot/et interface is not created when invalid speed is configured. [PR1403546](#)
- ACX 5448: TrTCM Policer configuration parameters are as per RFC4115. [PR1405798](#)
- The **show services inline stateful-firewall flow** or **show services inline stateful-firewall flow extensive** command might cause a memory leak. [PR1408982](#)
- ACX Series routers drop DNS responses that contain an underscore. [PR1410062](#)
- VPLS traffic might stop across ACX5000 with the aggregated Ethernet interface. [PR1412042](#)
- Junos PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- Number of inet-arp policers implemented on ACX5000 has been increased from 16 to 64. [PR1413807](#)
- Swap memory is not initialized on boot on ACX5048. [PR1415898](#)

- Commit error while configuring firewall with term having log/syslog and accept actions. [PR1417377](#)
- CoS table error can sometimes cause traffic outages and SNMP timeouts if the optic is plugged out and inserted back. [PR1418696](#)
- Slow copy image speed to ACX5448. [PR1422544](#)

SEE ALSO

What's New 11
What's Changed 16
Known Limitations 19
Open Issues 22
Documentation Updates 31
Migration, Upgrade, and Downgrade Instructions 32

Documentation Updates

IN THIS SECTION

- [Installation and Upgrade Guide | 31](#)

This section lists the errata and changes in Junos OS Release 19.2R3 for the ACX Series documentation.

Installation and Upgrade Guide

- **Veriexec explained (ACX Series)**—Verified Exec (also known as veriexec) is a file-signing and verification scheme that protects the Junos operating system (OS) against unauthorized software and activity that might compromise the integrity of your device. Originally developed for the NetBSD OS, veriexec was adapted for Junos OS and enabled by default from Junos OS Release 7.5 onwards.
[See [Veriexec Overview](#).]

SEE ALSO

What's New 11
What's Changed 16
Known Limitations 19
Open Issues 22
Resolved Issues 25
Migration, Upgrade, and Downgrade Instructions 32

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 32

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

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

SEE ALSO

What's New	 11
What's Changed	 16
Known Limitations	 19
Open Issues	 22
Resolved Issues	 25
Documentation Updates	 31

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

- What's New | 34
- What's Changed | 41
- Known Limitations | 44
- Open Issues | 45
- Resolved Issues | 49
- Documentation Updates | 63
- Migration, Upgrade, and Downgrade Instructions | 63

These release notes accompany Junos OS Release 19.2R3 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.


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

What's New

IN THIS SECTION

- What's New in Release 19.2R3 | 35
- What's New in Release 19.2R2 | 35
- What's New in Release 19.2R1-S1 | 36
- What's New in Release 19.2R1 | 36

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 19.2R2: EX2300, EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in Release 19.2R3

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.2R3.

What's New in Release 19.2R2

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 19.2R2.

What's New in Release 19.2R1-S1

Routing Protocols

- **Decouple RSVP for IGP-TE (MX Series, PTX Series, ACX Series, QFX Series, SRX Series, and EX Series)**—Starting in Junos OS Release 19.2R1-S1, 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.

What's New in Release 19.2R1

Authentication, Authorization, and Accounting (AAA)

- **802.1X authentication (EX4650 switches)**—Starting with Junos OS Release 19.2R1, EX4650 switches support port-based network access control using 802.1X authentication as defined in the IEEE 802.1X standard.

[See [802.1X for Switches Overview](#).]

Dynamic Host Configuration Protocol

- **Support for DHCP snooping and other access port security features on private VLANs (EX4300-MP switches and Virtual Chassis)**—Starting in Junos OS Release 19.2R1, you can enable DHCP snooping for security purposes on access ports that are in a private VLAN (PVLAN). You can also protect those ports with DHCP options, dynamic ARP inspection (DAI), IP source guard, and neighbor discovery inspection.

[See [Putting Access Port Security on Private VLANs](#).]

EVPN

- **Support for BFD, BGP, IS-IS, and OSPF on IRB interfaces in EVPN-MPLS networks (EX series)**—Starting with Junos OS Release 19.2R1, you can configure Bidirectional Forwarding Detection (BFD), BGP, IS-IS, and OSPF routing protocols on the IRB interface in an EVPN-MPLS network to route and forward EVPN traffic. This feature supports single-homed, single-active, and all-active multihomed networks.

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

- **EVPN support of VLAN ID ranges and lists in service provider style interface configurations (EX9200 switches)**—Starting in Junos OS Release 19.2R1, EX9200 switches, ACX5448 and MX Series routers, and vMX virtual routers support the use of VLAN ID ranges and lists in a service provider style interface configuration, which must be referenced in an EVPN routing instance. This configuration is supported with the following EVPN environments, services, and features:

- Environments:
 - EVPN with VXLAN encapsulation
 - EVPN with MPLS encapsulation
- VLAN bundle service:

- E-LAN
- E-Tree
- E-Line
- Features:
 - EVPN multihoming:
 - All-active
 - Single-active
 - Singlehoming

[See [VLAN ID Ranges and Lists in an EVPN Environment](#).]

- **Support for control word in EVPN-VPWS (EX9200 switches)**—Starting with Junos OS Release 19.2R1, Junos OS supports the insertion of a control word between the label stack and the MPLS payload in a network with EVPN-VPWS service. This feature prevents a transit device from delivering out-of-order packets as a result of the device's load-balancing hashing algorithm. When you enable the control word feature on a PE device, the PE device advertises support for a control word. If all the PE devices in an EVI on the EVPN-VPWS serviced network support control word, then the PE device inserts a control word between the label stack and the L2 header in the packet thus preventing the packet from being misidentified by transit devices.

[See [Control Word for EVPN-VPWS](#).]

JWeb

- **Support for EX4650 switches**—Starting in Junos OS Release 19.2R1, you can use J-Web to configure, monitor, and manage EX4650 switches.

To configure the EX4650 switch using the J-Web interface, you must connect the cable to the port labeled **CON** on the rear panel of the switch.

NOTE: In J-Web, the chassis viewer displays only the standalone EX4650 switches view. It does not display the Virtual Chassis configuration because the EX4650 switch does not support the Virtual Chassis configuration.

[See [Dashboard for EX Series Switches](#) and [Connecting and Configuring an EX Series Switch \(J-Web Procedure\)](#).]

Layer 2 Features

- **L2PT support (EX4300 multigigabit switches)**—Starting with Junos OS Release 19.2R1, you can configure Layer 2 protocol tunneling (L2PT) for the following protocols on EX4300 multigigabit switches (EX4300-48MP models): CDP, E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, STP (including RSTP and MSTP), UDLD, VSTP, and VTP.

[See [Layer 2 Protocol Tunneling](#).]

Multicast

- **Support for multicast traffic counters (EX4300, EX4300-MP, EX4300 Virtual Chassis)**—Starting with Junos OS Release 19.2R1, you can use firewall filters to count packets and check the bandwidth of multicast traffic received by a host from a particular source and group in a routing instance. To enable this feature, include the **multicast-statistics** statement at the **[edit system packet-forwarding-options]** hierarchy level. To check the packet count and bandwidth for each multicast route, use the **show multicast route extensive** command.

[See [multicast-statistics \(system-packet forwarding\)](#).]

- **IGMP snooping with private VLANs (EX4300 multigigabit switches)**—Starting in Junos OS Release 19.2R1, EX4300 multigigabit switches (EX4300-48MP models) support IGMP snooping with private VLANs (PVLANS). A PVLAN consists of secondary isolated and community VLANs configured within a primary VLAN. Without IGMP snooping support on the secondary VLANs, switches receive multicast streams on a primary VLAN and flood them to the secondary VLANs. This feature extends IGMP snooping on a primary VLAN to its secondary VLANs to further constrain multicast streams only to interested receivers on PVLANS. When you enable IGMP snooping on a primary VLAN, you implicitly enable it on all secondary VLANs, and the secondary VLANs learn the multicast group information on the primary VLAN.

NOTE: Ports in a secondary VLAN cannot be used as IGMP multicast router interfaces. Secondary VLANs can receive multicast data streams ingressing on promiscuous trunk ports or inter-switch links acting as multicast router interfaces.

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

Network Management and Monitoring

- **Support for displaying valid user input in the CLI for command options and configuration statements in custom YANG data models (EX Series)**—Starting in Junos OS Release 19.2R1, the CLI displays the set of possible values for a given command option or configuration statement in a custom YANG data model when you include the **action-expand** extension statement in the option or statement definition and reference a script that handles the logic. The **action-expand** statement must include the **script** child statement, which defines the Python action script that is invoked when a user requests context-sensitive help in the CLI for the value of that option or statement.

[See [Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules](#).]

Port Security

- **Stateless address autoconfiguration (SLAAC) snooping (EX2300, EX3400, EX4300, and Virtual Chassis)**—Starting in Junos OS Release 19.2R1, Junos OS supports SLAAC snooping on EX2300, EX2300 VC, EX3400, EX3400 VC, EX4300, and EX4300 VC. IPv6 clients using SLAAC for dynamic address assignment are validated against the SLAAC snooping binding table before being allowed access to the network.

[See [IPv6 Stateless Address Autoconfiguration \(SLAAC\) Snooping](#).]

- **Fallback PSK for Media Access Control Security (MACsec) (EX Series)**—Starting in Junos OS Release 19.2R1, fallback PSK for MACsec is supported on EX Series routers that support MACsec. The fallback PSK provides functionality to establish a secure session in the event that the primary PSKs on each end of a MACsec-secured link do not match.

[See [Configuring MACsec on EX, SRX and Fusion Devices](#).]

- **Support for 802.1X authentication on private VLANs (PVLANS) (EX4300-48MP switches and Virtual Chassis)**—Starting in Junos OS Release 19.2R1, you can enable 802.1X (dot1x) authentication for security purposes on access ports that are in a PVLAN.

PVLANS provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs. PVLANS are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts.

Authentication prevents unauthenticated devices and users from gaining access to your LAN. For 802.1X and MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a DHCP server.

On a switch that is configured with both 802.1X authentication and PVLANS, when a new device is attached to the PVLAN network, the device is authenticated and then is assigned to a secondary VLAN based on the PVLAN configuration or RADIUS profile. The device then obtains an IP address and receives access to the PVLAN network.

[See [Using 802.1X Authentication and Private VLANs Together on the Same Interface.](#)]

- **Media Access Control security with 256-bit cipher suite (EX4300)**—Starting in Junos OS Release 19.2R1, the GCM-AES-256 cipher suite for MACsec in static CAK mode is supported on the 2-port QSFP+/1-port QSFP28 uplink module for EX4300-48MP switches. The GCM-AES-256 cipher suite has a maximum key length of 256 bits and is also available with extended packet numbering (GCM-AES-XPB-256).

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

- **Support for MACsec PSK keychain (EX9253)**—Starting in Junos OS Release 19.2R1, EX9253 switches support MACsec PSK chains hitless rollover and Key Agreement Protocol Fail Open mode.

[See [Configuring MACsec on EX, SRX and Fusion Devices.](#)]

System Management

- **Support for transferring accounting statistics files and router configuration archives using HTTP URL (EX Series)**—Starting in Junos OS Release 19.2R1, you can transfer accounting statistics files and router configuration archives to remote servers by using an HTTP URL. In addition to SCP and FTP, the following HTTP URL will be supported under the **archive-sites** statement:

`http://username@host:url-path password password`

- To transfer accounting statistics files, configure **archive-sites** under **[edit accounting-options file <filename>]** hierarchy.
- To transfer router configuration archival, configure **archive-sites** under **edit system archival configuration** hierarchy.
- To view the statistics of transfer attempted, succeeded, and failed, use the **show accounting server statistics archival-transfer** command.
- To clear the statistics of transfer attempted, succeeded, and failed, use the **clear accounting server statistics archival-transfer** command.

[See [archive-sites](#), [Backing Up Configurations to an Archive Site](#), [show accounting server statistics archival-transfer](#), and [clear accounting server statistics archival-transfer](#)].

SEE ALSO

[What's Chnaged](#) | 41

[Known Behavior](#) | 44

[Open Issues](#) | 45

[Resolved Issues | 49](#)[Documentation Updates | 63](#)[Migration, Upgrade, and Downgrade Instructions | 63](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.2R3 | 41](#)
- [What's Changed in Release 19.2R2 | 41](#)
- [What's Changed in Release 19.2R1 | 43](#)

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

What's Changed in Release 19.2R3

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

See [traceoptions \(Services\)](#).

What's Changed in Release 19.2R2

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, 19.2R2, and later, EX Series switches support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

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

- **Logical Interface is created along with physical Interface by default (EX Series, MX Series, and QFX Series)**—In Junos OS Release 19.2R2 and later, logical interface is created on **ge**, **et**, **xe** interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces are created.

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

Layer 2 Feature

- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—In Junos OS Release 19.2R2, we have introduced the configuration statement **input-native-vlan-push** at the **[edit interfaces *interface-name*]** hierarchy level. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement **input-vlan-map** with a **push** operation is configured.

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

Multicast

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 19.2R2, EX4600, EX4650, and the QFX5000 line of switches provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **show multicast snooping route extensive** CLI command to see this count in the **Statistics: ... n packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). In earlier Junos OS releases, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** do not provide valid statistics for multicast traffic at Layer 2.

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

Routing Protocols

- **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 19.2R2, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.2R2, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

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

What's Changed in Release 19.2R1

Interfaces and Chassis

- **Deprecation of the [edit fabric protocols bgp] hierarchy level (EX Series)**—Starting in Junos OS Release 19.2R1, the **[edit fabric protocols bgp]** hierarchy level is deprecated.

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (EX Series)**—Starting in Junos OS Release 19.2R1, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier

releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.

- **Custom YANG RPC support for input parameters of type empty (EX Series)**—Starting in Junos OS Release 19.2R1, custom YANG RPCs support input parameters of type **empty** when executing the RPC's command in the Junos OS CLI, and the value passed to the action script is the parameter name. In earlier releases, input parameters of type **empty** are only supported when executing the RPC in a NETCONF or Junos XML protocol session, and the value passed to the action script is the string '**none**'.

[See [Creating Action Scripts for YANG RPCs on Devices Running Junos OS](#).]

VLAN Infrastructure

- **Specifying a descending VLAN ID range (EX9200 switches)**—In Junos OS releases prior to Junos OS Release 19.2R1, the system accepts a descending range—for example, 102-100, with the **vlan-id-range** configuration statement in the **[edit interfaces interface-name unit logical-unit-number]** hierarchy.

Starting with Junos OS Release 19.2R1, the system considers a descending range specified with **vlan-id-range** to be invalid and raises an error if you try to commit this configuration.

SEE ALSO

What's New 34
Known Behavior 44
Open Issues 45
Resolved Issues 49
Documentation Updates 63
Migration, Upgrade, and Downgrade Instructions 63

Known Limitations

IN THIS SECTION

- [EVPN | 45](#)
- [Platform and Infrastructure | 45](#)

Learn about known limitations in this release 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.

EVPN

- Commit blocks for VLAN-ID none with EVPN routing-instance and without routing-instance. [PR1287557](#)

Platform and Infrastructure

- On the EX4300 device, multicast traffic statistics do not get tracked for active s,g route streams even after all the 512 routes get timed out and programmed filter entries are cleared. [PR1419926](#)
- With 288000 MAC scale, the Routing Engine command output shows the learned scale entries after a delay of around 60 seconds. [PR1367538](#)
- Unable to ping peer IP over MACsec with AES-256 cipher suite. [PR1416499](#)
- Memory spike or leakage is seen after the image upgrades to Junos OS Release 19.2R1.8 in a mixed mode Virtual Chassis. [PR1464062](#)
- The following error message might appear: **Failed to complete DFE tuning**. This error message has no functional impact and can be ignored. [PR1473280](#)

SEE ALSO

What's New	 34
What's Chnaged	 41
Open Issues	 45
Resolved Issues	 49
Documentation Updates	 63
Migration, Upgrade, and Downgrade Instructions	 63

Open Issues

IN THIS SECTION

- [Authentication and Access Control](#) | 46
- [Infrastructure](#) | 46

- Interfaces and Chassis | 47
- Junos Fusion for Enterprise | 47
- Layer 2 Features | 47
- Network Management and Monitoring | 47
- Platform and Infrastructure | 47
- Routing Protocols | 48

Learn about the open issues in hardware and software in Junos OS Release 19.2R3 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

- The lldpd process might crash when the LLDP neighbor expires. [PR1408707](#)
- After reboot, the SSL certificates are not present. [PR1431086](#)

Infrastructure

- Some PIM groups are not able to send out native multicast traffic. [PR1209585](#)
- On the EX3400 and EX2300 devices during ZTP with configuration and image upgrade with FTP as file transfer, image upgrade is successful but sometimes VM generates core files. [PR1377721](#)
- Switch encapsulate protocol PDUs even if it is not configured for the Layer 2 PT tunneling. [PR1395638](#)
- The dot1x authentication might fail to be authenticated. [PR1408717](#)
- Error messages are seen at (pkt tx) ifd get failed 700 ,(brcm_port_learning_config:1375) Setting L2 learning unit:0, port_num:68, learn_flg 5, Disabling DHCP trapping on xe-0/0/40 dev:0, port:68 & routing_chip_output_packet:8001 (pkt tx) ifd get failed 700. [PR1422402](#)
- The FPC crashes with pfem generating core file might be seen if large-scale number of firewall filters are configured. [PR1434927](#)
- The following error message is observed: IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) failed continuously in AD with base configurations. [PR1485038](#)

Interfaces and Chassis

- The GRES and VSTP ports cost on the aggregated Ethernet interfaces might get changed, leading to topology change. [PR1174213](#)
- Traffic convergence is greater than 12 seconds for events like node reboot, ICCP and ICL flaps. [PR1371493](#)

Junos Fusion for Enterprise

- On the Junos fusion environment system, intermediate traffic drop is seen between AD and SD when sFlow is enabled on the ingress interface. [PR1450373](#)

Layer 2 Features

- GARPs are being sent from the switch once in every 10 minutes. [PR1192520](#)

Network Management and Monitoring

- The Packet Forwarding Engine process might crash in scenarios where there are large log messages. [PR1233050](#)

Platform and Infrastructure

- On the EX2300 and EX3400 devices, the upgrade might fail because of the shortage of space. [PR1464808](#)
- When 3000 ARPs are pushed to the EX Series switch with DAI enabled, the ARP inspection fails for valid hosts. [PR1165757](#)
- On the EX2300 device, the last reboot reason might display wrong values. [PR1331264](#)
- On the EX4650 device, the filter action to change VLAN does not work. [PR1362609](#)
- The EX4300 device might not update the **dhcp-security** binding upon renewal when the loopback filter and DHCP-security configuration are used together. [PR1376454](#)
- In the EX9208 device, a few xe interfaces go down with the following error message:
`f_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error.` [PR1377840](#)
- MACsec session might fail to reestablish after the interface flaps. [PR1378710](#)
- DCPFE does not come up in some instances of abrupt power-off or power-on. [PR1393554](#)
- The following error message is observed after multiple triggers: **JTASK_IO_CONNECT_FAILED.** [PR1408995](#)

- Traffic loss of approximately 26-32 seconds is observed after restarting the routing daemon on the EX9200 device with the MC-LAG configurations. [PR1409773](#)
- uRPF in strict mode does not work. [PR1417546](#)
- Committing the configurations that involve the **interface-range** defined over **wild-card range** such as **ge-*/*/*** are not supported. [PR1421446](#)
- IGMP transit query packets might not be flooded on VLAN. [PR1427542](#)
- On the EX9214 device, the following error message are observed after reboot and MACsec-enabled link flaps: **errorlib_set_error_log(): err_id(-1718026239)**. [PR1448368](#)
- In some cases, if we have an OSPF session on the IRB over LAG interface with 40-Gigabit Ethernet port as member, the session gets stuck in restart. [PR1498903](#)
- On the EX4600 device with VXLAN enabled, the ARP request might get dropped if the storm control is configured. [PR1515254](#)

Routing Protocols

- Packet loss is observed for the stream **bLock:irb_lacp_tr_ospf** while verifying traffic from access to core network for IPv4. [PR1520059](#)

SEE ALSO

[What's New | 34](#)

[What's Chnaged | 41](#)

[Known Behavior | 44](#)

[Resolved Issues | 49](#)

[Documentation Updates | 63](#)

[Migration, Upgrade, and Downgrade Instructions | 63](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3 | 49](#)
- [Resolved Issues: 19.2R2 | 51](#)
- [Resolved Issues: 19.2R1 | 59](#)

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

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

Resolved Issues: 19.2R3

Authentication and Access Control

- On the EX4600 and EX4300 switches, MAC entry is not present in the Ethernet-Switching table for the MAC-RADIUS client in a server fail scenario when tagged traffic is sent for two clients. [PR1462479](#)
- On the EX2300-48MP switch, client does not receive the captive-portal success page by downloading the ACL parameter, because the authentication failed. [PR1504818](#)
- The **DOT1XD_AUTH_SESSION_DELETED** event is not triggered with a single supplicant mode. [PR1512724](#)

EVPN

- The ESI of IRB interfaces does not update after autonomous-system number change if the interface is down. [PR1482790](#)
- The l2ald memory leakage might be observed in any EVPN scenario. [PR1498023](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)

Infrastructure

- On the EX2300 and EX3400 switches, kernel might generate core file when deactivating daemon. [PR1483644](#)

Junos Fusion Satellite Software

- On the EX4300 in Junos fusion scenario, temperature sensor alarm is observed. [PR1466324](#)

Layer 2 Ethernet Services

- Issues with DHCPv6 Relay processing confirm and reply packets are observed. [PR1496220](#)

MPLS

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

Platform and Infrastructure

- MAC learning under bridge-domain stops after MC-LAG interface flaps. [PR1488251](#)
- Packets get dropped when next hop is IRB over It interface. [PR1494594](#)
- On the EX4300 switches, the NSSU upgrade might fail due to storage issue on the `/var/tmp` directory. [PR1494963](#)
- IPv6 neighbor solicitation packets might be dropped in a transit device. [PR1493212](#)
- On the EX4300 device, traffic loss might be seen with framing errors or runs if MACsec is configured. [PR1502726](#)
- On the EX9200 device, IRB over VTEP unicast traffic might get dropped. [PR1436924](#)
- On the EX4300, EX3400, and EX2300 Virtual Chassis with NSB and xSTP enabled, continuous traffic loss might be observed while performing GRES. [PR1500783](#)
- On the EX4600 device, traffic loss might be seen with framing errors or runs if MACsec is configured. [PR1469663](#)
- On the EX4600 device, DSCP marking might not work as expected if the fixed classifiers are applied to interfaces. [PR1472771](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- DHCP binding might fail when the PVLAN is configured with a firewall to block or allow certain IPv4 packets. [PR1490689](#)
- On the EX4650 switch, traffic loss might be seen under MC-LAG scenario. [PR1494507](#)
- Outbound SSH connection flaps or leaks memory during push configuration to ephemeral database with high rate. [PR1497575](#)

- Traffic might get dropped if the aggregated Ethernet member interface is deleted or added, or a SFP of the aggregated Ethernet member interface is unplugged or plugged. [PR1497993](#)
- LLDP is not acquired when **native-vlan-id** and tagged VLAN-ID are the same on a port. [PR1504354](#)

Routing Protocols

- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

Resolved Issues: 19.2R2

Class of Service (CoS)

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

EVPN

- EVPN or MPLS IRB logical interfaces might not come up when local Layer 2 interfaces are down. [PR1436207](#)
- Configuring ESI on a single-homed 25-Gigabit Ethernet port might not work. [PR1438227](#)
- An ARP request or a Neighbor Solicitation message might be sent back to the local segment by the designated forwarder (DF) router. [PR1459830](#)
- The rpd might crash after the EVPN-related configuration is changed. [PR1467309](#)

Forwarding and Sampling

- Enable interface with input/output vlan-maps to be added to a routing instance configured with a VLAN ID or VLAN tags (instance type **virtual-switch/vpls**). [PR1433542](#)
- The l2ald process might experience memory leak on devices running Junos OS. [PR1455034](#)
- Type 1 ESI/AD route might not be generated locally on EVPN PE devices in all-active mode. [PR1464778](#)

General Routing

- Transit OSPF traffic over Q-in-Q tunneling might be dropped if a firewall filter is applied to loopback interface. [PR1355111](#)
- The l2ald process might crash and generate a core file on EX Virtual Chassis when converting a trunk port to dot1x access port with tagged traffic flowing. [PR1362587](#)
- The interface on a 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 exiting through another Virtual Chassis member if MLD snooping is enabled. [PR1423310](#)

- The MAC address pool might overlap between different switches. [PR1425123](#)
- Virtual Chassis split is seen after the network topology is changed. [PR1427075](#)
- The fxpc or the Packet Forwarding Engine process might crash on EX2300 and EX3400 switches. [PR1427391](#)
- Rebooting or halting a Virtual Chassis member might cause traffic on the RTG link to be down for about 30 seconds. [PR1427500](#)
- The l2ald process crashes after the dot1x configuration is deleted when dot1x and private VLAN (PVLAN) are enabled simultaneously on EX Series and QFX Series switches. [PR1428469](#)
- A client might fail to get an IP address from the DHCPv6 server. [PR1428769](#)
- The delay in transmission of BPDUs after GRES might result in loss of traffic on EX2300 and EX3400 Virtual Chassis. [PR1428935](#)
- The EX4300-48MP switch cannot learn MAC addresses through some access ports that are directly connected to a host when auto negotiation is used. [PR1430109](#)
- Disabling DAC QSFP ports might not work on MX204, MX10003, or EX9251. [PR1430921](#)
- Erroneous log messages and chassis environment output related to fan tray in EX4300MP-EX4300-48P Virtual Chassis. [PR1431263](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [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 transceiver into another port. [PR1435221](#)
- Traffic drop might be seen after MACsec session key rollover between primary and fallback for more than ten times. [PR1435277](#)
- The multichassis aggregated Ethernet (mc-ae) interface might get stuck in the **Waiting** state in a dual mc-ae scenario. [PR1435874](#)
- i40e NVM upgrade support for EX9200 platform. [PR1436223](#)
- The Gigabit Ethernet or multigigabit Ethernet SFP-T interface might not come up on EX2300, EX3400, and EX4300 switches. [PR1438078](#)
- Commit check error for VSTP on EX9200 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 might crash during the booting process in certain conditions. [PR1438597](#)
- The dot1x configuration might not work when captive port is also configured on the interface on a backup or non-master FPC. [PR1439200](#)
- LACP state might get stuck in **Attached** state after disabling peer active members. [PR1439268](#)

- On EX9200 DHCPv6 relay scenario, when DHCPv6 snooping and Neighbor Discovery Inspection (NDI) are enabled simultaneously on an IRB interface, the DHCPv6 relay binding does not come up. [PR1439844](#)
- The EX4600 and QFX5100 Virtual Chassis might not come up after you replace the Virtual Chassis port fiber connection with a DAC cable. [PR1440062](#)
- CPU might hang or an interface might be stuck down on a particular 1-Gigabit Ethernet port on MX Series, EX Series, and PTX Series devices. [PR1440526](#)
- MAC addresses learned on an RTG might not be aged out after a Virtual Chassis member is rebooted. [PR1440574](#)
- Clients in isolated VLANs might not get IP addresses after completing authentication when both **dhcp-security** and **dot1x** are configured. [PR1442078](#)
- EX3400 fan alarm (Fan X not spinning) appears and disappears repeatedly after the fantray (Absent) is removed. [PR1442134](#)
- The rpd might crash when BGP sends a notification message. [PR1442786](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- The port role might be incorrect in STP after the STP configuration is changed. [PR1443489](#)
- The **/var/host/motd does not exist** message is flooded every 5 seconds in chassisd logs. [PR1444903](#)
- On EX4300-MP, the following log message is generated continuously: **rp[6550]: task_connect: task AGENTD I/O.128.0.0.1+9500 addr 128.0.0.1+9500: Connection refused**. [PR1445618](#)
- Major alarm log messages are seen for temperature conditions at 56 degrees Celsius. [PR1446363](#)
- The traffic might be dropped when a firewall filter rule uses **then vlan** as the action in a Virtual Chassis scenario. [PR1446844](#)
- Phone home on EX Series devices because sysctl cannot read the device serial number. [PR1447291](#)
- EX3400 Virtual Chassis might go into hang state when a disk error occurs on EX3400. [PR1447853](#)
- Unicast ARP requests are not replied to with the no-arp-trap option. [PR1448071](#)
- On EX3400, IPv6 routes received through BGP do not show the correct age time. [PR1449305](#)
- Except one aggregated Ethernet member link, the other links do not send out sFlow sample packets for ingress traffic. [PR1449568](#)
- Tunneling encapsulated packets are dropped on the Layer 3 VPN MPLS PE-CE interface. [PR1451032](#)
- DHCP snooping static binding does not take effect after deleting and re-adding the entries. [PR1451688](#)
- The MAC pause frames will be incrementing in the Receive direction if half-duplex mode on 10-Mbps or 100-Mbps speed is configured. [PR1452209](#)
- The l2ald and eventd processes are hogging 100 percent CPU after the **clear ethernet-switching table** command is issued. [PR1452738](#)
- Configuration change in **VLAN all** option might affect the per-VLAN configuration. [PR1453505](#)

- Version compare in PHC might fail, making the PHC to download the same image. [PR1453535](#)
- You might not be able to apply a firewall filter in a particular VC/VCF member as TCAM space runs out. [PR1455177](#)
- Packet drop might be seen after removing and reinserting the SFP transceiver of the 40-Gigabit uplink module ports. [PR1456039](#)
- Link-up delay and traffic drop might be seen on mixed SP Layer 2 or Layer 3 and EP Layer 2 type configurations. [PR1456336](#)
- The syslog message **timeout connecting to peer database-replication** is generated when the **show version detail** command is issued. [PR1457284](#)
- Overtemperature SNMP trap messages are displayed after the software upgrade and update even though the temperatures are within the system thresholds. [PR1457456](#)
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used. [PR1458559](#)
- The fxpc process might crash due to several BGP IPv6 session flaps. [PR1459759](#)
- Storage space limitation leads to image installation failure when phone home is used on EX2300 and EX3400 platforms. [PR1460087](#)
- MAC addresses learned on an RTG might not be aged out after aging time. [PR1461293](#)
- RTG link faces nearly 20 seconds downtime when the backup node is rebooting. [PR1461554](#)
- Configuring any combination of VLANs and interfaces under VSTP/MSTP might cause VSTP/MSTP-related configurations that cannot be committed. [PR1463251](#)
- The Virtual Chassis function might be broken after upgrading on EX2300 and EX3400 switches. [PR1463635](#)
- On EX Series switches with ELS and on QFX Series switches, some command lines to disable MAC learning are not working. [PR1464797](#)
- The jdncpd might consume high CPU and no further subscribers can be brought up if more than 4000 DHCP-relay clients in the MAC-MOVE scenario. [PR1465277](#)
- The fxpc might crash after mastership election on EX2300 and EX3400 switches. [PR1465526](#)
- The broadcast and multicast traffic might be dropped over IRB or LAG interfaces in QFX Series and EX Series Virtual Chassis scenario. [PR1466423](#)
- The MAC move message might have an incorrect **from** interface when rapid MAC moves occurs. [PR1467459](#)
- In EX3400 FPCs get disconnected from Virtual Chassis briefly after the image upgrades or reboots. [PR1467707](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- FPC might be down when configuring **vxlan-routing**. [PR1468736](#)

- On the EX3400, traffic loss is observed between SFP-T connected interfaces because of auto negotiation failure. [PR1469750](#)
- EX3400 is advertising only 100 Mbps when configured with 100-Mbps speed with auto negotiation enabled. [PR1471931](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)
- The dhcpd process might crash in a Junos fusion environment. [PR1478375](#)
- TFTP installation from loader prompt might not succeed on the EX Series switches. [PR1480348](#)
- ARP request packets for unknown hosts might get dropped in the remote PE device in an EVPN-VXLAN scenario. [PR1480776](#)
- On the EX2300 Series, the SNMP traps are not generated when MAC addresses when the limit is reached. [PR1482709](#)

Infrastructure

- The operations on the console might not work if the **system ports console log-out-on-disconnect** configuration statement is configured. [PR1433224](#)
- On the EX4300 Series, the CLI configuration **on-disk-failure** is not supported. [PR1450093](#)
- Certain EX Series switches might panic and generate VM core files, leading to reboot. [PR1456668](#)
- Error messages related to soft reset of ports due to queue buffers being stuck could be seen on EX4600 and EX4300 Virtual Chassis. [PR1462106](#)
- Traffic drop is seen on an EX4300-48MP device that acts as a leaf node in a Layer 2 IP fabric EVPN-VXLAN environment. [PR1463318](#)
- Continuous dcpfe error messages and eventd process hog might be seen in an EX2300 Virtual Chassis scenario. [PR1474808](#)

Interfaces and Chassis

- On EX9200 devices, an unexpected **duplicate VLAN-ID** commit error might be seen. [PR1430966](#)
- The VRRP IPv6 state might flap between init and idle states after VLAN tagging is configured. [PR1445370](#)
- Traffic might be forwarded to wrong interfaces in an MC-LAG scenario. [PR1465077](#)
- Executing commit might hang because of a stuck dcd process. [PR1470622](#)

J-Web

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

Junos Fusion for Enterprise

- Reachability of the host connected to the satellite device might be affected in a Junos fusion for enterprise environment with EX9200 Series switches as aggregation devices. [PR1447873](#)

- Loop detection might not work on extended ports in a Junos fusion scenario. [PR1460209](#)

Layer 2 Features

- Ethernet ring protection switching (ERPS) nodes might not converge to the **Idle** state after failure recovery or reboot. [PR1431262](#)
- Physical layer and MAC/ARP learning might not work for copper base SFP-T transceivers on QFX5100, QFX5110, and EX4600. [PR1437577](#)
- The traffic leaving QFX5000 and EX4600 switches might not be properly load balanced over aggregated Ethernet interfaces. [PR1448488](#)
- The LLDP function might fail when a Juniper Networks device connects to a non-Juniper device. [PR1462171](#)
- An fxpc core file might be seen when committing the configuration all together. For example, after a reboot [PR1467763](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)

Layer 2 Ethernet Services

- The DHCP decline packets are not forwarded to the DHCP server when **forward-only** is set within DHCP reply. [PR1429456](#)
- The jdhcpd_era log files constantly consume 121 MB of space out of 170 MB, resulting into a full file system traffic impact. [PR1431201](#)
- DHCP request might get dropped in a DHCP relay scenario. [PR1435039](#)
- In EX9200 switches, DHCP relay is stripping the GIADDR field in messages toward the DHCP clients. [PR1443516](#)

Platform and Infrastructure

- LACP DDoS policer is incorrectly triggered by other protocol- traffic on all EX9200, T4000, and MX Series platforms. [PR1409626](#)
- On the EX4300-48MP running Junos OS Release 18.3R1.9, overtemperature SNMP trap is generated wrongly for line card (EX4300-48P) based on master Routing Engine (EX4300-48MP) temperature threshold value. [PR1419300](#)
- On the EX4300, the runt counter never increments. [PR1419724](#)
- SNMP (ifHighSpeed) value does not appear properly only for VCP interfaces; , it appears as zero. [PR1425167](#)
- Packet drops, replication failure, or ksyncd crashes might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- IPv6 traffic might be dropped when static /64 IPv6 routes are configured. [PR1427866](#)
- EX4300 does not drop FCS frames with CRC error on xe- interfaces. [PR1429865](#)

- Unicast ARP requests are not replied with the **no-arp-trap** option. [PR1429964](#)
- EX4300 without soft error recovery (parity check, correction and memscan) enabled. [PR1430079](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- An FPC/pfex crash might be observed due to DMA buffer leaking. [PR1436642](#)
- The **/var/db/scripts** directory might be deleted after the **request system zeroize** command is executed. [PR1436773](#)
- The laser TX might be enabled while the interface is disabled [PR1445626](#)
- The PoE might not work after the PoE firmware on EX4300 switches is upgraded. [PR1446915](#)
- The firewall filters might not be created due to TCAM issues. [PR1447012](#)
- NSSU causes a traffic loss after the backup-to-master transitions. [PR1448607](#)
- The Errors on certain MPCs are classified as major, which should be minor or non-fatal. [PR1449427](#)
- The REST API process becomes nonresponsive when a number of requests come at a high rate. [PR1449987](#)
- The IRB traffic might drop after a mastership switchover. [PR1453025](#)
- The traffic for some VLANs might not be forwarded when **vlan-id-list** is configured. [PR1456879](#)
- The OSPF neighbor might go down when mDNS or PTP traffic is received at a rate higher than 1400 pps. [PR1459210](#)
- ERP might not revert to the **Idle** state after reload or reboot of multiple switches. [PR1461434](#)
- Traffic loss might be observed longer than 20 seconds when performing NSSU on EX4300 Virtual Chassis. [PR1461983](#)
- IGMP reports are dropped with mixed enterprise/SP configuration styles on EX4300 switches. [PR1466075](#)
- The switch might not be able to learn MAC addresses with **dot1x** and **interface-mac-limit** configured. [PR1470424](#)
- On an EX4300, the input firewall filter attached to isolated or community VLANs is not matching dot1p bits on the VLAN header. [PR1478240](#)
- The Virtual Chassis VRRP peer drops packets to VRRP VIP after IRB is disabled. [PR1491348](#)

Routing Protocols

- Host-destined packets with the **filter log** action might not reach the Routing Engine if **log/syslog** is enabled. [PR1379718](#)
- BGP IPv4 or IPv6 convergence and RIB might delete and then install the time degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- The traffic with destination UDP port 520 (RIP) or 521 (RIPng) gets dropped on the QFX5000 and EX4600 switches. [PR1429543](#)

- The fxpc core file might be seen during the reboot of QFX5100 and EX4600 devices. [PR1432023](#)
- The **RPD_DYN_CFG_GET_PROF_NAME_FAILED: Get profile name for session XXX failed: -7** error message might be seen in syslog after restarting the routing process. [PR1439514](#)
- The bandwidth value of DDoS protection might cause packet loss after a device reboot. [PR1440847](#)
- Traffic might be dropped after the Q-in-Q-enabled interface is flapped or a change is made to the **vlan-id-list** configuration. [PR1441402](#)
- IPv6 connectivity between MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- Junos OS BFD sessions with authentication flap after a certain time. [PR1448649](#)
- Loopback address exported into other VRF instances might not work on ACX Series, EX Series, and QFX Series. [PR1449410](#)
- MPLS LDP might still use the stale MAC address of the neighbor even the LDP neighbor's MAC address changes. [PR1451217](#)
- The **other querier present interval** timer cannot be changed in the IGMP/MLD snooping scenario. [PR1461590](#)
- The MUX state in an LACP interface does not go to **collecting and distributing** and remains **attached** after enabling the aggregated Ethernet interface. [PR1484523](#)
- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

User Interface and Configuration

- The switch might be unable to commit baseline configuration after zeroization. [PR1426341](#)
- Problem with access to J-Web after update from Junos OS Release 18.2R2 to Junos OS Release 18.2R3. [PR1454150](#)
- The **umount: unmount of /.mount/var/val/chroot/packages/mnt/jweb-ex32-d2cf6f6b failed: Device busy** message is seen when Junos OS is upgraded with the **validate** option. [PR1478291](#)

Virtual Chassis

- The current MAC address might change when one of the multiple Layer 3 interfaces is being deleted. [PR1449206](#)

VPNs

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

Resolved Issues: 19.2R1

Authentication and Access Control

- Without configuring anything related to dot1x, the syslog **dot1xd[2192]: task_connect: task PNACAUTH./var/run/authd_control addr /var/run/authd_control: Connection refused** is generated repeatedly. [PR1406965](#)

EVPN

- The device might proxy the ARP probe packets in an EVPN environment. [PR1427109](#)
- ESI is configured on a single-homed 25G port might not work. [PR1438227](#)

General Routing

- On EX4650 switches, convergence delay between PE1 and P router link is more than the expected delay value. [PR1364244](#)
- OAM Ethernet **connectivity-fault-management** configured on aggregated Ethernet interfaces is not supported and no commit error is seen. [PR1367588](#)
- IPv6 router advertisement (RA) messages potentially increase internal kernel memory usage. [PR1369638](#)
- RIPV2 update packets might not be sent with IGMP snooping enabled. [PR1375332](#)
- Input rate PPS does not increase on EX2300-MP uplink ports when the packet is a pure L2 packet like non-etherII or non-EtherSnap. [PR1389908](#)
- EX3400VC - When an interface in a Virtual Chassis member switch that is not master, is flapped, IGMP query packets 224.0.0.1 are sent to all the ports of the members except the master FPC. [PR1393405](#)
- PTP over Ethernet traffic might be dropped when IGMP and PTP TC are configured together. [PR1395186](#)
- EX3400 might not learn 30,000 MAC addresses while sending MAC learning traffic. [PR1399575](#)
- MAC-limit with persistent MAC is not working after reboot. [PR1400507](#)
- After upgrading to Junos OS Release 18.1R3.3, **adt7470_set_pwm** output message is observed continuously. [PR1401709](#)
- The DHCP discover packets are forwarded out of an interface incorrectly when DHCP snooping is configured on that interface. [PR1403528](#)

- On EX4300-48MP devices, the packets drop when the traffic filter and the routing instance are configured. [PR1407424](#)
- The l2cpd might crash if the **vstp traceoptions** and **vstp vlan all** commands are configured. [PR1407469](#)
- MAC address movement might not happen in flexible Ethernet services mode when family inet/inet6 and vlan-bridge are configured on the same physical interface. [PR1408230](#)
- EX3400 PSU status is still taking "check" status even though PSU module has been removed. [PR1408675](#)
- On EX2300-24P switches, error message **dc-pfe: BCM_NH-,bcm_nh_resolve_get_nexthop(),346:Failed to find if family** is seen. [PR1410717](#)
- On EX Series devices, the PEM alarm for backup FPC remains on master FPC though the backup FPC is detached from Virtual Chassis. [PR1412429](#)
- On EX4300-48MP devices, the chassis status LED shows yellow instead of amber. [PR1413194](#)
- The chassisd output power budget is received continually per 5 seconds without any alarm after an upgrade to Junos OS Release 18.1R3. [PR1414267](#)
- VXLAN encapsulation next hop (VENH) does not get installed during BGP flap or when routing is restarted. [PR1415450](#)
- On EX3400 switches, the **show chassis environment** repeats **OK** and **Failed** at short intervals. [PR1417839](#)
- The EX3400 VC status might be unstable during the boot-up of the Virtual Chassis or after the Virtual Chassis port flaps. [PR1418490](#)
- Virtual Chassis might become unstable and FXPC crashes and generates a core file when there are a lot of configured filter entries. [PR1422132](#)
- On EX3400 auto-negotiation status shows incomplete on ge-0/2/0 using SFP-SX. [PR1423469](#)
- On EX4600 line of switches, MACsec might not connect when the interface disconnects while traffic is passing. [PR1423597](#)
- I2C read errors are seen when an SFP-T is inserted into a disabled state port configured with **set interface <*> disable** command. [PR1423858](#)
- Incorrect model information while polling through SNMP from Virtual Chassis. [PR1431135](#)

Infrastructure

- IfSpeed and IfHighSpeed erroneously reported as zero on EX2300. [PR1326902](#)
- Packet Forwarding Engine is flooded with messages **// pkt rx** on physical interface NULL unit 0. [PR1381151](#)

Interfaces and Chassis

- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. [PR1402606](#)
- EVPN aggregated Ethernet interface flaps followed by a commit. [PR1425339](#)

Junos Fusion for Enterprise

- PoE over LLDP negotiation is not supported on Junos fusion for enterprise setup. [PR1366106](#)
- New satellite device cannot be added to the Fusion scenario. [PR1374982](#)
- Cascade port might go down after SD reboot in Junos fusion for enterprise environment. [PR1382091](#)
- Cannot log in to SD cluster though it is recognized by AD properly. [PR1395570](#)
- The l2ald might crash when **clear ethernet-switching table persistent-learning** command is issued. [PR1409403](#)
- Extended ports in Junos fusion for enterprise do not adjust MTU when VoIP is enabled. [PR1411179](#)
- The traffic might silently drop and get discarded in Junos fusion for enterprise scenario with dual-AD. [PR1417139](#)

Layer 2 Ethernet Services

- The malfunction of core isolation feature in EVPN VXLAN scenarios might cause traffic to get silently dropped and discarded. [PR1417729](#)

Network Management and Monitoring

- Over temperature trap is not sent out even when there is a temperature-hot-alarm. [PR1412161](#)

Platform and Infrastructure

- Ping does not go through the device after WTR timer expires in Ethernet ring protection switching (ERPS) scenario. [PR1132770](#)
- EX4300 upgrade fails during validation of SLAX script. [PR1376750](#)
- Unicast DHCP request gets misforwarded to backup RTG link on EX4300 Virtual Chassis. [PR1388211](#)
- EX4300 OAM LFM might not work on extended-vlan-bridge interface with **native vlan** configured. [PR1399864](#)
- Traffic drop is seen on EX4300 when 10-Gigabit fiber port is using 1-Gigabit Ethernet SFP optics with auto-negotiation enabled. [PR1405168](#)
- On EX4300, when power supply (PEM) is removed, alarm is not generated. [PR1405262](#)
- The policer might not work when it is applied through the dynamic filter. [PR1410973](#)
- The traffic to the NLB server might not be forwarded if the NLB cluster works on multicast mode. [PR1411549](#)
- EX4300 QinQ - untagged UNI traffic egress as single-tagged on NNI interface. [PR1413700](#)
- Runt counter never incremented. [PR1419724](#)
- EX4300 does not send fragmentation needed message when MTU is exceeded with DF bit set. [PR1419893](#)
- The pfex process might crash and core files might be generated when SFP is reinserted. [PR1421257](#)

- Traffic might get silently dropped when one of logical interfaces on LAG is deactivated or deleted. [PR1422920](#)
- Auditd crashes when accounting RADIUS server is not reachable. [PR1424030](#)
- The native VLAN ID of packets might fail when leaving out. [PR1424174](#)
- Interface flapping scenario might lead to ECMP next-hop install failure on EX4300 switches. [PR1426760](#)
- VIP might not forward the traffic if VRRP is configured on an aggregated Ethernet interface. [PR1428124](#)
- EX4300 does not drop FCS frames on XE interfaces. [PR1429865](#)
- The ERPS failover does not work as expected on EX4300 device. [PR1432397](#)

Routing Protocols

- Host-destined packets with filter log action might reach the Routing Engine. [PR1379718](#)
- The rpd crashes on static route configuration for multicast source. [PR1408443](#)
- Host-generated ICMPv6 RA packets might be dropped on the backup member of Virtual Chassis if **igmp-snooping** is configured. [PR1413543](#)
- The EX Series switches might not install all IRB MAC addresses in the initialization. [PR1416025](#)
- After restarting multicast-snooping process, **igmp-snooping** might not work. [PR1420921](#)

Software Installation and Upgrade

- Configuration loss and traffic loss might be seen if backup Routing Engine is zeroized and is then switched over to master within a short time. [PR1389268](#)

Subscriber Access Management

- authd reuses address quickly before jdhcpd completely cleans up the old subscriber that gives the following error log **DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add x.x.x.x as it is already used by xxx.** [PR1402653](#)
- On EX4300 /var showing full /var/log/dfcd_enc file grows in size. [PR1425000](#)

SEE ALSO

[What's New | 34](#)

[What's Chnaged | 41](#)

[Known Behavior | 44](#)

[Open Issues | 45](#)

[Documentation Updates | 63](#)

[Migration, Upgrade, and Downgrade Instructions | 63](#)

Documentation Updates

IN THIS SECTION

- [Installation and Upgrade | 63](#)

This section lists the errata and changes in Junos OS Release 19.2R3 for the EX Series switches documentation.

Installation and Upgrade

- **Veriexec explained (EX Series)**—Verified Exec (also known as veriexec) is a file-signing and verification scheme that protects the Junos operating system (OS) against unauthorized software and activity that might compromise the integrity of your device. Originally developed for the NetBSD OS, veriexec was adapted for Junos OS and enabled by default from Junos OS Release 7.5 onward.
[See [Veriexec Overview](#).]

SEE ALSO

What's New 34
What's Chnaged 41
Known Behavior 44
Open Issues 45
Resolved Issues 49
Migration, Upgrade, and Downgrade Instructions 63

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 64](#)

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	 	34
What's Chnaged	 	41
Known Behavior	 	44
Open Issues	 	45
Resolved Issues	 	49
Documentation Updates	 	63

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

- What's New | 65
- What's Changed | 67
- Known Limitations | 68
- Open Issues | 68
- Resolved Issues | 69
- Documentation Updates | 71
- Migration, Upgrade, and Downgrade Instructions | 71

These release notes accompany Junos OS Release 19.2R3 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

IN THIS SECTION

- What's New in Release 19.2R3 | 66
- What's New in Release 19.2R2 | 66
- What's New in Release 19.2R1 | 66

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

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

What's New in Release 19.2R3

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 19.2R3.

What's New in Release 19.2R2

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 19.2R2.

What's New in Release 19.2R1

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 19.2R1.

SEE ALSO

What's Changed	 67
Known Limitations	 68
Open Issues	 68
Resolved Issues	 69
Documentation Updates	 71
Migration, Upgrade, and Downgrade Instructions	 71

What's Changed

IN THIS SECTION

- [What's Changed in 19.2R3 | 67](#)
- [What's Changed in 19.2R2 | 67](#)
- [What's Changed in 19.2R1 | 67](#)

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

What's Changed in 19.2R3

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 19.2R3 for Junos fusion for enterprise.

What's Changed in 19.2R2

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 19.2R2 for Junos fusion for enterprise.

What's Changed in 19.2R1

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 19.2R1 for Junos fusion for enterprise.

SEE ALSO

What's New 65
Known Limitations 68
Open Issues 68
Resolved Issues 69
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 19.2R3 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	 65
What's Changed	 67
Open Issues	 68
Resolved Issues	 69
Documentation Updates	 71
Migration, Upgrade, and Downgrade Instructions	 71

Open Issues

IN THIS SECTION

- [Junos fusion for enterprise](#) | 68

This section lists the known issues in hardware and software in Junos OS Release 19.2R3 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

- On EX4300 when 10G fiber port is using 1G Ethernet SFP optics, auto-negotiation is enabled by default. To bring up the satellite device, BCM recommends to disable the auto-negotiation for PHY84756 ports. [PR1420343](#)

- In Junos fusion for enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, loop-detect feature does not work for ports 0-23, since the loop detect filter is not properly applied.
[PR1426757](#)
- In a Junos fusion system, intermediate traffic drop might be seen between the aggregation and satellite device when Sflow is enabled on the ingress interface. When Sflow is enabled, the original packet is corrupted for those packets which hit the Sflow filter. This is because the packets egressing the aggregation device are short 4 bytes of FCS and 2 bytes of data. Normal data packets are 128 bytes (4 bytes for FCS, 14 bytes for Ethernet header, 20 bytes for IP header and 90 bytes for data). The corrupted packets are 122 bytes (14 bytes for Ethernet header, 20 bytes for IP header, and 88 bytes for data).
[PR1450373](#)

SEE ALSO

What's New 65
What's Changed 67
Known Limitations 68
Resolved Issues 69
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3 | 70](#)
- [Resolved Issues: 19.2R2 | 70](#)
- [Resolved Issues: 19.2R1 | 70](#)

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

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

Resolved Issues: 19.2R3

There are no resolved issues in Junos OS Release 19.2R3 for Junos fusion for enterprise.

Resolved Issues: 19.2R2

- Reachability issue of the host connected to the satellite device might be affected in a Junos Fusion Enterprise environment with EX9200 series devices as aggregation devices. [PR1447873](#)
- Loop detection might not work on extended ports in Junos fusion scenarios. [PR1460209](#)
- The dpd process might be generate a core file on satellite devices in a Junos fusion for enterprise setup. [PR1460607](#)
- In a Junos fusion service provider setup the EX4300 acting as a satellite device is triggering the temperature sensor alarm on multiple satellite device modules connected to same aggregation device. [PR1466324](#)

Resolved Issues: 19.2R1

- PoE over LLDP negotiation is not supported on Junos Fusion Enterprise setup. [PR1366106](#)
- New satellite device cannot be added to the Fusion scenario. [PR1374982](#)
- Cascade port might go down after SD reboot in Junos Fusion Enterprise environment. [PR1382091](#)
- Cannot login to SD cluster though it is recognized by AD properly. [PR1395570](#)
- The l2ald might crash when **clear ethernet-switching table persistent-learning** command is issued. [PR1409403](#)
- Extended ports in JFE do not adjust MTU when VoIP is enabled. [PR1411179](#)
- The traffic might silently drop and get discarded in Junos Fusion Enterprise scenario with dual-AD. [PR1417139](#)

SEE ALSO

[What's New | 65](#)

[What's Changed | 67](#)

[Known Limitations | 68](#)

[Open Issues | 68](#)

[Documentation Updates | 71](#)

[Migration, Upgrade, and Downgrade Instructions | 71](#)

Documentation Updates

There are no errata or changes in Junos OS Release 19.2R3 for documentation for Junos fusion for enterprise.

SEE ALSO

[What's New | 65](#)[What's Changed | 67](#)[Known Limitations | 68](#)[Open Issues | 68](#)[Resolved Issues | 69](#)[Migration, Upgrade, and Downgrade Instructions | 71](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 71](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 73](#)
- [Preparing the Switch for Satellite Device Conversion | 74](#)
- [Converting a Satellite Device to a Standalone Switch | 75](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 75](#)
- [Downgrading from Junos OS | 76](#)

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

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the **junos-install** package. Use other packages (such as the **jbundle** package) only when so instructed by a Juniper Networks support

representative. For information about the contents of the **junos-install** package and details of the installation process, see the [Installation and Upgrade Guide](#).

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

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

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

To download and install Junos OS:

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

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

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

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

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

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

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

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos fusion for enterprise. See [Configuring or Expanding a Junos Fusion 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 for enterprise is first supported in Junos OS Release 16.1, although you can downgrade a standalone EX9200 switch to earlier Junos OS releases.

NOTE: You cannot downgrade more than three releases.

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

To downgrade a Junos fusion for enterprise from Junos OS Release 19.2R3, 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 | 65](#)

[What's Changed | 67](#)

[Known Limitations | 68](#)

[Open Issues | 68](#)

[Resolved Issues | 69](#)

[Documentation Updates | 71](#)

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- [What's New | 77](#)
- [What's Changed | 78](#)
- [Known Limitations | 79](#)
- [Open Issues | 79](#)
- [Resolved Issues | 80](#)
- [Documentation Updates | 82](#)
- [Migration, Upgrade, and Downgrade Instructions | 82](#)

These release notes accompany Junos OS Release 19.2R3 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

IN THIS SECTION

- [What's New in Release 19.2R3 | 78](#)
- [What's New in Release 19.2R2 | 78](#)
- [What's New in Release 19.2R1 | 78](#)

Learn about new features introduced in the main and maintenance releases for Junos Fusion Provider Edge.

What's New in Release 19.2R3

- There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.2R3.

What's New in Release 19.2R2

- There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 19.2R2.

What's New in Release 19.2R1

Spanning-Tree Protocols

- **Support for Multiple Spanning Tree Protocol (MSTP) (Junos Provider Edge)**—Starting with Junos OS Release 19.2R1, you can configure MSTP on MX480 devices. MSTP scales better than other types of spanning-tree protocols and enables load balancing.

[See [Configuring MSTP Protocol](#).]

SEE ALSO

[What's Changed | 78](#)

[Known Limitations | 79](#)

[Open Issues | 79](#)

[Resolved Issues | 80](#)

[Documentation Updates | 82](#)

[Migration, Upgrade, and Downgrade Instructions | 82](#)

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 19.2R3 for Junos Fusion Provider Edge.

SEE ALSO

[What's New | 77](#)

Known Limitations	79
Open Issues	79
Resolved Issues	80
Documentation Updates	82
Migration, Upgrade, and Downgrade Instructions	82

Known Limitations

There are no known behaviors, system maximums, or limitations in hardware and software in Junos OS Release 19.2R3 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	77
What's Changed	78
Open Issues	79
Resolved Issues	80
Documentation Updates	82
Migration, Upgrade, and Downgrade Instructions	82

Open Issues

IN THIS SECTION

- [Junos Fusion Provider Edge](#) | [80](#)

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 Provider Edge

- In a Junos fusion system, intermediate traffic drop might be seen between the aggregation and satellite device when Sflow is enabled on the ingress interface. When Sflow is enabled, the original packet is corrupted for those packets which hit the Sflow filter. This is because the packets egressing the aggregation device are short 4 bytes of FCS and 2 bytes of data. Normal data packets are 128 bytes (4 bytes for FCS, 14 bytes for Ethernet header, 20 bytes for IP header and 90 bytes for data). The corrupted packets are 122 bytes (14 bytes for Ethernet header, 20 bytes for IP header, and 88 bytes for data). [PR1450373](#)

SEE ALSO

What's New	 77
What's Changed	 78
Known Limitations	 79
Resolved Issues	 80
Documentation Updates	 82
Migration, Upgrade, and Downgrade Instructions	 82

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3](#) | [81](#)
- [Resolved Issues: 19.2R2](#) | [81](#)
- [Resolved Issues: 19.2R1](#) | [81](#)

Learn which issues were resolved in Junos OS main and maintenance releases for Junos Fusion Enterprise.

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

Resolved Issues: 19.2R3

Junos Fusion Provider Edge

- The extended ports on cluster MX480 always show half duplex from the aggregate device. [PR1490101](#)

Junos Fusion Satellite Software

- In Junos Fusion scenario with EX4300 acting as SD, the temperature sensor alarms and logs might be seen as a result of the incorrect temperature reading of about 2-3 degrees lower than the expected. There is no CLI command to clear the alarm and logs. [PR1466324](#)

Resolved Issues: 19.2R2

Junos Fusion Satellite Software

- The dpd crash might be observed on satellite devices in Junos Fusion Enterprise. [PR1460607](#)
- Temperature sensor alarm is seen on EX4300 in Junos Fusion scenario. [PR1466324](#)

Resolved Issues: 19.2R1

Junos Fusion 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 to it. [PR1441869](#)
- If a default shaper is applied to a cascade interface of an aggregation device (AD), the displayed value of "Guaranteed rate" is greater than the value of "Shaping rate" in the output of the **show class-of-service scheduler-hierarchy interface** command. [PR1415502](#)
- Auto-negotiation is not disabled in the hardware after the no-auto-negotiation option is set using the CLI. [PR1411852](#)

Junos Fusion Satellite Software

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

SEE ALSO

[What's New | 77](#)

[What's Changed | 78](#)

[Known Limitations | 79](#)

[Open Issues | 79](#)

[Documentation Updates | 82](#)

[Migration, Upgrade, and Downgrade Instructions | 82](#)

Documentation Updates

There are no errata or changes in Junos OS Release 19.2R3 documentation for Junos Fusion Provider Edge.

SEE ALSO

[What's New | 77](#)

[What's Changed | 78](#)

[Known Limitations | 79](#)

[Open Issues | 79](#)

[Resolved Issues | 80](#)

[Migration, Upgrade, and Downgrade Instructions | 82](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading an Aggregation Device | 83](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 85](#)
- [Preparing the Switch for Satellite Device Conversion | 86](#)
- [Converting a Satellite Device to a Standalone Device | 87](#)
- [Upgrading an Aggregation Device | 90](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 90](#)
- [Downgrading from Junos OS Release 19.2 | 90](#)

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for Junos Fusion 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 **bundle** 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 19.2R3 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/jinstall64-19.2R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot  
source/jinstall-19.2R3.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/jinstall64-19.2R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-19.2R3.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 19.2R3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

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

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

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

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

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

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

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

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

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

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See [Configuring Junos Fusion 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 19.2R1, 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 19.2

To downgrade from Release 19.2 to another supported release, follow the procedure for upgrading, but replace the 19.2 **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 77
What's Changed 78
Known Limitations 79
Open Issues 79
Resolved Issues 80
Documentation Updates 82

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

IN THIS SECTION

- [What's New | 92](#)
- [What's Changed | 113](#)
- [Known Limitations | 123](#)
- [Open Issues | 127](#)
- [Resolved Issues | 139](#)
- [Documentation Updates | 191](#)
- [Migration, Upgrade, and Downgrade Instructions | 192](#)

These release notes accompany Junos OS Release 19.2R3 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [What's New in 19.2R3 | 92](#)
- [What's New in 19.2R2 | 93](#)
- [What's New in 19.2R1-S4 | 93](#)
- [What's New in 19.2R1-S1 | 94](#)
- [What's New in 19.2R1 | 95](#)

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

What's New in 19.2R3

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 19.2R3.

What's New in 19.2R2

Junos OS XML, API, and Scripting

- **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 19.2R2 supports 64-bit architecture. Junos OS operating scripts (op scripts) and on-box JET applications can now 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 \(Python\)](#).]

Network Management and Monitoring

- **Implement new MIBs using telemetry-based model (MX Series)**—Starting in Junos OS Release 19.2R2, new MIBs `mplsMldpInterfaceStatsEntry` and `mplsMldpFecUpstreamSessTable` are introduced. The Routing Engine uses a telemetry-based approach to collect statistics to provide MIB data for these MIBs. A new statement, **sensor-based-stats** at the `[edit protocols ldp traffic-statistics]` hierarchy level, enables telemetry-based collection. You must configure this statement to enable MIB data collection for `mplsMldpInterfaceStatsEntry` and `mplsMldpFecUpstreamSessTable`.

Routing Protocols

- **ECMP nexthop update rate throttling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R2, you can choose to defer multipath computation for all families during a BGP peering churn. In very large-scale network deployments during BGP peering churn there is a temporary spike in multipath computation, which takes a toll on the Packet Forwarding Engine resources. This feature allows you to pause the multipath computation and to resume after the peering churn settles down. Note that if there is no BGP peering churn, then multipath computation is not paused.

To enable the pause option for BGP multipath computation during BGP peering churn, include the **pause computation** statement at the `[edit protocols BGP multipath]` hierarchy level.

Subscriber Management and Services

- **CoA messages support Session-Timeout attribute (MX Series)**—Starting in Junos OS Release 19.2R2, you can apply a session timeout for subscriber sessions with a RADIUS CoA message that includes the Session-Timeout attribute (27). This capability is useful, for example, when subscribers purchase Internet access for a specific period of time and must log out when the session expires. In earlier releases, the router does not recognize the attribute if it is included in a CoA message.

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

What's New in 19.2R1-S4

Interfaces and Chassis

- **Support for 1-Gbps speed on QFX-60S line card on PTX10008 and PTX10016 Routers**—Starting in Junos OS Release 19.2R1-S4, QFX10000-60S-6Q line card supports 1-Gbps speed on its ports (0 through 59). The QFX10000-60S-6Q line card contains 60 SFP+ ports that support 10-Gbps, two dual-speed QSFP28 ports that support either 40-Gbps or 100-Gbps, and four QSFP+ ports that support 40-Gbps. You can individually configure ports 0 to 59 for 10-Gbps or 1-Gbps port speed. Use the **set chassis fpc fpc-slot-number pic pic-number port port-number speed 1G** command to change the mode of a port from 10-Gbps to 1-Gbps. The transceivers supported for 1-Gbps are QFX-SFP-1GE-LX, QFX-SFP-1GE-SX, and QFX-SFP-1GE-T.

[See [QFX10000 Line Cards](#) for details on the combination of modes supported on the ports.]

Services Applications

- **Support for Two-Way Active Measurement Protocol (TWAMP) and hardware timestamping of RPM probe messages (MX10000 and PTX10000 routers)**—Starting in Release 19.2R1-S4, Junos OS supports TWAMP and hardware timestamping of RPM probe messages on the MX10008, MX10016, PTX10008 and PTX10016 routers. You can use TWAMP to measure IP performance between two devices in a network. By enabling hardware timestamping of RPM you can account for the latency in the communication of probe messages and also generate more accurate timers in the Packet Forwarding Engine.

[See [Understanding Two-Way Active Measurement Protocol on Routers](#) and [Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers](#).]

What's New in 19.2R1-S1

MPLS

- **Distributed CSPF for segment routing LSPs (MX Series)**—Starting in Junos OS Release 19.2R1-S1, 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 19.2R1-S1, 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 19.2R1-S1, 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](#).]

Routing Protocols

- **Decouple RSVP for IGP-TE (MX Series, PTX Series, ACX Series, QFX Series, SRX Series, and EX Series)**—Starting in Junos OS Release 19.2R1-S1, 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.

What's New in 19.2R1

Hardware

- **New fixed-configuration Modular Port Concentrator (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.2R1, the MPC10E-10C-MRATE is a new Modular Port Concentrator (MPC) that is supported on the MX240, MX480, and MX960 routers.

The MPC10E-10C-MRATE features the following:

- Line-rate throughput of up to 1.0 Tbps when installed with an enhanced midplane and 800 Gbps when installed with a standard midplane.
- Eight QSFP28 ports—Port numbers 0/0 through 0/3 and 1/0 through 1/3. The ports can be configured as 10-Gbps, 40-Gbps, or 100-Gbps Ethernet ports.
- Two QSFP56-DD ports—Port numbers 0/4 and 1/4. The ports can be configured as 10-Gbps, 40-Gbps, 100-Gbps Ethernet ports.

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

- **MX10016 Universal Routing Platform**—The MX10016 router provides 10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet modular solutions that support up to 2.4 Tbps per slot. The MX10016 router provides redundancy and resiliency. All major hardware components including the power system, the cooling system, the control board and the switch fabrics are fully redundant. MX10016 enables cloud and data center operators to transition from 10-Gigabit Ethernet and 40-Gigabit Ethernet networks to 100-Gigabit Ethernet high-performance networks. The 21 rack unit (21 U) modular chassis can provide 38.4 Tbps of throughput. The MX10016 router has 16 slots for the line cards that can support a maximum of 1536 10-Gigabit Ethernet ports, 384 40-Gigabit Ethernet ports, or 384 100-Gigabit Ethernet ports.

You can deploy the MX10016 router in an IP edge network using an MX10K-LC2101 line card (ordering model number is JNP10K-LC2101).

[See [MX10016 Hardware Guide](#).]

- **Advanced Cooling and Power Components (MX10008 Routers)**—Starting in Junos OS Release 19.2R1, MX10008 routers offer 5.5 KW power supplies, new high performance fan tray, and compatible fan tray controller. The JNP10K-PWR-AC2 power supply supports AC, high-voltage alternating current (HVAC), DC, or high-voltage direct current (HVDC). The JNP10K-PWR-DC2 provides a 5.5 KW upgrade for DC users. The JNP10008-FAN2 offers increased air flow through the chassis. The JNP10008-FAN2 offers 1793 cubic feet per minute (CFM) per fan tray. The new fan tray controller, JNP10008-FTC2 supports the new fan tray.

[See [MX10008 Hardware Guide](#).]

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Option to enable and disable SCP per user level independent of SSH (MX Series)**—Starting in Junos OS 19.2R1, you can enable and disable SCP for a certain login class user independent of SSH. By default, SCP is not allowed for users added to the system defined classes read-only, operator and unauthorized and is only allowed to the system defined class super-user. SCP is allowed for any login class user belonging to a user defined class. You can deny SCP request for a user assigned to a user defined class by using the **no-scp-server** configuration statement. Prior to 19.2R1, SCP was enabled and disabled when SSH was enabled and disabled.

To disable SCP for a certain login class, use set **no-scp-server** at the **[edit system login class <class_name>]** hierarchy level.

[See [no-scp-server](#).]

- **Option to enable and disable SFTP per user level (MX Series)**—Starting in Junos OS 19.2R1, you can enable and disable SFTP for a certain login class user. By default, SFTP is not allowed for users added to the system defined classes read-only, operator and unauthorized and is only allowed to the system defined class super-user if SFTP is enabled globally. For a user assigned to a user defined class, by default SFTP requests are allowed if **set system services ssh sftp-server** is configured. You can now deny SFTP requests for a user assigned to a user defined class by using the **no-sftp-server** configuration statement.

To disable SFTP for a certain login class, use set **no-sftp-server** at the **[edit system login class <class_name>]** hierarchy level.

[See [no-sftp-server](#).]

EVPN

- **Support for BFD, BGP, IS-IS, and OSPF on IRB interfaces in EVPN-MPLS networks (MX Series and vMX)**—Starting with Junos OS Release 19.2R1, you can configure Bidirectional Forwarding Detection (BFD), BGP, IS-IS, and OSPF routing protocols on the IRB interface in an EVPN-MPLS network to route and forward EVPN traffic. This feature supports single-homed, single-active, and all-active multihomed networks.

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

- **EVPN support of VLAN ID ranges and lists in service provider style interface configurations (MX Series routers, and vMX virtual routers)**—Starting in Junos OS Release 19.2R1, EX9200 switches, ACX5448 and MX Series routers, and vMX virtual routers support the use of VLAN ID ranges and lists in a service provider style interface configuration, which must be referenced in an EVPN routing instance. This configuration is supported with the following EVPN environments, services, and features:
 - Environments:
 - EVPN with VXLAN encapsulation
 - EVPN with MPLS encapsulation

- VLAN bundle service:
 - E-LAN
 - E-Tree
 - E-Line
- Feature:
 - EVPN multihoming:
 - All-active
 - Single-active
 - Singlehoming

[See [VLAN ID Ranges and Lists in an EVPN Environment](#).]

- **Connectivity fault management support in EVPN-VPWS (MX Series)**—Starting with Junos OS Release 19.2R1, you can configure Up maintenance association end points (MEPs) and maintenance association intermediate point (MIPs) on attachment circuits in support of connectivity fault management (CFM) in EVPN-VPWS networks. With the MEPs, you can monitor connectivity between two points on the EVPN-VPWS network. Junos OS supports the continuity check messages (CCM), loopback and link trace messages (LTMs) as defined in IEEE 802.1AG CFM, and delay measurements (DM) and synthetic loss measurements (SLMs) as defined in Y.1731 on a single-active homing network.

[See [Connectivity Fault Management Support for EVPN and Layer 2 VPN Overview](#).]

- **Support for control word in EVPN-VPWS (MX Series and vMX)** —Starting with Junos OS Release 19.2R1, Junos OS supports the insertion of a control word between the label stack and the MPLS payload in a network with EVPN-VPWS service. This feature prevents a transit device from delivering out-of-order packets as a result of the device's load-balancing hashing algorithm. When you enable the control word feature on a PE device, the PE device advertises support for a control word. If all the PE devices in an EVI on the EVPN-VPWS serviced network support control word, then the PE device inserts a control word between the label stack and the L2 header in the packet thus preventing the packet from being misidentified by transit devices.

[See [Control Word for EVPN-VPWS](#).]

Forwarding and Sampling

- **Support for local preference when selecting forwarding next-hops for ECMP traffic (MX Series)**—Starting in Junos OS Release 19.2R1, you can have equal cost multi-path (ECMP) traffic flows prefer local forwarding next-hops over remote ones. This feature supports BGP prefixes that are directly reachable with IPv4 MPLS ECMP next-hops. Use **ecmp-local-bias** to direct ECMP traffic towards local links, for example, to ensure that the overall load on the fabric is reduced. [See [ecmp-local-bias](#) for usage details.]

High Availability (HA) and Resiliency

- **ISSU support for MX2008 (MX Series)**—Starting in Junos OS Release 19.2R1, MX2008 routers support ISSU.

[See [Understanding In-Service Software Upgrade \(ISSU\)](#)]

Interfaces and Chassis

- **Support for local preference when selecting forwarding next-hops for load balancing (MX Series)**—Starting in Junos OS Release 19.2R1, you can have traffic flows across aggregated Ethernet or logical-tunnel interfaces prefer local forwarding next-hops over remote ones, for example to ensure that the overall load on the fabric is reduced. [See [local-bias](#) for usage details.]
- **Support to collect and display PRBS statistics (MX10003 and MX204)**—Starting in Junos OS Release 19.2R1, on MX10003 and MX204 routers, you can check the physical link connectivity by issuing the **test interfaces ifd-name prbs-test-start pattern-type type direction (0|1) flip (0|1)** that starts collecting the PRBS statistics.

The output of the **show interfaces interface-name prbs-stats** command displays the PRBS statistics while the test is in progress. These statistics are cleared after the test is complete or if it is stopped. You can stop collecting the statistics by issuing the **test interfaces ifd-name prbs-test-stop direction (0|1)** command.

NOTE: While running PRBS statistics, the link will be down.

[See [prbs-test-start](#), [prbs-test-stop](#), [show interfaces prbs-stats](#), [Collecting Pseudo Random Bit Sequence \(PRBS\) Statistics](#).]

- **Domain Name System (DNS) is VRF aware (MX Series)**—Starting in Junos OS Release 19.2R1, when the **management-instance** statement is configured at the **[edit system]** hierarchy level, you can use the non-default management routing instance **mgmt_junos** as the routing instance through which the DNS name server is reachable. To specify the routing instance **mgmt_junos**, configure our new configuration statement **routing-instance mgmt_junos**, at the **[edit system name-server server-ip]** hierarchy level.

[See [Management Interface in a Nondefault Instance](#), [Configuring a DNS Name Server for Resolving a Hostname into Addresses](#), [name-server](#), and [show host](#).]

- **SCBE3-MX interoperates with MPC10E-10C (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.2R1, the Enhanced Switch Control Board SCBE3-MX (model number: SCBE3-MX-S) supports fabric management on the MPC10E-10C line card on the MX240, MX480, and MX960 routers. The

SCBE3-MX-S supports a pluggable Routing Engine and provides a control plane and data plane interconnect to each line card slot. The MPC10E-10C supports a bandwidth of up to 1 Tbps (800 Gbps with four planes and 1 Tbps with 5 or 6 planes). With MPC10E 15C line card, in a non-redundant configuration the SCBE3-MX provides fabric bandwidth of up to 1 Tbps per slot with four fabric planes and 1.5 Tbps per slot when all six fabric planes are used. Starting in this release, the MPC10E line cards support the standard midplane, which supports a bandwidth up to 800 Gbps per slot. Support for the enhanced midplane, which provides a bandwidth of 1.5 Tbps with MPC10E-15C and 1 Tbps with MPC10E-10C, is already available.

[See [SCBE3-MX Description](#) and [MPC10E-15C-MRATE](#)]

- **Support for QSFP-100GE-DWDM2 transceiver (MX204, MX10003, MX10008, and MX10016)**—Starting in Junos OS Release 19.2R1, the MX204, MX10003, MX10008, and MX10016 routers support the QSFP-100GE-DWDM2 transceiver. The 100-Gbps bidirectional transceiver has a dual transmitter/receiver that enables it to transmit and receive data through a single optical fiber. You can perform the following actions when this transceiver is installed:
 - View the diagnostics data, warnings, and alarms for interfaces. [See [show interfaces diagnostics optics](#).]
 - Clear the bit error rate (BER) counters. [See [clear interfaces statistics](#).]
 - Obtain the transport, performance monitoring, and threshold crossing alert (TCA) information for interfaces. [See [show interfaces transport pm](#).]
 - Clear the optics information from transport performance monitoring data. [See [clear interfaces transport pm](#).]
 - Enable or disable TCAs. [See [tca](#).]
 - Enable or disable loopback mode. [See [optics-options](#).]
- **MPC10 distributed LACP support in PPM AFT (MX Series)**—Starting in Junos OS Release 19.2R1, the MPC10E-15C-MRATE and MPC10E-10C-MRATE MPCs support distributed LACP in Periodic Packet Manager (ppman) Advanced Forwarding Toolkit (AFT).
- **Support for Routing Engine hard disk smart check (MX240, MX480, MX204, MX960, MX10008, MX2008, MX2020, MX10016, MX10000, MX2010, MX10002, and MX10003)**—Starting in Junos OS Release 19.2R1, you can configure the device to perform certain health checks on the Routing Engine solid-state drive (SSD) and log a health event or raise an alarm in case a predefined health attribute threshold is breached. You can use the **set chassis routing-engine disk smart-check** command to instruct the system to raise an alarm when an SSD health attribute threshold is breached. You can view the alarm by using the command **show chassis alarms**.

[See [smart-check](#)]

Junos OS XML API and Scripting

- **Automation script library additions and upgrades (MX Series)**—Starting in Junos OS Release 19.2R1, devices running Junos OS that support the Python extensions package include new and upgraded Python modules. Python automation scripts can leverage new on-box Python modules, including the **requests**, **chardet**, and **urllib3** modules, as well as upgraded versions of the **idna**, **ipaddress**, and **six** modules. The Requests library provides additional methods for supporting initial deployments as well as for performing routine monitoring and configuration changes on devices running Junos OS.

[See [Overview of Python Modules Available on Devices Running Junos OS](#) and [Using the Requests Library for Python on Devices Running Junos OS](#).]

Junos Telemetry Interface

- **Inline active flow monitoring support using JTI (MPC10E-15C-MRATE line cards)**—Starting in Junos OS Release 19.2R1, Junos Telemetry Interface (JTI) supports streaming inline active flow monitoring service-related statistics and errors counters for export to outside collectors at configurable intervals using remote procedure call (gRPC) services.

Use the following resource path to export statistics:

`/junos/system/linecard/services/inline-jflow/`

To provision the sensor to export data through gRPC services, 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 JTI.

[See [Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates](#), [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#) and [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Packet Forwarding Engine support for JTI (MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.2R1, Junos telemetry interface (JTI) supports streaming of Packet Forwarding Engine statistics for MX2010 and MX2020 routers using Remote Procedure Calls (gRPC). gRPC is a protocol for configuration and retrieval of state information.

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

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

- **Sensor- level statistics support on JTI (MX960, MX2008, MX2010, MX2020, PTX5000, PTX1000, and PTX10000 routers and QFX5100 and QFX5200 switches)**—Starting with Junos OS Release 19.2R1, you can issue the Junos operational mode command **show network-agent statistics** to provide more information on a per-sensor level for statistics being streamed to an outside collector by means of remote procedure calls (gRPC) and Junos telemetry interface (JTI). Only sensors exported with gRPC are supported. The command does not support UDP-based sensors.

[See [show network-agent statistics](#) and [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **ONCE mode supported using gNMI services and JTI (MX Series)**—Starting in Junos OS Release 19.2R1, you can include the "ONCE" mode with the **Subscribe** RPC when subscribing to gRPC Network Management Interface (gNMI) services to export statistics for telemetry monitoring and management using Junos telemetry interface (JTI). ONCE mode ensures that the collector is only streamed telemetry information one time.

The Subscribe RPC and subscription parameters are defined in the [gnmi.proto](#) file.

Streaming telemetry data through gNMI 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 JTI.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Packet Forwarding Engine statistics export using gNMI and JTI (MX960, MX2008, MX2010 and MX2020 routers)**—Starting in Junos OS Release 19.2R1, you can stream Packet Forwarding Engine statistics to an outside collector using gRPC Management Interface (gNMI) version 0.7.0 and Junos telemetry interface (JTI). Prior to this, these statistics were exported using OpenConfig gRPC and UDP protocol buffer (gpb) format. OpenConfig gRPC and gNMI are both protocols used to modify and retrieve configurations as well as export telemetry streams from a device in order to manage and monitor it

To provision Packet Forwarding Engine sensors to export data through gNMI, use the Subscribe RPC defined in the [gnmi.proto](#) to specify request parameters. This RPC already supports Routing Engine statistics to be exported by means of gNMI. Now, Packet Forwarding Engine sensors will also stream KV pairs in gNMI format for a majority of Packet Forwarding Engine sensors.

Streaming telemetry data through gNMI 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 JTI.

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

- **Broadband edge statistics support through JTI (MX Series)**—Starting in Junos OS Release 19.2R1, subscriber-based telemetry streaming is enabled when an MX Series router is configured for Broadband Network Gateway (BNG) and Junos fusion where subscribers are connected through Junos fusion Satellite devices. You can use remote procedure calls (gRPC) to export broadband edge (BBE) telemetry statistics to external collectors. gRPC is a protocol for configuration and retrieval of state information.

You can stream all BBE resource paths except for the following:

- `/junos/system/subscriber-management/access-network/ancp`
- `/junos/system/subscriber-management/client-protocols/l2tp`
- `/junos/system/subscriber-management/infra/network/l2tp/`

To stream BBE statistics, include a resource path starting with `/junos/system/subscriber-management/` in your gRPC subscription.

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

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

- **gRPC-based streaming telemetry support for subscriber service accounting statistics for JTI (MX Series 5G Universal Routing Platform)**—Starting in Junos OS Release 19.2R1, you can enable service filter accounts statistics for subscribers using Junos telemetry interface (JTI) and remote procedure calls (gRPC). Service accounting statistics include IP protocol IPv4 family, IPv6 family, or both, as well as transmit and receive packets and bytes for subscriber service sessions.

To enable these statistics from an MX Series router, include the **service-statistics** statement at the **[edit dynamic-profiles my-service-profile telemetry]** hierarchy level.

To stream these statistics, include the resource path `/junos/system/subscriber-management/dynamic-profiles/interfaces/services/` in your gRPC subscription to export the statistics to an outside collector.

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

[See [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#) **service-statistics**, and [Enable Service Filter Accounting Statistics for Subscribers](#).]

- **FPC and optics support for JTI (MX Series)**—Starting in Junos OS Release 19.2R1, Junos telemetry interface (JTI) supports streaming of Flexible PIC Concentrator (FPC) and optics statistics for the MX Series router using remote procedure calls (gRPC). gRPC is a protocol for configuration and retrieval of state information. This feature effort includes the addition of a new process (SensorD daemon) to export telemetry data for integration with AFTTelemetry and LibTelemetry libraries in the OpenConfig model called AFT platform.

The following base resource paths are supported:

- `/junos/system/linecard/environment/`
- `/junos/system/linecard/optics/`
- `/junos/system/linecard/optics/optics-diag[if-name =]`
- `/junos/system/linecard/optics/optics-diag/if-name`

- `/junos/system/linecard/optics/optics-diag/snmp-if-index`
- `/junos/system/linecard/optics/lane[lane_number=]/`

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

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

- **Specify Routing Instance for JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R1, you can specify the routing instance to use for remote procedure call (gRPC) services. Include the **routing-instance** *instance-name* at the **[edit system services extension-service request-response grpc]** hierarchy level. The routing instance name specified should match the name of the existing routing instance, such as a name configured under the **[routing-instances]** hierarchy level or **mgmt_junos** if **system management-instance** is configured (the dedicated management routing instance).

Configuring the routing instance lets you choose the VRF for gRPC services. When the routing instance is not configured, the default behavior is that all gRPC-related services are available through the management **fxp0/em0** interface.

Layer 2 VPN

- **Support for group key acknowledgment messages (MX Series)**—Starting with Junos OS Release 19.2R1, Junos OS supports group members sending acknowledgment messages as defined in RFC 8263 in response to group key push messages sent by group controllers and key servers. The group member sends acknowledgment messages when it receives a group key push message with a standard KEK_ACK_REQUESTED value of 9 in the SA KEK payload as defined in RFC 8263 or a KEK_ACK_REQUESTED value of 129 that is used in older key servers. No additional configuration is required.

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

Layer 2 Features

- **Support for basic Layer 2 features on MPC10E-15C-MRATE line card (MX Series)**—Starting in Junos OS Release 19.2R1, MPC10E-15C-MRATE line card supports the following basic Layer 2 features:
 - Layer 2 bridging with trunk and access modes
 - MAC learning and aging
 - Handling BUM (broadcast, unknown unicast and multicast) traffic, including split horizon
 - MAC move
 - Layer 2 forwarding and flooding statics
 - Mesh groups

- Static MAC addresses
- MAC learning and forwarding on AE interfaces
- Bridging on untagged interfaces
- Basic Q-n-Q tunneling (without VLAN-translation and VLAN map operations)

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

Layer 3 Features

- **MPC10E-10C and MPC10E-15C support layer 3 routing features (MX240, MX480, and MX960)**—Starting in Junos OS Release 19.2R1, MPC10E-10C and MPC10E-15C line cards support the following features in hyper-mode:
 - Configuring ICMP redirects and generating ICMP redirect messages.
 - Padding VLAN packets to a minimum frame size of 68 bytes, by using the existing command **set interfaces *interface-name* gigether-options pad-to-minimum-frame-size**.
 - Collecting interface family statistics for IPv4 and IPv6, by using the existing command **show interfaces statistics detail *interface-name***.

See [Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches](#)

MPLS

- **Dynamic creation of segment routing LSPs using BGP protocol next hops (MX Series)**—Starting in Junos OS Release 19.2R1, you can configure tunnel templates on colored and non-colored segment routing traffic-engineered (SR-TE) paths. These templates enable dynamic creation of segment routing tunnels using protocol next hops with BGP prefixes to resolve destination segment identifiers (SIDs).

With this feature, you can benefit from reduced configuration, especially when the network deployment requires connectivity from each provider edge (PE) device to every other PE device.

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

- **CSC support for MPLS-over-UDP tunnels (MX Series with MPC and MIC and VMX)**—Starting in Junos Release 19.2R1, carrier supporting carrier (CSC) architecture can be deployed with MPLS-over-UDP tunnels carrying MPLS traffic over dynamic IPv4 UDP tunnels that are established between supporting carrier's provider edge (PE) devices. With this enhancement, the scaling advantage that the MPLS-over-UDP tunnels provided is further increased. This feature is not supported on IPv6 UDP tunnels.

[See [Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels](#).]

Network Management and Monitoring

- **Support for displaying valid user input in the CLI for command options and configuration statements in custom YANG data models (MX Series)**—Starting in Junos OS Release 19.2R1, the CLI displays the set of possible values for a given command option or configuration statement in a custom YANG data

model when you include the **action-expand** extension statement in the option or statement definition and reference a script that handles the logic. The **action-expand** statement must include the **script** child statement, which defines the Python action script that is invoked when a user requests context-sensitive help in the CLI for the value of that option or statement.

[See [Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules](#).]

- **Support for Synchronous Ethernet with ESMC on JNP10K-LC2101 (MX10008 and MX10016)**—Starting in Junos OS Release 19.2R1, the JNP10K-LC2101 line card supports Synchronous Ethernet (SyncE) with ESMC. Synchronous Ethernet is a physical layer technology that is used to transfer clock signals over Ethernet interfaces. ESMC transmits Synchronization Status Message (SSM) information, which is the quality level of the transmitting synchronous Ethernet equipment clock (EEC), by using ESMC protocol data units (PDUs). This support allows you to configure BITS-0 (external-0) and BITS-1 (external-1) ports as clock sources or outputs on master Routing and Control Board (JNP10K-RE1). You can also configure a GPS (external-2) port as a clock source on master Routing and Control Board. This feature also supports SyncE over aggregated Ethernet (AE).

NOTE: Only the GPS port and BITS ports that are configured on master RCB are active.

[[Centralized Clocking Overview](#) and [Understanding ESMC Quality Level Mapping](#)]

- **Support for optimizing the SNMP walk execution time for IPsec statistics (MX Series)**—Starting in Junos OS Release 19.2R1, you can optimize the SNMP walk execution time for IPsec statistics. To achieve this optimization, increase the cache lifetime of the IPsec related information (for example statistics and SA information) so that a single SNMP walk request is served for N number of IPsec Security Associations (SAs) with N number of queries made to the service PIC. IPsec statistics are now fetched by the burst mode, thereby reducing the load on the Routing Engine daemon, kmd. For different scale needs, we may have to tweak the hidden SNMP knob parameters, for example, with Dead Peer detection (DPD) having more number of tunnels without traffic and simultaneous SNMP walks.

Port Security

- **Fallback PSK for Media Access Control Security (MACsec) (MX Series)**—Starting in Junos OS Release 19.2R1, fallback PSK for MACsec is supported on MX Series routers that support MACsec. The fallback PSK provides functionality to establish a secure session in the event that the primary PSKs on each end of a MACsec-secured link do not match.

[See [Configuring Media Access Control Security \(MACsec\) on MX Series Routers](#).]

Routing Policy and Firewall Filters

- **Support for CCC and Layer 3 firewall forwarding on MPC10E-15C-MRATE line cards (MX Series)**—Starting with Junos OS Release 19.2R1, circuit cross-connect (CCC) traffic and Layer 3 firewall forwarding features are supported on MPC10E-15C-MRATE line cards.

[See [CCC Overview](#) and [Protocols and Applications Supported by the MPC10E-15C-MRATE](#).]

Routing Protocols

- **MPC10 Inline BFD support (MX Series)**—Starting in Junos OS Release 19.2, MPC10 MPCs support inline BFD features, excluding micro BFD and BFD sessions with authentication.

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

- **Support for IPv6 fragment reassembly for v4ov6 dynamic tunnels**—Starting in Junos OS Release 19.2R1, you can configure an additional attribute, **dynamic-tunnel-reassembly-enable** for reassembling IPv6 fragments before the termination of v4ov6 tunnels. The fragment reassembly feature is disabled by default. IPv6 fragments are discarded when this feature is not enabled.
- **IPv6 reassembly for v4ov6 tunnels (MX Series)**—Starting in Junos OS 19.2R1, you can enable the MX chassis to perform IPV6 fragment reassembly for forwarding Ipv4 traffic. When the **dynamic-tunnel-reassembly** is configured, the tunnels using the attribute would be setup for reassembling the IPv6 fragments before the termination of v4ov6 tunnels. By default, this attribute is turned off and the tunnels are set up to discard the IPv6 fragments.

To enable IPv6 fragment reassembly for forwarding Ipv4 traffic, use **set dynamic-tunnel-reassembly on** statement at the **[edit routing-options dynamic-tunnels tunnel-attributes <dynamic-tunnel-name>]** hierarchy level.

[See [dynamic-tunnel-reassembly](#).]

- **Map single IPv6 anycast address on multiple anchor Packet Forwarding Engines (MX240, MX480, MX960, MX2020)**—Starting in Junos OS Release 19.2R1, you can assign the same IPv6 anycast address to multiple anchor Packet Forwarding Engines to manage high traffic from CPE to internet. By default, this feature is disabled. Prior to Junos OS Release 19.2R1, you can assign an anycast address only to a single Packet Forwarding Engine and the maximum v4ov6 tunnel scale per Packet Forwarding Engine in MX Series is 150k. This restricts a single anycast address to be used for 150k tunnels.

To configure the same source address over multiple tunnel-attributes, use **set v4ov6 ipv6-anycast-source-duplication** statement at the **[edit routing-options dynamic-tunnels]** hierarchy level.

If v4ov6 packets are fragmented, the fragmented packets get steered to one of the anchor Packet Forwarding Engines for IPv6 reassembly processing. To steer the traffic to the correct anchor, Packet Forwarding Engine needs information about the range of IPv4 prefixes that goes over a particular tunnel. To get the range of IPv4 prefixes that goes over a particular tunnel, use set **get-route-range** statement at the **[edit policy-options policy-statement <policy-name> term <term-name> from route-filter <route-filter-value> <range>]** hierarchy level.

[See [v4ov6](#) and [get-route-range](#).]

- **Support for export of BGP Local RIB through BGP Monitoring Protocol (BMP) (MX Series)**—Starting in Junos OS Release 19.2R1, BMP is enhanced to support monitoring of local RIB (**loc-rib**) policy. The **loc-rib** policy is added to RIB types under the **bmp route-monitoring** statement.

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

- **Support for BGP routes with N-Multipath primary and 1-Protection backup gateway (MX Series)**—Starting in Junos OS 19.2R1, the following enhancements are made to the Junos OS:
 - Support N+1 formation for BGP labelled unicast protection (LU).
 - Support N+1 formation for BGP PIC (IPv4, IPv6, LU).
 - Support for hetero-nexthops (ListNH) in such N+1 formations.
 - Support for KRT to defer fib-update if BGP-multipath is in progress.
 - Removed restriction to use **delay-route-advertisement** statement for IPv4 labeled-unicast.
 - Four new options **import**, **install-address <address>**, **no-install**, and **rib (inet.0 | inet6.0)** are added under the **egress-te** statement.
 - A new configuration statement **allow-protection** is introduced to allow protection for multipath legs. To allow protection for multipath legs, use **set allow-protection** statement at the **[edit protocols bgp multipath]** hierarchy level.
 - A new option **always-wait-for-krt-drain** is introduced under **delay-route-advertisement** statement to make more-specific BGP-routes re-advertisement to wait for KRT-queue to drain. To configure this, use **set always-wait-for-krt-drain** at the **[edit protocols bgp family inet unicast delay-route-advertisements]** hierarchy level.

[See [allow-protection \(Multipath\)](#), [delay-route-advertisements](#) and [egress-te](#).]

Security

- **Juniper Malware Removal Tool**—Starting in Junos OS Release 19.2R1, the Juniper Malware Removal Tool (JMRT) can be used to scan and remove malware running on Junos OS devices. To run JMRT, use the operational commands under the **request system malware-scan** hierarchy. There are 2 types of scans you can perform with JMRT:

Quick—Scan each running program file.

Veriexec check—Check if verified execution is enabled.

[See [request system malware-scan](#).]

Services Applications

- **Support for IPv6 BGP next-hop address in IPv6 and MPLS-IPv6 inline flow record templates(MX Series)**—Starting in Junos OS Release 19.2R1, a new element, IPv6 BGP NextHop Address, is available in the the IPv6 inline flow record template and the MPLS-IPv6 inline flow record template to add support for IPv6 BGP NextHop information element. The new element is supported on both version 9 and version 10 (IPFIX) export formats. The element ID is 63 and the element size is 16 bytes.

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

- **IPv4 and IPv6 version 9 templates for inline active flow monitoring (MPC10E-15C-MRATE on MX Series)**—Starting in Junos OS Release 19.2R1, while configuring inline active flow monitoring, you can apply version 9 flow templates to define a flow record template suitable for IPv4 or IPv6 traffic.

[See [Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates](#).]

- **Support for Two-Way Active Measurement Protocol (TWAMP) on MPC10E-15C-MRATE line card**—Starting in Junos OS Release 19.2R1, TWAMP is supported on MPC10E line card on the MX240, MX480, and MX960 routers. TWAMP defines a standard for measuring IPv4 performance between two devices in a network. You can use the TWAMP-Control protocol to set up performance measurement sessions between a TWAMP client and a TWAMP server, and use the TWAMP-Test protocol to send and receive performance measurement probes.

Configuring the TWAMP client instance to use *si-x/y/z* as the destination interface (which enables inline services) is not supported if the router has an MPC10E-15C-MRATE installed in the chassis. You can configure only the **none** authentication mode on the line card.

[See [Understanding Two-Way Active Measurement Protocol on Routers](#).]

- **DS-Lite support on MX Virtual Chassis and MX BNG**—Starting in Junos OS Release 19.2R1, the MX Series Virtual Chassis and MX Series broadband network gateway (BNG) support dual-stack lite (DS-Lite). DS-Lite uses IPv4-over-IPv6 tunnels to traverse an IPv6 access network to reach a carrier-grade IPv4-IPv4 NAT. DS-Lite enables the phased introduction of IPv6 on the Internet by providing backward compatibility with IPv4.

DS-Lite on the MX Series Virtual Chassis and MX Series BNG does not support the following:

- Application Layer Gateways (ALGs)
- Limits per subnet
- Clearing NAT mappings and flows for a specific subscriber, for a basic bridging broadband device (B4), or for a specific service set
- Port Control Protocol

[See [Tunneling Services for IPv4-to-IPv6 Transition Overview](#).]

- **Hardware timestamping of RPM probe messages**—Starting in Junos OS Releases 19.2R1, you can enable timestamps on RPM probes messages in the Packet Forwarding Engine host processor for the following line cards:
 - MPC10E-15C-MRATE line card on MX240, MX480, and MX960 routers
 - MPC11E line card on MX2008, MX2010, and MX2020 routers

You can use the following configuration statements at the `[edit services rpm probe owner test test-name]` hierarchy level:

- **hardware-timestamp**—Enables timestamping of RPM probe messages in the Packet Forwarding Engine host Processor.
- **one-way-hardware-timestamp**—Enables timestamping of RPM probe messages for one-way delay and jitter measurements.

These configuration statements are supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types.

See [\[hardware-timestamp\]](#)

[\[one-way-hardware-timestamp\]](#)

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers](#)

- **Increased number of AMS members supported on single chassis (MX2020)**—Starting in Junos OS Release 19.2R1, you can configure up to 60 MS-PICs as part of aggregated multiservices (AMS) bundles on a single chassis. The configuration supports backup and load-balancing mode (N:1) and all active mode (N:0) with both next-hop style services and interface style services of configurations.

See [\[Understanding Aggregated Multiservices Interfaces\]](#).

- **IPFIX flow-cache support (MX150)** —Starting in Junos OS Release 19.2R1, the flow cache infrastructure support is extended to IPFIX to provide improved throughput with IPFIX service enabled. In earlier releases, without flow cache support for IPFIX, all data traffic would take the microcode path which is much slower than flow cache. With this feature, the unsampled traffic gets forwarded using flow cache which results in better throughput.

Software Defined Networking

- **PCE-initiated bypass LSPs (MX Series)**—Starting in Junos OS Release 19.2R1, the Path Computation Element Protocol (PCEP) functionality is extended to allow a stateful Path Computation Element (PCE) to initiate, provision, and manage bypass label-switched paths (LSPs) for a protected interface. Multiple bypass LSPs with bandwidth reservation can be initiated by the PCE to protect a resource.

With this feature, you can benefit from the LSP state synchronization of manual, dynamic, and PCE-initiated bypass LSPs from a PCE, and leverage on the PCE's global view of the network, resulting in better control over traffic at the time of a failure, and deterministic path computation of protection paths.

[See [Support of the Path Computation Element Protocol for RSVP-TE Overview](#).]

- **Support for unified ISSU on abstracted fabric interfaces (MX480, MX960, MX2010, MX2020, MX2008)**—Starting in Junos OS Release 19.2R1, abstracted fabric (af) interfaces, configured for Junos Node Slicing, support unified in-service software upgrade (ISSU). Unified ISSU enables an upgrade between two Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

NOTE: Since the af interface traffic is load balanced across all available Packet Forwarding Engines, the traffic loss on an AF interface during ISSU might be higher, compared to the traffic loss on a regular interface.

An af interface is a pseudo interface that represents a first class Ethernet interface behavior. An AF interface facilitates routing control and management traffic between guest network functions (GNFs) through the switch fabric.

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

- **Centralized assignment of unique MAC addresses to GNFs (MX960, MX2008, MX2010, and MX2020)**—Starting in Junos OS Release 19.2R1, Junos node slicing supports the assignment of a globally unique MAC address range (supplied by Juniper Networks) for GNFs. To receive the globally unique MAC address range for the GNFs, contact your Juniper Networks representative and provide your GNF license SSRN (Software Support Reference Number), which will have been shipped to you electronically upon your purchase of the GNF license. For each GNF license, you will then be provided an 'augmented SSRN', which includes the globally unique MAC address range assigned by Juniper Networks for that GNF license. You must then configure this augmented SSRN at the JDM CLI as follows:

set system vnf-license-supplement vnf-id *gnf-id* license-supplement-string *augmented-ssrn-string*.

[See [Assigning MAC Addresses to GNF](#)]

- **Support for IPSec, stateful firewall, and CGNAT services on MS-MPCs over abstracted fabric interfaces (MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 19.2R1, guest network functions (GNF) support Layer 3 services such as Carrier-Grade Network Address Translation (CGNAT), stateful firewall, and IP Security (IPsec) on Multiservices MPCs (MS-MPCs) over abstracted fabric (af) interfaces.

[See [Abstracted Fabric Interface](#)]

- **MX2008 routers support in-chassis Junos node slicing (MX Series)**—Starting in Junos OS Release 19.2R1, MX2008 routers support the in-chassis model of Junos node slicing deployment. In the in-chassis model, the base system (BSYS), Juniper Device Manager (JDM), and all guest network functions (GNFs) run within the Routing Engine of the MX Series router. To support in-chassis Junos node slicing, the MX2008 must have the routing engine REMX2008-X8-128G installed.

[See [Configuring MX Series Router to Operate in In-Chassis Mode](#)]

Software Installation and Upgrade

- **The curl binary is packaged and made available on all Junos OS variants (MX Series)**—The curl binary is a command-line utility, used from the shell, that you can use to perform operations over several transport protocols, including the following: dict, file, ftp, gopher, http, imap, pop3, rtsp, smtp, telnet, tftp. The features enabled on Junos OS are curl version 7.59, libcurl version 7.59.

Subscriber Management and Services

- **Support for M:N subscriber redundancy on BNGs (MX Series)**—Starting in Junos OS Release 19.2R1, you can configure broadband network gateways (BNGs) to provide interface-level redundancy for DHCP subscribers that are on the same static VLAN and use the same access interface. Failover from master to backup BNG is transparent to the clients because the subscriber sessions remain up. You must configure DHCP active leasequery with topology discovery on peer DHCP relay agents on the master and backup BNGs to support the redundancy.

[See [M:N Subscriber Redundancy](#).]

- **Support for Interface-Level Redundancy with DHCP Topology Discovery (MX Series)**—Starting in Junos OS Release 19.2R1, you can configure DHCP active leasequery with topology discovery to provide interface-level subscriber redundancy between peer relay agents. Topology discovery enables master and backup peer relay agents to determine the access interfaces on peers that correspond to their own local access interfaces for servicing subscriber redundancy groups. During synchronization, DHCP translates the subscriber binding information to use the local interface on the backup instead of the interface on the master. You must use topology discovery when you configure M:N subscriber redundancy.

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

- **Support for fixed wireless access subscribers on BNGs (MX Series)**—Starting in Junos OS Release 19.2R1, you can configure the broadband network gateway (BNG) to support subscribers that use a fixed wireless network. Providers use a wireless network for subscriber access over the air instead of than running fiber to the home. The wireless infrastructure saves costs and reduces complexity compared to the fiber network. The BNG acts as the Third-Generation Partnership Project (3GPP) System Architecture Evolution Gateway (SAEGW). The SAEGW incorporates the functions of both the Serving Gateway (SGW) and the Packet Data Network Gateway (PGW). The SGW function routes and forwards user data packets. The PGW function provides connectivity to external packet data networks

[See [Fixed Wireless Access Networks](#).]

System Management

- **Support for transferring accounting statistics files and router configuration archives using HTTP URL (MX Series)**—Starting in Junos OS Release 19.2R1, you can transfer accounting statistics files and router configuration archives to remote servers by using an HTTP URL. In addition to SCP and FTP, the following HTTP URL will be supported under the **archive-sites** statement:

`http://username@host:url-path password password`

- To transfer accounting statistics files, configure **archive-sites** under **[edit accounting-options file <filename>]** hierarchy.
- To transfer router configuration archival, configure **archive-sites** under **edit system archival configuration** hierarchy.
- To view the statistics of transfer attempted, succeeded, and failed, use the **show accounting server statistics archival-transfer** command.
- To clear the statistics of transfer attempted, succeeded, and failed, use the **clear accounting server statistics archival-transfer** command.

[See [archive-sites](#), [Backing Up Configurations to an Archive Site](#), [show accounting server statistics archival-transfer](#), and [clear accounting server statistics archival-transfer](#)].

Timing and Synchronization

- **Support for Synchronous Ethernet with ESMC on MPC10E-15C-MRATE (MX240, MX480, MX960)**—Starting in Junos OS Release 19.2R1, MPC10E-15C-MRATE supports Synchronous Ethernet with ESMC. Synchronous Ethernet is a physical layer technology that is used to transfer clock signals over Ethernet interfaces. It supports hop-by-hop frequency transfer, where all interfaces on the trail must support Synchronous Ethernet.

ESMC is a logical communication channel. It transmits Synchronization Status Message (SSM) information, which is the quality level of the transmitting synchronous Ethernet equipment clock (EEC), by using ESMC protocol data units (PDUs).

[See [Synchronous Ethernet Overview](#)].

SEE ALSO

[What's Changed | 113](#)

[Known Limitations | 123](#)

[Open Issues | 127](#)

[Resolved Issues | 139](#)

[Documentation Updates | 191](#)

[Migration, Upgrade, and Downgrade Instructions | 192](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.2R3-S1 | 114](#)
- [What's Changed in Release 19.2R3 | 114](#)
- [What's Changed in Release 19.2R2 | 115](#)
- [What's Changed in Release 19.2R1 | 120](#)

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

What's Changed in Release 19.2R3-S1

General Routing

- **New commit check for MC-LAG (MX Series)**—We've introduced a new commit check to check the values assigned to the redundancy group identification number on the MC-AE interface (**redundancy-group-id**) and ICCP peer (**redundancy-group-id-list**) when you configure multichassis aggregation groups (MC-LAGs). If the values are different, the system reports a commit check error. In previous releases, if the configured values were different, the l2ald process would crash.

[See [iccp](#) and [mc-ae](#).]

Infrastructure

- **Change in support for interface-transmit-statistics statement (MX Series)**—You cannot configure aggregated Ethernet interfaces to capture and report the actual transmitted load statistics by using the **interface-transmit-statistics** statement. Aggregated Ethernet interfaces do not support reporting of the transmitted load statistics. The **interface-transmit-statistics** statement is not supported in the aggregated Ethernet interfaces hierarchy. In earlier releases, the **interface-transmit-statistics** statement was available in the aggregated Ethernet interfaces hierarchy but not supported.

[See [interface-transmit-statistics](#).]

What's Changed in Release 19.2R3

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](#) and [Resource Monitoring for Subscriber Management and Services](#).

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**—You can set the verbosity of the trace log to only show error messages using the error option at the [edit system services extension-service traceoptions level] hierarchy.

[See [traceoptions \(Services\)](#).]

Network Management and Monitoring

- **Enhancement to the show snmp mib command**— Starting in Junos OS Release 19.2R3, a new option, **hex**, is supported to display the SNMP object values in the hexadecimal format. In earlier releases, the **show snmp mib** command displays the SNMP object values in ASCII and decimal format only.

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

Routing Protocols

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

Services Applications

- **New option for configuring delay in IPsec SA installation**—In Junos OS Release 19.2R3, you can configure the `natt-install-interval seconds` option at the `[edit services ipsec-vpn rule rule-name term term-name then dynamic]` hierarchy level to specify the duration of delay in installing IPsec SA in a NAT-T scenario soon after the IPsec SA negotiation is complete. The default value is 0 seconds.

Subscriber Management and Services

- **Improved tunnel session limits display (MX Series)**—Starting in Junos OS Release 19.2R3, the `show services l2tp tunnel extensive` command displays the configured value for maximum tunnel sessions. On both the LAC and the LNS, this value is the minimum from the global chassis value, the tunnel profile value, and the value of the Juniper Networks VSA, Tunnel-Max-Sessions (26–33). On the LNS, the configured host profile value is also considered.

In earlier releases, the command displays the value 512,000 on the LAC and the configured host profile value on the LNS.

[See [Limiting the Number of L2TP Sessions Allowed by the LAC or LNS.](#)]

What's Changed in Release 19.2R2

General Routing

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

In earlier Junos OS releases, the confirmation prompt is available for only the main options.

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on `ge`, `et`, `xe` interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for `ge` interfaces, earlier when you view the `show interfaces` command, by default,

only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

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

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

Junos OS XML API and Scripting

- **Root XML tag change for show rsvp pop-and-forward | display xml command (MX480)**—We've changed the root XML tag for the show rsvp pop-and-forward | display xml command to rsvp-pop-and-fwd-information to make it consistent with the XML tag convention. In earlier releases, the command output displays rsvp-pop-and-fwd-info XML tag. Update the scripts with the rsvp-pop-and-fwd-info XML tag to reflect the new rsvp-pop-and-fwd-information XML tag.
- [See [Junos XML API Explorer - Operational Tags](#)].

Interfaces and Chassis

- **Change in error severity (MX960, MX240, MX2020, MX480, MX2008, and MX2010)**—Starting in Junos OS Release 19.2R2, we have reduced the severity of the CRC errors (XR2CHIP_ASIC_JGCI_FATAL_CRC_ERROR) from Fatal to Major. Earlier, these errors caused the line card to be reset, if the CLI command **interasic-linkerror-recovery-enable** was configured. Now, these errors only disable the Packet Forwarding Engines that are affected. With this change, the **interasic-linkerror-recovery-enable** configuration has no effect when these errors occur because of the reduced severity.

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

- **Logical Interface created along with physical Interface by default (MX Series routers)**—In Junos OS Release 19.2R2 and later, 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, when you view the **show interfaces** command in earlier releases, by default, only the physical interface (for example, ge-0/0/0), is displayed. Now, the logical interface (for example, ge-0/0/0.16386) is also displayed.

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, 19.2R2, and later, MX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

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

MPLS

- **Deprecated statement (MX Series)**—Starting in Junos OS Release 19.2R2, we have deprecated the **preference** statement at the **[edit protocols source-packet-routing source-routing-path name]** hierarchy level. This is because you could have two different sequences of the same route, wherein the active route entry that is selected can be different.

Network Management and Monitoring

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

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

OAM

- **Performance monitoring history data is lost when a change in number of supported history records is detected (ACX Series and MX Series)**—In Junos OS Release 19.2R2, when Ethernet connectivity fault management starts, it detects the number of history records supported by the existing Performance Monitoring history database and if there is any change from the number of history records supported (that is, 12) in Release 19.2R2, then the existing Performance Monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

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

19.2R2, we have included an XML RPC equivalent for the **show bgp output-scheduler | display xml rpc** CLI command. In Junos OS releases before Release 19.2R2, the **show bgp output-scheduler | display xml rpc** CLI command does not have an XML RPC equivalent.

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

- **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.

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 (MX Series Services Applications)**—In Junos OS Release 19.2R2, the **version-3** option under the **[edit services softwire softwire-concentrator map-e]** hierarchy is optional. The version number helps distinguish between the currently supported version of the Internet draft draft-ietf-softwire-map-03 is optional. In earlier Junos OS releases, if you do not configure the **version-3** option, the configuration results in an error.

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

- **Change in NAT port block syslog message display (MX Series routers)**—When you configure a softwire prefix other than 128, all the JSERVICES_NAT_PORT_BLOCK logs now display the prefixed B4 address. We have modified the following JSERVICES_NAT_PORT_BLOCK logs:
 - JSERVICES_NAT_PORT_BLOCK_ALLOC
 - JSERVICES_NAT_PORT_BLOCK_RELEASE
 - JSERVICES_NAT_PORT_BLOCK_ACTIVE

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

Software Defined Networking (SDN)

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

With the increase in the 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 may disrupt the PCEP session with the main active stateful PCE.

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

Subscriber Management and Services

- **Enhancement to commands to display reason for Routing Engine disconnect (MX Series)**—Starting in Junos OS Release 19.2R2, several commands display the reason when the master and standby Routing Engines disconnect because of a memory mismatch error. On a chassis with two Routing Engines, a DRAM size mismatch error can result when both of the following are true:
 - The Routing Engines have different amounts of DRAM.
 - A 64-bit Junos OS image is loaded on the chassis.

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

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

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

- **Prevent queue-based throttling from stopping subscriber login (MX Series)**—Starting in Junos OS Release 19.2R2, 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.
- **XML output format change for test aaa type user commands (MX Series)**—Starting in Junos OS Release 19.2R2, the XML output format changes for the **test aaa authd-lite user**, **test aaa dhcp user**, and **test aaa ppp user** commands. Each RADIUS server attribute name has an associated attribute value. Each of these pairs is now enclosed by the `<radius-server-data>` tag. The new tag makes it easier to recognize the name/value pairs, both for operators and API clients. You may have to change any scripts that use the XML output to work properly with the new format.

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

- **Support for Pseudowire Physical Interface for ANCP Autoconfiguration (MX Series)**—Starting in Junos OS Release 19.2R2, you can associate an ANCP neighbor with a subscriber-facing pseudowire physical interface for ANCP autoconfiguration of VLANs. When configured, ANCP Port Up and Port Down

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

What's Changed in Release 19.2R1

EVPN

- **Support for disabling automatic ESI generation (MX Series and QFX Series)**—Starting with Junos OS Release 19.2R1, Junos OS supports disabling the automatic ESI generation for virtual gateway addresses. We recommend that you disable the automatic ESI generation for EVPN networks with edge-routed bridging to improve performance. To disable automatic ESI generation, include the **no-auto-virtual-gateway-esi** statement at the **[edit interfaces name irb unit logical-unit-number]** hierarchy level.

Interfaces and Chassis

- **Deprecation of the [edit fabric protocols bgp] hierarchy level (MX Series)**—Starting in Junos OS Release 19.2R1 and later, the **[edit fabric protocols bgp]** hierarchy level is deprecated.
- **Support to get Optics Loopback Status for QSFP-100GE-DWDM2 transceivers (MX Series)**—In Junos OS Release 19.2R1, and later, on MX Series routers, you can get the optics loopback status of QSFP-100GE-DWDM2 transceivers along with the regular ethernet loopback status by issuing the **show interfaces interface-name** or **show interfaces interface-name brief** command. New Output field **Optics Loopback** is added under **Link-level type** when **show interfaces interface-name** CLI command is executed.
- **Monitoring information available only in Trace log (MX Series)**—In Junos OS Release 19.2R1 and later, the Ethernet link fault management daemon (lfmd) in the peer router stops monitoring the locally occurred errors until ISSU completes. You can view the monitoring-related details only through the trace log file.
- **Health check for power supplies (MX10008 and MX10016)**—Starting in Junos OS Release 19.2R1, on the MX10008 and MX10016 routers, the **show chassis environment pem** command displays the health check information about the DC or AC Power supplies. For any power supply that does not support health check, the status is shown as **Unsupported**. The system starts health check of a power supply only if the power consumption exceeds 7 KW.

[See [show chassis environment pem](#)]

MPLS

- **New debug statistics counter (MX Series)**—The **show system statistics mpls** command has a new output field, called **Packets dropped, over p2mp composite nexthop**, to record the packet drops over composite point-to-multipoint next hops.
- **IPv4 explicit-null label retained from the merged protocol MPLS label stack**—The IPv4 explicit-null label is retained from the merged protocol MPLS label stack, if the IPv4 explicit-null is at the bottom of the MPLS label stack.

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (MX Series)**—Starting in Junos OS Release 19.2R1, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.
- **Custom YANG RPC support for input parameters of type empty (MX Series)**—Starting in Junos OS Release 19.2R1, custom YANG RPCs support input parameters of type **empty** when executing the RPC's command in the Junos OS CLI, and the value passed to the action script is the parameter name. In earlier releases, input parameters of type **empty** are only supported when executing the RPC in a NETCONF or Junos XML protocol session, and the value passed to the action script is the string '**none**'.

[See [Creating Action Scripts for YANG RPCs on Devices Running Junos OS](#).]

- **Change in power supply alarms (MX10003)**—Starting in Junos OS Release 19.2R1, the MX10003 routers do not raise an alarm if a Power Entry Module (PEM) slot is empty. However, when the number of operational PEMs available is less than 2, the router raises a major alarm. This alarm is cleared when the required number of PEMs are made available.

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

Routing Policy and Firewall Filters

- **Fixed an issue with certain combination of match conditions**—In Junos OS Release 19.2R1, fixed a temporary issue wherein configuring a firewall filter with a match condition for **port** along with **source-port** and/or **destination-port** in the same filter term would cause a commit error. Any valid combination of the filter terms is now supported.

Services Applications

- **Support for host generated traffic on a GRE over GRE tunnel (MX Series)**—In Junos OS Release 19.2R1, you can send host generated traffic on a GRE over GRE tunnel. However, when path maximum transmission unit (PMTU) is updated for the outer GRE tunnel, MTU for inner GRE tunnel is not corrected.
- **New syslog message displayed during NAT port allocation error (MX Series Routers with MS MPC)**—With address pooling paired (APP) enabled, an internal host is mapped to a particular NAT pool address. In case, all the ports under a NAT pool address are exhausted, further port allocation requests from the internal host results in a port allocation failure. The following new syslog message is displayed during such conditions:

JSERVICES_NAT_OUTOF_PORTS_APP

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

Software Defined Networking

- **Deprecated CLI commands and options for JDM (MX480, MX960, MX2010, MX2020, and MX2008)**—Starting in Junos OS Release 19.2R1, in Junos Node Slicing, Juniper Device Manager (JDM) does not support the following CLI commands or options:
 - **show system visibility**
 - **show system inventory**
 - the **jinventoryd** option in the **restart** command

Subscriber Management and Services

- **Changing attributes of physical interface with active subscribers (MX Series)**—Starting in Junos OS Release 19.2R1, the commit check fails when you change any attribute of the physical interface, such as the MTU, when subscribers are active. This affects only aggregated Ethernet physical interfaces with targeted distribution configured. In earlier releases, the commit check does not fail and the attribute change brings down the physical interface and all subscribers using that interface.

[See [CoS for Aggregated Ethernet Subscriber Interfaces Overview](#).]

- **Out-of-address SNMP trap requires thresholds to be configured (MX Series)**—Starting in Junos OS Release 19.2R1, the behavior has changed for generating an out-of-address SNMP trap for an address pool. You must now configure both the high-utilization and abated-utilization thresholds. When the number of assigned addresses surpasses the high-utilization threshold, a high-utilization trap is generated. If all the addresses are assigned from the pool, an out-of-address trap is generated and an out-of-address syslog message is sent.

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

[See [Configuring Address-Assignment Pool Usage Threshold Traps](#).]

- **juniper-access-line-attributes option replaces juniper-dsl-attributes (MX Series)**—Starting in Junos OS Release 19.2R1, the **juniper-access-line-attributes** option replaces the **juniper-dsl-attributes** option at the **[edit access profile *profile-name* radius options]** hierarchy level. For backward compatibility with existing scripts, the **juniper-dsl-attributes** option redirects to the new **juniper-access-line-attributes** option. We recommend that you use **juniper-access-line-attributes** from now on.

VLAN Infrastructure

- **Specifying a descending VLAN ID range (MX Series routers, and vMX virtual routers)**—In Junos OS releases prior to Junos OS Release 19.2R1, the system accepts a descending range—for example, 102-100, with the **vlan-id-range** configuration statement in the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy.

Starting with Junos OS Release 19.2R1, the system considers a descending range specified with **vlan-id-range** to be invalid and raises an error if you try to commit this configuration.

SEE ALSO

[What's New | 92](#)

[Known Limitations | 123](#)

[Open Issues | 127](#)

[Resolved Issues | 139](#)

[Documentation Updates | 191](#)

[Migration, Upgrade, and Downgrade Instructions | 192](#)

Known Limitations

IN THIS SECTION

- [General Routing | 124](#)
- [Interfaces and Chassis | 126](#)
- [Platform and Infrastructure | 127](#)
- [Routing Protocols | 127](#)

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

- The Routing Engine boots from the secondary disk when you: a) press the reset button, on the RCB front panel, while the Routing Engine is booting up but before Junos OS is up. b) Upgrade software, by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network. c) Upgrade BIOS and the upgrade fails. d) Reboot and the system hangs before Junos OS is up. [PR1344342](#)
- During a unified ISSU that warrants host upgrade, if the router is configured with 8 million IPv4/IPv6 routes or more, the unified ISSU might fail resulting in FPC restart. [PR1348825](#)
- The commit is successful when the configured MTU value is greater than 9500, which is the maximum permissible value. However, the actual value is set back to 1518B without any error. Check the DCD log to verify the occurrence. [PR1372690](#)
- The MIC-MACSEC-20G supports 10-Gigabit speed through the **set chassis fpc x pic y pic-mode 10G** configuration applied to both the PICs in that MIC. Any other PIC mode configuration should be removed and then the 10-Gigabit PIC mode configuration is to be applied. [PR1374680](#)
- In Junos OS, most daemons underwent architectural change in transition from Junos OS Release 14.1X53 to Junos OS Release 17.X (4 years) and many new features were added. These changes caused an increase in memory footprint in Junos OS Release 17.X compared to Junos OS Release 14.1X53. Unless we see system instability or any adverse performance impact, or a daemon crash due to low memory, this increased memory footprint should not be an issue, and functionality should work fine. The increased memory footprint is a Junos OS property. [PR1390226](#)
- On MX2008 platform with MPC9E, in line rate traffic with a redundant SFB2 scenario, if offline one redundant SFB2, there might be tail or sometimes WRED drops in MPC9E, resulting in partial traffic loss. Under normal circumstances, the SFBs should be auto fail-over if one of them fails, and there should be only a little packet dropped momentarily. [PR1395591](#)
- The dfe tuning failing at times is a known issue on MX10003, the only recovery option in this situation is to restart the FPC. [PR1413233](#)
- The MX104 router has the following limitations in error management: The **show chassis fpc error** 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 to disable the Packet Forwarding Engine when major errors occur, Junos OS does not disable the Packet Forwarding Engine on encountering a major error. [PR1413314](#)
- In Next Gen Services and non-Next Gen Services cases, the monitor interface is MS or VMS. When chassisd restarts, all FPCs are restarted. The service redundancy daemon (srd) also gets restarted and the ICCP connection goes down. If the FPC hosting the ICL goes down first before srd receives the information about the down physical interface for the monitored interface, it will not do switchover immediately. The same behavior is observed in the Next Gen Services and the non-Next Gen Services as well. [PR1416064](#)

- In the following scenario Device 1 Remote Device **MX10003-mx1ru-h** <-----> **MX10003-mx3ru-i et-0/0/2 et-1/0/1**. If PRBS is started on simultaneously as TX and RX on both the devices, there will be errors seen at remote device because when PRBS is started as TX on remote device, it attempts to df tune the line again but PRBS is already running as RX which causes the error. So first start As Tx on Device 1 and as Rx on Remote device, then stop the test on both the ends and start as TX on remote device and as Rx on Device 1. [PR1416124](#)
- Names of user-defined applications are always displayed in the sessions output if they match the traffic criteria defined in the application definition. This happens irrespective of whether the match conditions in the rule has these applications as one of the match condition or not. [PR1416365](#)
- Since creating the loopback at the MacSec port (remote end) in this specific situation, the link itself is down at the EA port hence PRBS test fails with incrementing error counts. [PR1421432](#)
- Due to a race condition between the creation of logical interfaces and sending out of GARP when a logical interface is configured, there is an issue of logical interface statistics incrementing by one output packet. [PR1430431](#)
- FLT will not support source-port and port combination match due to the limitation. [PR1432201](#)
- Dynamic spring-te tunnel creation to LDP (non SR) speaking nodes are not supported even in the presence of mapping server configurations. Spring-te internally converts the tunnel-hop IP addresses (prefix/adjacency) into corresponding labels through auto-translate feature. This feature internally makes use of Traffic Engineering Database (TED); where at present the mapping server entries are not present. [PR1432791](#)
- On MPC2 Junos telemetry interfaces services statistics might not be available after the unified ISSU. [PR1433589](#)
- 128k source-ip addresses as match condition should be configured under couple terms. After commit the configuration, it will take 10 minutes to effect. [PR1433974](#)
- On MX10003 platform with no MSATA device, xSTP topology change is seen during FRU upgrade state in unified ISSU. [PR1435397](#)
- When the Junos telemetry interface collector runs for a longer duration, the iLatency will be negative. [PR1436126](#)
- With scaled inline single-hop BFD sessions, and events such as restart of FPC, ppm, drpd, and some of the BFD sessions might flap. [PR1436543](#)
- In a large-scale setup (such as large number of routing instances or interfaces), if there are frequent changes in configuration and interface flapping when the rpd is restarted by deactivating and then activating the logical system or restarting routing, the rpd might crash. [PR1438049](#)
- MX Series routers report Routing Engine and FPC policer violation when DDoS violation occurs. [PR1439427](#)
- 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 will be up again. After the timer expires and the primary path is still not up, DTM asks SRTE to remove the tunnel.

Also, if there are multiple paths to reach the tunnel endpoint, bgp routes will resolve over the other route, for example a L-ISIS path. Later even if the primary path comes up, bgp routes will remain resolved over the other secondary route and does not change. No re-resolution is happening because the SRTE tunnel is resolving with more than one indirection (SRTE over MPLS over IS-IS in this case). Because of the whole design of how resolution happens and multiple dependencies, there is no simple fix for this. The same issue is applicable to RSVP tunnels also. The issue is applicable to uncolored tunnels only. [PR1439557](#)

- Interworking between MPC10E and SCBE3 is not supported. [PR1440073](#)
- A privilege escalation vulnerability might occur in devices running Junos OS configured with dual Routing Engines, Virtual Chassis or high-availability cluster might allow a local authenticated low-privileged user with access to the shell to perform unauthorized configuration modification. [PR1441795](#)
- The jinsightd might display constant CPU utilization levels (for example, 5-6 percent) with no health monitor due to the presence of default fault monitoring telemetry sensors (check with the **show agent sensors** command). This is expected and there is no service impact due to this. The utilization level depends on the number of FPCs in the chassis. [PR1451057](#)
- Syslog error message **Failed to complete DFE tuning** is generated. This message has no functional impact and can be ignored. [PR1473280](#)

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after upgrade. This is because of the presence of an old version of `/var/db/cfm.db`. [PR1281073](#)
- In a large-scale subscriber environment, changing aggregated Ethernet member link configuration might generate core files for the two Routing Engines. [PR1375638](#)
- When disabling physical interface with JNP-100G-AOC-xM AOC cables, port LED could turn red or go off depending on vendor. JNP-100G-AOC-xM cables sourced by Finisar will cause port LED to turn red when physical interface is disabled. Cables sourced by Innolight will cause the port LED to turn off in contrary. Transceiver vendor information can be obtained from the **show chassis pic fpc-slot <fpc slot> pic-slot<pic slot>** CLI command. Transceiver vendor field contains 'JUNIPER-FINISAR' for Finisar and 'JUNIPER-INNO' for Innolight. [PR1415958](#)
- Firmware upgrade for nPhi Madison optics is not supported on MX10008/16 platform. [PR1424408](#)

Platform and Infrastructure

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

Routing Protocols

- When 32,000 SR-TE policies are configured at once, during configuration time there might be scheduler slips. [PR1339829](#)
- A mis-design in the route selection code with respect to BGP multiple exit discriminator (MED) grouping is done. This interacts poorly with families such as EVPN and Layer 2 VPN which create their own routes rather than doing rib-leaking similar to Layer 3 VPN. However, this issue might trigger even without those families if routes are made comparable to BGP (for example, preference 170). This is a corner case with racing condition within rpd code. [PR1352697](#)

SEE ALSO

[What's New | 92](#)

[What's Changed | 113](#)

[Open Issues | 127](#)

[Resolved Issues | 139](#)

[Documentation Updates | 191](#)

[Migration, Upgrade, and Downgrade Instructions | 192](#)

Open Issues

IN THIS SECTION

- [EVPN | 128](#)
- [Forwarding and Sampling | 128](#)
- [General Routing | 129](#)
- [Infrastructure | 134](#)
- [Interfaces and Chassis | 134](#)
- [Junos Fusion for Provider Edge | 135](#)
- [Layer 2 Features | 135](#)

- Layer 2 Ethernet Services | 135
- MPLS | 135
- Network Management and Monitoring | 136
- Platform and Infrastructure | 136
- Routing Protocols | 137
- User Interface and Configuration | 138
- VPNs | 138

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.

EVPN

- Duplicate packets in EVPN scenario are seen because a nondesignated forwarder is sending an inclusive multicast packet to the PE-CE interface after MAC lookup. [PR1245316](#)
- In an EVPN-VXLAN core isolation scenario, the server is multihomed to the leaf devices through LACP interfaces. If graceful restart is enabled, when you reboot the system or restart routing on the leaf device, the core isolation does not work. If you reboot the system, the issue results in the leaf device silently dropping the traffic sent from the server during the time window between LACP coming up and BGP coming up. If you restart routing, there might be no traffic drop because of the graceful restart. [PR1461795](#)
- In an EVPN-MPLS scenario with the **proxy-macip-advertisement** enabled on an IRB, ARP for remote CE device on local PE device might fail and forwarding-table entry always remains in hold state. The **proxy-macip-advertisement** enables the proxy advertisement feature on a device that can function as a Layer 3 gateway. With this feature enabled, the Layer 3 gateway advertises the MAC and IP routes (MAC+IP type 2 routes) on behalf of Layer 2 VXLAN gateways. [PR1506343](#)

Forwarding and Sampling

- The **skip-service** configuration does not work with IPv6 NDP negotiation or ping. [PR1074853](#)
- If IPv4 prefix is added to a prefix list referred by IPv6 firewall filter, then the log message **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized** is not seen. [PR1395923](#)
- Restart the firewall process in both Routing Engines when firewall error logs are noticed along with SSD hardware failure logs. [PR1397171](#)

- In the case of a physical interface policer for **ip-option traffic**, the traffic rate is found to be more than 10 percent. [PR1398728](#)
- Verify event CP down/up is long enough to trigger EP timeout for CoS hierarchy model 2, failing as expected when DHCP subscribers are not bound. [PR1505409](#)

General Routing

- SIP session fails when the IPv4 SIP client in a public network initiates a SIP call with the IPv6 SIP client in a private network. [PR1139008](#)
- If a VM host snapshot is taken on an alternate disk and there is no further VM host software image upgrade, the expectation is that if the current VM host image gets corrupted, the system boots from the alternate disk so the user can recover the primary disk to restore the state. However, if the host root file system is corrupted, the node boots with the previous VM host software instead of booting from the alternate disk. [PR1281554](#)
- The deletion of the **oneset/leaf-list** configuration through JSON might not be successful when the **delete** attribute is passed in the JSON string. [PR1287342](#)
- The **chain-composite** statement does not bring in a lot of gain because TCNH is based on an ingress rewrite premise. Without this statement things work fine. [PR1318984](#)
- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections. The reactions to failure situations might not be handled gracefully. This results 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 infrastructure in the MX Series Virtual Chassis setup. [PR1332765](#)
- The first packet pertaining to the J-Flow Packet Forwarding Engine sensor in UDP mode is missing after the line card reboots on an MX150 platform. [PR1344755](#)
- With GRES enabled in a subscriber environment, if subscribers are logging in or logging out quickly, the service sessions in the session database (SDB) of the backup Routing Engine might leak. If the problem is not detected for a long time, the backup Routing Engine might not be able to synchronize with the master Routing Engine and will not be ready for GRES. [PR1346300](#)
- Backup Routing Engine might crash after more than 10 continuous GRES switchovers. [PR1348806](#)
- For configurations of bridging routing instances with aggregated Ethernet logical interfaces (6400 logical interfaces) and IRB instances, all from a single FPC, the CPU utilization of the FPC stays at 100 percent for 4 minutes. The behavior from PFEMAN of the FPC has the processing time spiked on IF IPCs, and this seems to be the case of MPC7E line cards from Junos OS Release 16.1R1 (or even earlier). After 4 minutes, the CPU utilization comes down and the FPC is normal. Therefore, scaled configurations on MPC7E line card takes settling time of more than 4 minutes. [PR1359286](#)

- In rare circumstances, a faulty SFP transceiver installed in an MX104 might cause the AFEB to go offline. The backup Routing Engine and fan tray might also show an alarm. [PR1360426](#)
- When an FPC is booting up (either during unified ISSU, router reboot, or FPC restart), I2C timeout errors for an SFP transceiver are seen and the I2C action is not completed because the device is busy. When the FPC is up, all the I2C transactions to the device are normal, so no periodic failure is observed. There is no functional impact and these errors can be ignored. [PR1369382](#)
- If any of the log messages continue to appear in the MPC console, it indicates the presence of a faulty SFP/SFP+ transceiver is causing I2C transaction from the main board CPU. There is no software recovery available to recover from this situation. These logs also indicate potential I2C transaction failure with any of the 10 ports available with GMIC2 in PIC 0, resulting in an unexpected behavior. For example, links do not come up or the MIC does not boot when restarted. **I2C Failed device: group 0xa0 address 0x70 Failed to enable PCA9548(0x70):grp(0xa0)->channel(0) mic_sfp_select_link:MIC(0/0) - Failed to enable PCA9548 channel, PCA9548 unit:0, channel ID: 0, SFP link: 0 mic_sfp_id_read: Failed to select link 0.** As a workaround, detect and replace the faulty SFP/SFP+ transceiver plugged into the GMIC2 ports. [PR1375674](#)
- On MX480 routers, a few xe- interfaces are going down with the error message **if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error.** [PR1377840](#)
- The virtio throughput remains the same for multi-queue and single-queue deployments. [PR1389338](#)
- In a BGP Prefix-Independent Convergence (PIC) case, if a route R1 resolves on top of the multipath route R2, where R2 has primary and backup indirect next hops, it will be better if the backup leg is not used for the resolution of R1. There is no impact on any existing CLI commands. The backup path is never used when the primary path is available. [PR1401322](#)
- Parity error detection and correction is not supported. [PR1402455](#)
- On MX150 and vMX-based platforms, when the **clear pim join instance instance-name all** command is issued, it might result in stopping of the riot process on the system. [PR1409527](#)
- A small number of tunneled subscribers might terminate during the unified ISSU because of momentary loss of IP connectivity between the LAC and LNS devices. [PR1414928](#)
- FPC core files are generated on multiple additions or deletions of hierarchical CoS from pseudowire devices. As a workaround, remove the pseudowire device without changing the hierarchical CoS configuration. [PR1414969](#)
- In Next Gen Services and non-Next Gen Services cases, the monitor interface is MS or VMS. When chassisd restarts, all FPCs are restarted. The service redundancy daemon (srd) also gets restarted and the ICCP connection goes down. If the FPC hosting the ICL goes down first before srd receives the information about the down physical interface for the monitored interface, it will not do switchover immediately. The same behavior is observed in the Next Gen Services and the non-Next Gen Services as well. [PR1416064](#)
- The MX Series Packet Forwarding Engine does not account for the labels pushed on to the packet on the egress Packet Forwarding Engine, while the PTX Series Packet Forwarding Engine does. This results

in a slight difference in the byte count for the same traffic stream across these two platforms. The packet count is still the same across the platforms. Currently, this issue is noticed for uncolored SR-TE policies. [PR1416738](#)

- System ID of an old master Routing Engine is reported by FPCs even after GRES. [PR1417366](#)
- Traffic statistics are not displayed for the hybrid access gateway session and tunnel traffic. [PR1419529](#)
- The ROUTING_LOOP_DETECTED subcode is not generated under the PATHERR_RECV code when a strict path loop is created for LSP event telemetry notifications. [PR1420763](#)
- If the HTTP header enrichment function is used, the traffic throughput decreases when the traffic passes through header enrichment. [PR1420894](#)
- Because the loopback was created at the MACsec port (remote end) in this specific situation, the link itself is down at the EA port. Therefore, the PRBS test fails with incrementing error counts. [PR1421432](#)
- For ALGs with out-to-in sessions, if the data sessions come from an IP address that is different from the IP address available in the control sessions with the NAT rule matching, such ALGs should have the match condition for the **destination-address** as **any** and **not a specific IP**, or you must add all possible IP addresses from where the data sessions for the ALGs can come. [PR1421555](#)
- The Junos OS Releases 19.1 and later support RFC8231 and RFC8281 compliance by default. However, if the controller is not compliant with RFC8231 and RFC8281, a backward compatibility can be configured to fall back to pre-RFC 8231/8281 behavior. [PR1423894](#)
- Due to a race condition between the creation of logical interfaces and sending out of gratuitous ARP when a logical interface is configured, there is an issue of logical interface statistics incrementing by one output packet. [PR1430431](#)
- On MPC10 line card, the error message **failed, Return code: 500** is seen with baseline. [PR1431552](#)
- After restarting the router, changing the anti-spoof status causes tunnel duplication. However, half of those tunnels are up because they have not been cleaned up. [PR1433930](#)
- On MX Series platforms, if the clock frequency is slowly changing on CB0 (slow drift), the clock source for MPC-3D-16XGE-SFPP might not be changed to CB1, which causes interfaces on it to go down and remain in the down state. [PR1433948](#)
- When you reboot or power off the backup Routing Engine, a trap message is reported. This is the generic design for the MX10003 platform. [PR1436212](#)
- 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](#)
- With scaled inline single-hop BFD sessions, and events such as restart of FPC, pppmd, rpd, and some of the BFD sessions might flap. [PR1436543](#)
- FPC might crash when the Packet Forwarding Engine memory usage for a partition such as NH/DFW is high. Under low Packet Forwarding Engine memory condition, log message **Safety Pool below 25% Contig Free Space** or **"Safety Pool below 50% Contig Free Space"** might be observed. [PR1439012](#)

- The interface-specific filters do not work on the MPC10E line card, and both count and policer actions are affected. It is advisable not to use interface-specific firewall filters in this release. [PR1439327](#)
- MX Series routers report Routing Engine and FPC policer violation when DDoS violation occurs. [PR1439427](#)
- Before switching mastership of Routing Engine, you need to wait at least 4 minutes after enabling the GRES configuration for both the Routing Engines to come up in dual Routing Engine mode. Check GRES readiness by executing the **request chassis routing-engine master switch check** command from the master Routing Engine and the **show system switchover** command from the backup Routing Engine. [PR1439884](#)
- There is a change in the way egress topology is being set up for the control packets in MPC10 line cards from the way it is set up in legacy MX Series routers. In legacy MX Series routers, the control packets (ARP) are not subject to family **any** firewall next hops, whereas in MPC10 line card they will be. Thus, if the firewall does not have the **ACCEPT** default term, it is expected to drop the ARP packet. [PR1440792](#)
- On MX204 and MX10003 routers, egress stream flush failure and silent dropping of traffic could occur in a rare occasion for a repeatedly flapping link on MPC7E, MPC8E, and MPC9E line cards. [PR1441816](#)
- The BGP session establishing over the GRE tunnel fails when the router receives the BGP packets encapsulated as GRE and uses the firewall filter action to de-encapsulate the GRE header. [PR1443238](#)
- When an xe- interface working in 1-Gigabit mode is added to a member link of an aggregated Ethernet interface, the speed of the aggregated Ethernet interface is incorrectly shown as 10 Gbps. There is no functional impact. This is a display issue. [PR1449887](#)
- On MX Series, the dropped packets are seen on MQ/XM-based MPCs, although there is no traffic flowing through the system. [PR1451958](#)
- When you use the **replace pattern** command to replace the name in the apply-group, the mgd crashes. [PR1452136](#)
- SDK upgrade from 5.10 to 6.5.16 is done in this release. [PR1454144](#)
- When you edit a command and run the command from CLI command history, the timestamp might not appear. [PR1454387](#)
- The firewall filter might be incorrectly updated in the MPC10E Packet Forwarding Engine when a change (for example, add, delete, deactivate, or activate) of firewall filter terms occurs in some scenarios, such as large-scale term changes or changes happening during MPC reboot. The incorrect firewall filter might cause the traffic to be silently dropped or discarded and even lead to an MPC crash. It is a timing issue. [PR1458499](#)
- The commit script does not apply changes in private mode unless a **commit full** operation is performed. [PR1465171](#)
- With BGP RIB sharding and update threading, traffic drops by 100 percent in the BGP Layer 3 VPN streams after removal or restoration of the configuration. [PR1469873](#)

- When enhanced-mode ISSU is in progress, the MPC7E, MPC8E, and MPC9E line cards lose control connection, chassisd incorrectly marks FPCs offline reason as "Bad Voltage". [PR1473722](#)
- When you reboot the external server, the SNMP values configured within `/etc/snmp/snmpd.conf` at the server get overwritten with content from the JDM SNMP stanza. Trap configuration changes get completely removed. JDM restart or stop/start does not change the `host /etc/snmp/snmpd.conf` file; only system reboot of the server does this change. [PR1474349](#)
- Changing framing modes on a CHE1T1 MIC between E1 and T1 on an MPC3E NG HQoS line card causes the PIC to go offline. [PR1474449](#)
- In VPLS configurations, ARP resolution over an IRB interface might fail if the hosts are behind a vt-tunnel. As a workaround, you can use **no-tunnel-services** statement. [PR1477005](#)
- When specific hardware failure conditions occur in MX2000 platforms, fabric healing attempts to auto-heal the fault location in three phases to prevent traffic get silently dropped and discarded. In such fault conditions, fabric healing process in last phase-3 might not be able to decide which FPC slot should be marked faulty and off-line all the FPCs in the system reporting fabric destination time out. [PR1482124](#)
- BFD over Layer 2 VPN or Layer 2 circuit does not work because of the SDK upgrade to version 6.5.16. [PR1483014](#)
- Viewing a large file from vFPC console using cat application might hog the console. [PR1493805](#)
- On deactivating and activating routing instance, packets from nonexisting source on GRE or UDP designated tunnel are accepted where they are supposed to be dropped. [PR1503421](#)
- On all Junos OS platforms with the Juniper Telemetry Interface configured, the rpd might crash when there is telemetry streaming is in progress and meanwhile there is a network churn. This is a timing issue, and the rpd recovers automatically. [PR1505425](#)
- In an EVPN scenario with VRRPv6 is used, the Ethernet source MAC address might be used for IPv6 mac-ip binding when the NA is sent from VRRPv6 master. As this unexpected behavior is triggered on regular intervals, it causes the entries to keep refreshing in the EVPN database because NS from VRRPv6 master changes the mac-ip binding. This impacts the traffic. [PR1505976](#)
- A 10-Gigabit Ethernet interface configured with WAN-PHY framing might flap continuously if the **hold-down** timer is set to 0 (which is the default). This is not applicable to an interface with the default framing LAN-PHY. [PR1508794](#)

Infrastructure

- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory).** [PR1315605](#)

Interfaces and Chassis

- Out-of-sequence packets are seen with the LSQ interface. [PR1258258](#)
- With a connectivity fault management (CFM) configuration, if you execute an upgrade between releases that uses a different db format, the continuous cfmd crashes might be seen after upgrade. [PR1281073](#)
- After GRES, 1-Gigabit Ethernet speed calculation changes to 10-Gigabit Ethernet. [PR1326316](#)
- In MX Series Virtual Chassis, flooding of the error message **CHASSISD_CONFIG_ACCESS_ERROR: pic_parse_ifname: Check fpc rname failed** can be seen when MPC7, MPC8, and MPC9 line cards are inserted to member 1. The errors only impact DWDM PICs, and do not affect the MPC7, MPC8, and MPC9 line cards. Hence this syslog message can be safely suppressed. [PR1349277](#)
- There are multiple failures when events such as node reboot, ICL flaps, and ICCP flaps occur, and even with **enhanced convergence** configured there is no guarantee that subsecond convergence will be achieved. [PR1371493](#)
- Some routers index the SFP transceivers starting at 1, while interface numbering starts from 0; thus, reading the Packet Forwarding Engine-level output can be confusing. [PR1412040](#)
- Syslog error **scchassisd[]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x** is observed after MX Virtual Chassis local or global switchover. This syslog is harmless in this scenario. [PR1428254](#)
- If an aggregated Ethernet interface has VRRP configuration, in the following use cases, member logical interfaces are not created after a member physical interface comes up and the aggregated Ethernet interface is in down state:
 - FPC restart (**request chassis fpc restart slot <>**)
 - Chassis-control restart (**restart chassis-control**)
 - Reboot of both Routing Engines (**request system reboot both-routing-engines**).

So, before performing these operations, we recommend that you remove the VRRP configuration from the aggregated Ethernet interface. [PR1429045](#)

Junos Fusion for Provider Edge

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

Layer 2 Features

- If VLAN ID lists are configured under a single physical interface, Q-in-Q might stop working for certain VLAN ID lists. [PR1395312](#)

Layer 2 Ethernet Services

- DHCPV6_LEASEQUERY counter might not be as expected in the **show dhcpv6 relay statistics** command output. [PR1521227](#)

MPLS

- With nonstop active routing (NSR), when the routing protocol process (rpd) restarts on the master Routing Engine, the rpd on the backup Routing Engine might restart. [PR1282369](#)
- The 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](#)
- The default behavior of local reversion has changed from Junos OS Release 16.1 and that impacts the LSPs for which the ingress router does not perform make-before-break. Junos OS does not perform make-before-break for no-CSPF LSPs. [PR1401800](#)
- Packet drop might be seen if SR-TE and mapping server are configured. Dynamic SPRING-TE tunnel creation to LDP (non-SR) speaking nodes is not supported even in the presence of mapping-server configurations. SPRING-TE internally converts the tunnel hop IP addresses (prefix/adjacency) into corresponding labels through the auto-translate feature. This feature internally makes use of traffic engineering database where currently the mapping-server entries are not present. [PR1432791](#)
- When an interface in an MVPN routing instance is changed from a virtual tunnel (VT) interface to a label-switched interface (LSI), the P2MP LSP might get stuck in an incorrect state due to no-tear-down message created from the LSP egress side. In the end, MVPN traffic will be lost. [PR1454987](#)

Network Management and Monitoring

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

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log the error as **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- An accuracy issue occurs with three-color policers of both types, single rate and two rate, in which the policer rate and burst-size combination of the policer accuracy vary. This issue is present starting in Junos OS Release 11.4 on all platforms that use the MX Series ASIC. [PR1307882](#)
- If scaling logical-interface-set members and aggregated Ethernet members are configured on the same FPC, the FPC might crash when it restarts. [PR1380527](#)
- On MX Series routers with MPCs, the 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](#)
- On MX Series routers with MS-MPC cards, when an FPC restarts or the routing-instance type is changed (for example, virtual-router to VRF), or route distinguisher is changed, traffic from a group virtual private network (GVPN) tunnel to MPLS over UDP tunnel might fail to get decrypted on the MS-MPC, and this causes complete service loss. [PR1422242](#)
- On all platforms running Junos OS, with NSR enabled, the BGP session with a hold time of 6 seconds or smaller flaps after if the backup Routing Engine is powered off ungracefully. [PR1428518](#)
- For the bridge domains configured under an EVPN instance, ARP suppression is enabled by default. This enables the EVPN to proxy the ARP, and reduces the flooding of ARP in the EVPN networks. Because of that, storm control is not taking effect to the ARP packets on the ports under such bridge domains. [PR1438326](#)
- A dual Routing Engine Juniper node slicing GNF with no GRES configured and with **system internet-options no-tcp-reset drop-all-tcp** configured might enter dual backup Routing Engine state after a manual GNF Routing Engine mastership switchover attempt with the **request chassis routing-engine master [acquire|release|switch]** CLI command from either GNF Routing Engine CLI. [PR1456565](#)
- Expected PIM joins are not learned after performing GRES. [PR1457166](#)
- In NTP with the boot-server scenario, when the router or switch boots, the NTP daemon sends an ntpdate request to poll the configured NTP boot-server to determine the local date and time. If the ntpdate is not activated correctly while the device is booting, the ntpdate might not work successfully.

Then, some cosmetic error messages of time synchronization might be seen, but there is no impact with time update because the NTP daemon will update the time eventually. [PR1463622](#)

Routing Protocols

- When interoperating with other vendors in a draft-rosen multicast VPN, by default the Junos OS attaches a route target to multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities will be excluded from propagating if the BGP route target filtering is enabled on a device running Junos OS. Note that **draft-rosen-idr-rtc-no-rt** has been created in IETF to document this issue and carry the proposed fix through standards. [PR993870](#)
- Certain BGP traceoption flags (for example, open, update, and keepalive) might result in (trace) logging of debugging messages that do not fall within the specified traceoption category, which results in some unwanted BGP debug messages being logged to the BGP traceoption file. [PR1252294](#)
- LDP and OSPF are in the in-sync state because an IGP interface is down with LDP synchronization enabled for OSPF. According to the current analysis, the **IGP interface down** message is seen as the reason because although LDP notified OSPF that LDP synchronization is achieved, OSPF is not able to take note of the LDP synchronization notification because the OSPF neighbor is not up yet. [PR1256434](#)
- In rare cases, RIP replication might fail as a result of performing NSR Routing Engine switchovers when the system is not NSR ready. [PR1310149](#)
- BGP I/O threading was added in Junos OS Release 16.1R1, whereby BGP writes were batched to improve efficiency. This might sometimes lead to some latency in sending BGP update while reacting to certain network events. [PR1332301](#)
- SCP command with routing option (-JU) is not supported. [PR1364825](#)
- It is possible for a GNF with rosen6 multicast to display stuck KRT queue entries after recovery from a dual Routing Engine reboot at the BSYS. [PR1367849](#)
- Performance improvement with **addpath-optimization** statement configured will vary across releases because of variability of baseline convergences without the configuration statement. [PR1395684](#)
- During NSR initial state replication on a scaled setup, while BGP state replication is still ongoing, the BGP task replication might get marked as completed. This is because BGP replication is triggered and controlled through the backup Routing Engine. You must check the output of the **show bgp replication** command to confirm whether replication has actually completed. This corner case scenario is valid only on a scaled setup and during initial state synchronization. [PR1404470](#)
- The issue is seen when a direct change of route distinguisher is done on a routing instance. As a workaround, deactivate the instance before changing route distinguisher and then reactivate. [PR1433913](#)
- On the MX2010 Series routers, the BFD session on the IS-IS step up flaps during the ISSU - FRU upgrade stage. [PR1453705](#)

- When the system is in transient state (that is, it is learning new routes), the number of routes and the active routes in inet.0 and junosmain::inet.0 are not supposed to be same. When the system is in stable state, the number of routes and the active routes in inet.0 and junos-main::inet.0 converge eventually. [PR1453981](#)
- When the system is in transient state (that is, it is learning new routes), the number of routes and the active routes in inet.0 and junosmain::inet.0 are not supposed to be same. When the system is in stable state, the number of routes and the active routes in inet.0 and junos-main::inet.0 converge eventually. [PR1453981](#)
- The order of the statement displayed for the **show configuration** command is changed. [PR1457240](#)

User Interface and Configuration

- Changing nested apply groups does not take effect. [PR1427962](#)

VPNs

- In an MVPN environment with the SPT-only option, if the source or receiver is connected directly to the c-rp PE device and the MVPN data packets arrive at the c-rp PE device before its transition to SPT, the MVPN data packets might be dropped. [PR1223434](#)
- The LSP might stay down if you configure both the virtual tunnel (VT) interface and vrf-table-label in an MVPN scenario. In this case, VT is preferred over LSI. Later when the VT interfaces are deleted, there is no notification to MVPN indicating that LSI is still available. Hence traffic loss might be seen. [PR1474830](#)

SEE ALSO

[What's New | 92](#)

[What's Changed | 113](#)

[Known Limitations | 123](#)

[Resolved Issues | 139](#)

[Documentation Updates | 191](#)

[Migration, Upgrade, and Downgrade Instructions | 192](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3 | 139](#)
- [Resolved Issues: 19.2R2 | 146](#)
- [Resolved Issues: 19.2R1 | 175](#)

This section lists the issues fixed in Junos OS main and maintenance releases for MX Series routers.

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

Resolved Issues: 19.2R3

Class of Service (CoS)

- MX Series platforms with MPC1-Q and MPC2-Q line cards might report memory errors. [PR1500250](#)

EVPN

- Remote MAC address present in EVPN database might be unreachable. [PR1477140](#)
- The ESI of IRB interface does not update after autonomous-system number change if the interface is down. [PR1482790](#)
- Dead next hops might flood in a rare scenario after remote PE devices are bounced. [PR1484296](#)
- The rpd core file might be generated when doing Routing Engine switchover after disabling BGP protocol globally. [PR1490953](#)
- VXLAN bridge domain might lose VTEP logical interface after restarting chassisd. [PR1495098](#)
- The l2ald memory leak might be observed in an EVPN scenario. [PR1498023](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)
- The MAC address of the LT interface might not be installed in the EVPN database. [PR1503657](#)

Forwarding and Sampling

- Issue of UTC timestamp in the flat-file-accounting files when profile is configured. [PR1509467](#)

General Routing

- The cosmetic error **mshpmand[190]: msvcs_session_send: Plugin id 3 not present in the svc chain for session** is observed in the output. [PR1258970](#)
- The **show security group-vpn member ipsec security-associations detail | display xml** command output is not in the expected format. [PR1349963](#)
- On MX2010 routers, error message **Failed to get xfchip which_xfchip 7, CMTY_FM_HSL2: [cmt_y_fm_hsl2_adc_xbar_set_cell_limit_per_link] xbar limit set failed pb_slot[3], pfe[1], sfb_slot[7], sub_plane[2], link_index[83]** is observed during unified ISSU. [PR1388076](#)
- The configuration database might not be unlocked automatically if the related user session is disconnected when the commit operation in progress. [PR1410322](#)
- Core files are generated with reference to **ddl_access_check_sequence** during software upgrade. [PR1414118](#)
- **FPC x Voltage Tolerance Exceeded** alarm is raised and cleared upon bootup of JNP10K-LC2101. [PR1415671](#)
- Resetting the playback engine log messages are seen on MPC5E line card. [PR1420335](#)
- FPC crash might be observed after GRES when you commit changes in the firewall filter with the **next term** statement in subscriber scenario. [PR1421541](#)
- Configuration commit might fail when the file system gets into full state. [PR1423500](#)
- IRB over VTEP unicast traffic might get dropped on MX Series platforms. [PR1436924](#)
- FPC19 (MPC9E) might not go offline due to unreachable destinations in phase 3 stage. [PR1443803](#)
- Interfaces shut down by **disable-pfe** action might not be up using MIC offline or online command. [PR1453433](#)
- FPC or Packet Forwarding Engine might crash with ATM MIC installed in the FPC. [PR1453893](#)
- Multiple daemons might crash when you commit configuration changes related to groups. [PR1455960](#)
- 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 MPC2E-NG or MPC3E-NG line card with specific MIC might crash after a high rate of interface flap. [PR1463859](#)
- Daemons might not be started if **commit** is executed after **commit check**. [PR1468119](#)
- IPv6 dynamic subscribers might be unable to access on Junos OS Release 18.2R3 and later. [PR1468414](#)
- NSD core file generates after committing the configuration successfully if the destination NAT rule matches the destination address. [PR1469613](#)

- When both MSTP and ERP are enabled on the same interface, then ERP does not come up properly. [PR1473610](#)
- Routing Engine might not be able to send packets after some time when traffic-manager enhanced-priority-mode configuration is enabled on MPC10 or MPC11 line cards. [PR1476683](#)
- Error log message **chassisd[7836]: %DAEMON-3-CHASSISD_IOCTL_FAILURE: acb_get_fpga_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device)** is observed after every commit. [PR1477941](#)
- MPCs might crash when there is bulk route update failure in a corner case. [PR1478392](#)
- FPC with **vpn-localization vpn-core-facing-only** configuration might get stuck in ready state. [PR1478523](#)
- PPPoE subscribers might not log in after FPC restart. [PR1479099](#)
- The downstream IPv4 packets greater than BR MTU are getting dropped in MAP-E. .
- Interface input error counters are not increasing on MX150 routers. [PR1485706](#)
- Kernel core files are generated if you delete an ifstate. [PR1486161](#)
- The unified ISSU is not supported on next-generation MPC cards. [PR1491337](#)
- MS-MIC is down after loading some releases in MX Virtual Chassis scenario. [PR1491628](#)
- FPCs might stay down or restart when you swap the MPC7, MPC8, and MPC9 line cards with the MPC10 and MPC11 line cards or vice versa in the same slot. [PR1491968](#)
- User-configured MTU might be ignored after the unified ISSU upgrade uses **request vmhost software in-service-upgrade**. [PR1491970](#)
- On MX10008 platform, SNMP table entPhysicalTable does not match the PICs shown for the **show chassis hardware** command. [PR1492996](#)
- MPC10 and MPC11 line cards might crash if the interface is configured with firewall filter referencing shared bandwidth policer. [PR1493084](#)
- Flapping of one port on MIC-3D-2XGE-XFP might cause the other port on the same MIC to flap. [PR1493699](#)
- In node slicing setup after GRES, RADIUS interim updates might not carry actual statistics. [PR1494637](#)
- B4 devices might not be able to establish software with an AFTR device. [PR1496211](#)
- The MPC10E line card might restart with sensord crash on it due to a timing issue. [PR1497343](#)
- Outbound SSH connection flap or a memory leak issue might be observed during the high rate of pushing the configuration to the ephemeral DB. [PR1497575](#)
- In a node slicing scenario, when a FPC with aggregated Ethernet member port comes online, subscribers might be disconnected. [PR1498024](#)
- SNMP polling does not show correct PSM jnxOperatingState when one of the PSM inputs failed. [PR1498538](#)

- The rpd might crash when multiple VRFs with IFLs link-protection are deleted at a single time. [PR1498992](#)
- On MX2020 and MX2010, the **pem_tiny_power_remaining** message will be continuously logged in chassisd log. [PR1501108](#)
- Traffic might be silently dropped and discarded in a fast reroute scenario. [PR1501817](#)
- The chassisd process might get stuck. [PR1502118](#)
- Core files are generated if you add or delete ERP configuration multiple times and restart l2cpd. [PR1505710](#)
- The disabled QSPF transceiver might fail to turn on. [PR1510994](#)
- Static subscribers logs out after creating a unit under demux0 interface. [PR1511745](#)
- The l2ald memory leak might be seen after you add or delete VXLAN routing instances and interfaces. [PR1512802](#)
- The wavelength configured through CLI might not be set on SFP+-10G-T-DWDM-ZR optics when the optics is placed on MPC7E line card. [PR1513321](#)

Infrastructure

- An alarm might be seen if the PEM's serial number starts with 1F1. [PR1398128](#)

Interfaces and Chassis

- The **vrrp-inherit-from** change operation leads to packet loss when traffic is forwarded to the VIP gateway. [PR1489425](#)
- Unexpected dual VRRP backup state might happen after performing two subsequent Routing Engine switchovers with the **track priority-hold-time** statement is configured. [PR1506747](#)

J-Web

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

Junos Fusion Provider Edge

- The stats of extended ports on satellite device cluster might show wrong values from the aggregation device. [PR1490101](#)

Junos Fusion Satellite Software

- Temperature sensor alarm is seen in Junos fusion scenario. [PR1466324](#)

Layer 2 Features

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

Layer 2 Ethernet Services

- On MX204 platform, the Vendor-ID is set as MX10001 in factory-default configuration and DHCP client messages. [PR1488771](#)
- Issues with DHCPv6 relay processing confirm and reply packets. [PR1496220](#)
- The MC-LAG might become down after disabling and then enabling the **force-up** statement. [PR1500758](#)
- DHCPV6 leasequery is not as expected while verifying the DHCPV6 server statistics. [PR1506418](#)

MPLS

- Traffic loss might be seen if p2mp with NSR enabled. [PR1434522](#)
- The RSVP interface bandwidth calculation might be incorrect when the RSVP subscription percentage is configured under the RSVP interface. [PR1458527](#)
- The rpd might crash in PCEP for the RSVP-TE scenario. [PR1467278](#)
- The fast reroute detour next-hop down event might cause primary LSP go in the **Down** state in a particular scenario. [PR1469567](#)
- High CPU utilization for rpd might be seen if RSVP is implemented. [PR1490163](#)
- The rpd might crash when BGP with FEC 129 VPWS enabled flaps. [PR1490952](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. [PR1493431](#)
- The rpd process might crash and generate a core file when SNMP polling is done using OID jnxMplsTeP2mpTunnelDestTable. [PR1497641](#)
- Traffic loss might happen after a switchover when RSVP point-to-multipoint is configured. [PR1500615](#)
- CSPF job might get stalled for new or existing LSP in a high scale LSP setup. [PR1502993](#)
- The rpd process might crash with the RSVP configured in a rare timing case. [PR1505834](#)

- Activating or deactivating LDP-sync under OSPF might cause the LDP neighborship to go down and stay down. [PR1509578](#)
- The rpd might crash after upgrading from pre Junos OS Release 18.1 to Junos OS Release 18.1 and later. [PR1517018](#)

Platform and Infrastructure

- Core.vmx.mpc0 is seen at 5 0x096327d5 in l2alm_sync_entry_in_pfes (context=0xd92e7b28, sync_info=0xd92e7a78) at `../../../../src/pfe/common/applications/l2alm/l2alm_common_hw_api.c:1727`. [PR1430440](#)
- Traffic loss might be seen in case of Ethernet frame padding with VLAN. [PR1452261](#)
- The **show system buffer** command displays all zeros in the MX104 chassis. [PR1484689](#)
- MAC learning under bridge domain stops after MC-LAG interface flaps. [PR1488251](#)
- In node slicing setup MPLS TTL might be set to zero when the packet goes through af interface configured with CCC family. [PR1492639](#)
- Packets get dropped when next hop is IRB over lt- interface. [PR1494594](#)
- Routing Engine crash might be seen when a large number of next hops are quickly deleted and readded in large ARP/ND scale scenario. [PR1496429](#)
- Python or Slax script might not be executed. [PR1501746](#)
- Traffic originated from another subnet is sent out with 0x8100 instead of 0x88a8. [PR1502867](#)
- MPCs might crash when there is a change on routes learned on IRB interface configured in VPLS and EVPN instances. [PR1503947](#)
- Traffic loss might be seen in certain conditions under MC-LAG setup. [PR1505465](#)
- During route table object fetch failure, FPC might crash. [PR1513509](#)

Routing Protocols

- PIM RPF selection for the specific multicast group might get incorrectly applied to other multicast groups. [PR1443056](#)
- The rpd crash might be observed because of specific BGP update packets. [PR1481641](#)
- Multicast traffic loss might be seen in certain conditions while enabling IGMP snooping under EVPN-VXLAN ERB scenario. [PR1481987](#)
- RIPv2 packets stop transmitting when changing interface-type configuration from p2mp to broadcast. [PR1483181](#)
- Rpd memory leak is observed in certain looped MSDP scenario. [PR1485206](#)
- The rpd process might generate soft core files after **always-compare-med** is configured for BGP path-selection. [PR1487893](#)

- The rpd crashes when reset OSPF neighbors. [PR1489637](#)
- The BGP route target family might prevent route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- The rpd might crash because of rpd resolver problem of INH. [PR1494005](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)
- The rpd might crash if the import policy is changed to accept more routes that exceed the teardown function threshold. [PR1499977](#)
- The **show bgp neighbors** command shows change in x-path output for **input-updates** value. [PR1504399](#)
- The rpd crash might occur due to RIP updates being sent on an interface in down state. [PR1508814](#)
- The rpd process might crash on backup Routing Engine if BGP (standby) received a route from peer which is rejected due to invalid target community. [PR1508888](#)
- The rpd process might crash if there is a huge number of SA messages in MSDP scenario. [PR1517910](#)

Services Applications

- The FPC might crash and generate NPC core file if the service interface is configured under **service-set** in USF mode. [PR1502527](#)

Subscriber Access Management

- Syslog message **pfe_tcp_listener_open_timeout: Peer info msg not received from addr: 0x6000080. Socket 0xfffff804ad23c2e0 closed** is seen. [PR1474687](#)
- The authd logs events might not be sent to syslog host when **destination-override** is used. [PR1489339](#)
- MX Series platforms are not compliant with RFC 2868 and sending RADIUS access request includes tunnel assignment ID for LTS client. [PR1502274](#)

User Interface and Configuration

- The version information under the configuration changes from Junos OS Release 19.1. [PR1457602](#)

VPNs

- Traffic loss might be observed when the inter-AS next-generation MVPN VRF is disabled on one of the ASBRs. [PR1460480](#)
- The l2circuit neighbor might be stuck in RD state at one end of MG-LAG peer. [PR1498040](#)
- The rpd fork process might crash on removing last l2circuit neighbor interface. [PR1511783](#)
- The rpd might crash when you delete l2circuit configuration in a specific sequence. [PR1512834](#)

Resolved Issues: 19.2R2

Application Layer Gateways (ALGs)

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

Authentication and Access Control

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

Class of Service (CoS)

- The host-inbound packets might be dropped if you configure **host-outbound** FC. [PR1428144](#)
- The dfwd crashes for the **forwarding-class** configuration in policers. [PR1436894](#)
- MX Series generated OAM/CFM LTR messages are sent with a different priority than the incoming OAM/CFM LTM messages. [PR1466473](#)
- Unexpected traffic loss might be discovered under certain conditions in a Junos fusion scenario. [PR1472083](#)
- The MX10008 and MX100016 routers might generate cosd core file after executing **commit/commit check** command if **policy-map** configuration is set. [PR1475508](#)

EVPN

- The RA packets might be sent out without using the configured virtual gateway address. [PR1384574](#)
- Unexpected next-hop operation error from kernel to l2ald in a Layer 2 gateway during the MAC movement operation is observed. [PR1430764](#)
- Incorrect MAC count is observed with the **show evpn/bridge statistics** command. [PR1432293](#)
- Asynchronous state between ARP table and Ethernet switching table occurs if EVPN ESI link flaps multiple times. [PR1435306](#)
- EVPN/MPLS IRB logical interface might not come up when local Layer 2 interface is down. [PR1436207](#)
- Configuring ESI on a single-homed 25-Gigabit Ethernet port might not work. [PR1438227](#)
- On EVPN/VXLAN setup, the specific source-ports of UDP packet are dropped. [PR1441047](#)
- Restarting Layer 2 learning might cause some remote MAC addresses to move into forwarding dead state. [PR1441565](#)
- Traffic drop might be observed in an EVPN Layer 3 gateway scenario. [PR1442319](#)
- On MX Series platform, the core-isolation feature does not work after setting and then deleting the **no-core-isolation** statement. [PR1442973](#)
- The EVPN type 2 routes might not be advertised properly in logical-systems. [PR1443798](#)
- The local host address is not present in the EVPN database and the **mac-ip-table** table. [PR1443933](#)

- The bridge mac-table age timer does not expire for rbeb interfaces. [PR1453203](#)
- Instance type is changed from VPLS to EVPN, which results in loss of packet. [PR1455973](#)
- Preference-based DF Election algorithm does not work on the LT interface. [PR1458056](#)
- ARP request or NS might be sent back to the local segment by the DF router. [PR1459830](#)
- Traffic received from VTEP gets dropped if the VNI value used for Type-5 routes is greater than 65,535. [PR1461860](#)
- The rpd might crash after changing EVPN-related configuration. [PR1467309](#)
- Dead next hops might flood in a rare scenario after remote PE devices bounce. [PR1484296](#)
- The ARP entry gets deleted from the kernel after adding and deleting the virtual-gateway-address. [PR1485377](#)
- The rpd core might be generated when you do Routing Engine switchover after disabling the BGP protocol globally. [PR1490953](#)
- On an MX10003, VTEP interface are not installed under the VXLAN bridge domain after a chassis restart. [PR1495098](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)

Fault Management

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

Forwarding and Sampling

- The SRRD might crash when memory corruption occurs. [PR1414568](#)
- DT_BNG: **rt-delay-threshold** can be set below one second. However, **rt-marker-interval** is limited to one second. [PR1425544](#)
- The device is in amnesiac mode after ISSU with **mgd: error: configuration check-out failed** error generated. [PR1432664](#)
- Enable interface with input or output VLAN maps to be added to a routing instance configured with a VLAN ID or VLAN tags with the virtual-switch instance type and VLPS protocol. [PR1433542](#)
- The **test aaa ppp** command gets timed out with **Client session Activate: no response** error. [PR1435689](#)
- Sampling might return incorrect ASN for BGP traffic. [PR1439630](#)
- High CPU utilization of l2ald is observed after replacing the EVPN configuration. [PR1446568](#)
- On MX204, input/output counters of an aggregated Ethernet bundle or member links configured on nondefault logical systems are not updated. [PR1446762](#)
- The ARP packets get dropped by Packet Forwarding Engine after chassis-control is restarted. [PR1450928](#)
- Crafted packets traversing in a BNG configured with IPv6 NDP proxy could lead to denial-of-service. [PR1451959](#)

- On the PTX Series or TVP platforms, the pfd might crash and might not be able to come up. [PR1452363](#)
- Commit error and dfwd core files might be observed when applying a firewall filter with **then traffic-class** or **then dscp** action. [PR1452435](#)
- On devices running Junos OS, the l2ald process might experience memory leak. [PR1455034](#)
- The following syslog error messages are seen at pfd: **rtplib: ERROR received async message with no handler: 28**. [PR1458008](#)
- A problem with statistics on some interfaces of a router might be observed after FPC or PIC reboot. [PR1458143](#)
- With the MX Series devices, 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](#)
- Type 1 ESI/AD route might not be generated locally on EVPN PE device in the all-active mode. [PR1464778](#)
- On the MX10008 and MX10016 routers, policer bandwidth-limit cannot be set higher than 100-Gigabit Ethernet. [PR1465093](#)
- An output bandwidth-percent policer with **logical-bandwidth-policer** applied to an aggregated Ethernet bundle along with an **output-traffic-control-profile** has incorrect effective policing rate. [PR1466698](#)
- Traffic might not be forwarded into the correct queue instead of the default queue when the VPLS traffic has three or more VLAN tags with VLAN priority 5. [PR1473093](#)
- The filter might not be installed if the **policy-map xx** is present under the filter. [PR1478964](#)

General Routing

- On MX Series Virtual Chassis, suboptimal aggregated Ethernet load balancing occurs when an aggregated Ethernet bundle is part of an ECMP path. [PR1255542](#)
- **PFEIFD: Could not decode media address with length 0** syslog error messages might be generated by the Packet Forwarding Engine. [PR1341610](#)
- Default credentials supplied in vMX configuration. [PR1344858](#)
- SFP stop forwarding traffic after unified ISSU upgrade. [PR1379398](#)
- The severity of the error is reduced from fatal to major [PR1390333](#)
- The **high-cos-queue-threshold** range is changed to [uint 0 .. 90;]. [PR1390424](#)
- Commit error might be observed after adding additional sites to existing group and **routing-instance** configuration. [PR1391668](#)
- Layer 3 gateway does not update ARP entries IP or MAC quickly move from one router to another router in an EVPN-VXLAN environment. [PR1395685](#)
- NAPT66 pool split is not supported with AMS has commit fail with IPv6 pool in AMS. [PR1396634](#)
- The PPPoE subscribers are unable to reconnect after FPC reboot. [PR1397628](#)
- Confirmation message is missing when issuing **request vmhost reboot re***. [PR1397912](#)

- The rpd might crash when **condition-manager** policy is configured for routing table and the same routing table is repeatedly deleted then re-added. [PR1401396](#)
- The na-grpcd log file is not rotated and keeps growing until the Routing Engine is out of disk space. [PR1401817](#)
- Kernel memory leak in virtual memory because of interface flap. [PR1407000](#)
- FPC crash and slow convergence upon HMC fatal error condition is seen when inline J-Flow is used [PR1407506](#)
- For the initial packet, which is specific to MPC10 and onward, the ICMP redirects are not seen at the source and packets are sent to the better next hop. [PR1409346](#)
- Nonexistent subscribers might appear in the output of the **show system resource-monitor subscribers-limit chassis extensive** command. [PR1409767](#)
- Slow SNMP on entityMIB during subscribers load test. [PR1411062](#)
- Parity error might cause FPC alarm. [PR1411610](#)
- Log severity level change is seen on MX150. [PR1411846](#)
- Egress monitored traffic is not mirrored to destination for analyzers on MX Series router. [PR1411871](#)
- Redirect IP is not supported for BGP FlowSpec filters. If such an action is programmed for BGP FlowSpec rules, then it will not be reflected in filter term action. [PR1413371](#)
- J-Flow gets disabled when you reduce the maximum flow table size by using the **flex-flow-sizing** statement enabling the **bandwidth** command. [PR1413513](#)
- On PowerPCbased MX Series platforms, the DHCP/DHCPv6 subscribers might fail to establish sessions. [PR1414333](#)
- The **FPC x Voltage Tolerance Exceeded** alarm is raised and cleared after the JNP10K-LC2101 is powered on. [PR1415671](#)
- 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](#)
- Multiple ANs are created when you configure or unconfigure PSK and last for a longer duration. [PR1418448](#)
- Resetting the Playback Engine generates log files on the MPC5E line card. [PR1420335](#)
- Core voltage of ASIC chip in SIBs is not set as per the required e-fuse value and remains to default value of 0.9V on JNP10008-SF and JNP10016-SF SIBs. [PR1420864](#)
- 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 LNS might fail to forward the traffic on the subscriber access route. [PR1421314](#)

- PTP might not work on MX104 if **phy-timestamping** is enabled. [PR1421811](#)
- After control plane event, a few IPsec tunnels failed to send traffic through the tunnel [PR1421843](#)
- RSI bloat VM host-based log collection. [PR1422354](#)
- Packet Forwarding Engine wedge might be observed after running the **show forwarding-options load-balance** command. [PR1422464](#)
- The XML output might not be hierarchically structured if you issue the **show security group-vpn member ipsec statistics** command. [PR1422496](#)
- Ports might get incorrectly channelized if they are 10-Gigabit Ethernet already and they are channelized to 10-Gigabit Ethernet again. [PR1423496](#)
- The MPC10line card might crash once multiple filters are configured in a scaled environment. [PR1423709](#)
- PTP asymmetry change needs PTP bouncing. [PR1423860](#)
- The system does not reboot or halt as configured when encountering the disk error. [PR1424187](#)
- The rpd keeps crashing after changing the configuration. [PR1424819](#)
- Interface with FEC disabled might flap after Routing Engine mastership switchover. [PR1425211](#)
- The mspmand process might crash and restart with a mspmand core file created after you do a commit change to deactivate and activate service set. [PR1425405](#)
- One hundred percent CPU usage is seen on route monitor of static routes after the client disconnected from prpd server. [PR1425559](#)
- MPC reboot or Routing Engine mastership switchover might occur on MX204 and MX10003. [PR1426120](#)
- The host-bound traffic might be dropped after performing a change configuration related to **prefix-list**. [PR1426539](#)
- Some CFM and BFD sessions might flap while collecting MPLS statistics. [PR1426727](#)
- The **show lldp neighbors interface** command does not display all interface information. [PR1426793](#)
- The decoding of telemetry data at collector might not be proper if you configure the sensors. [PR1426871](#)
- ENTITY MIB has incorrect containedIn values for some fixed MPCs with built-in PICs. [PR1427305](#)
- Rebooting or halting a Virtual Chassis member might cause the RTG link to be down for 30 seconds. [PR1427500](#)
- When broadband edge PPPoE and DHCP subscribers coming up over Junos fusion satellite ports are active, the **commit full** and the **commit synchronization full** commands fail. [PR1427647](#)
- When installing YANG package without **proxy-xml** statement, the CLI environment does not work well. [PR1427726](#)
- The PPP session does not work properly on MX Series platform. [PR1428212](#)
- The **global-mac-limit** and **global-mac-ip-limit** configuration statements might allow more entries than the configured values. [PR1428572](#)

- Fabric drops might be seen on MX10003 platform when two FPCs come online together. [PR1428854](#)
- Incorrect IGMP interface counter for dynamic PPP interfaces. [PR1429018](#)
- The emitted XML output INVALID is thrown for the **show virtual-network-functions** command. [PR1429090](#)
- A race condition vulnerability might cause RPD daemon to crash when processing a BGP notification message. [PR1429719](#)
- Extended ukern thread (PFEBM task) priority to support BBE performance tuning. [PR1429797](#)
- The aggregated Ethernet interface does not come up after rebooting the FPC or device although the physical member link is up. [PR1429917](#)
- **Protect core** configured router might send IPFIX sampling packets with incorrect next-hop information. [PR1430244](#)
- Performance degradation is observed for about 20 seconds after the fabric board on MX10008 or MX100016 is taken offline. [PR1430739](#)
- Disabling DAC QSFP port might not work on MX204 and MX10003. [PR1430921](#)
- Inline LSQ might not work when it is configured on the same FPC where MIC-3D-16CHE1-T1 is slotted. [PR1431069](#)
- Error might be observed when you use a script to load the configuration. [PR1431198](#)
- The l2cpd process might crash and generate a core file when interfaces flap. [PR1431355](#)
- **SIB Link Error** is detected on a specific Packet Forwarding Engine might cause complete service impact. [PR1431592](#)
- Allow installation of three identical framed routes in the same routing instance. [PR1431891](#)
- Line card might be offline when Packet Forwarding Engine is powered off. [PR1432019](#)
- Dual-stack subscriber accounting statistics are not baselined when one stack logs out. [PR1432163](#)
- Traffic might be sent on the standby link of aggregated Ethernet bundle and lost with LACP **fast-failover** enabled. [PR1432449](#)
- Change to in-use parameterized filter prefix-list result in bbe-smgd core files on the backup Routing Engine. [PR1432655](#)
- Output traffic statistics might be incorrect with Routing Engine generated traffic. [PR1432724](#)
- In BBE configurations, receipt of a specific MPLS or IPv6 packet causes a denial of service. [PR1432957](#)
- After deleting the CLI configuration **chassis license bandwidth**, the bandwidth value is not defaulting to maximum bandwidth value. [PR1433157](#)
- A few entries specific to **show dynamic-tunnels database** output are not getting populated while testing the functionality after both PICs are taken offline and then one PIC is brought online. [PR1433247](#)
- Traffic drop **sa-multicast** is configur. [PR1433306](#)

- The gNMI set RPC with replace field does not work and the mgd-api might crash. [PR1433378](#)
- RSI and RSIbrief should not include **show route forwarding-table** when is enabled. [PR1433440](#)
- Junos telemetry interface-firewall ollected service statistics all 0 after unified ISSU for MPC2. [PR1433589](#)
- Lawful intercept for subscriber traffic is not programmed in Packet Forwarding Engine if it activated by Access-Accept. [PR1433911](#)
- URL case sensitivity support is needed. [PR1434004](#)
- Incorrect PLUGGABLE ID 17 on MX10003-LC2103. [PR1434183](#)
- rpd crashes during the route flash when the policy is removed. [PR1434243](#)
- Packet Forwarding Engine memory leak might be seen if MLPPP links are flapped. [PR1434980](#)
- Micro BFD 3x100 ms flap is seen upon inserting a QSFP to another port. [PR1435221](#)
- Traffic drops when session key rolls over between primary and fallback for more than 10 times. [PR1435277](#)
- DHCPv6 advertise to client might use incorrect destination MAC address. [PR1435694](#)
- Total number of packets mirrored after adding the DTCP trigger and enabling DTC is not in the expected range while verifying traffic on mirror port after DTCP drop policy is enabled. [PR1435736](#)
- MPC7/8/9/MX10003 MPC/EX9200-12QS/EX9200-40XS line card might crash in a scaling setup. [PR1435744](#)
- The mc-ae interface might get stuck in waiting state in a dual mc-ae scenario. [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](#)
- Control logical interface is not created by default for LLDP. [PR1436327](#)
- A few static PPP/PPPoE get subscribers stuck in the initialization state permanently and the **Failed to create client session, err=SDB data corrupted** error message is seen. [PR1436350](#)
- The subscriber interim statistics might reset to zero and idle-timeout might not work in the statistics setup. [PR1436419](#)
- The device might not be reachable after a downgrade from some releases. [PR1436832](#)
- On MPC10E, micro BFD sessions do not come up in centralized mode. [PR1436937](#)
- MX10003 FPCs show high CPU because the JGCI_Background thread runs for a long period. [PR1437283](#)
- Schema XSDs are missing objects/commands from Junos OS Release 19.1. [PR1437469](#)
- The CPU utilization on a daemon might be around 100 percent or backup Routing Engine might crash in race conditions. [PR1437762](#)

- The **LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** messages is seen while you commit the configurations. [PR1437824](#)
- LNS router might send the router-advertisement packet with NULL source link-layer option field. [PR1437847](#)
- The chassisd might crash after you enable **hash-key**. [PR1437855](#)
- Unified ISSU fails from 19.1R1 legacy Junos OS release images. [PR1438144](#)
- The rpd crash might be seen during the router startup file pointer issue. [PR1438597](#)
- Subscriber flows might not be synchronized between aggregated Ethernet members on MX Series Virtual Chassis platforms. [PR1438621](#)
- Carrier-grade NAT logs are not received by the syslog server over TCP-based-syslog when data traffic is sent at 10,000 sessionssec. [PR1438928](#)
- Incorrect values are seen in JUNIPER-TIMING-NOTFNS-MIB table. [PR1439025](#)
- The **show jdaf service cmd statistics / clients** command is not available. [PR1439118](#)
- FPC on Virtual Chassis backup router might reboot in an MX Series Virtual Chassis scenario. [PR1439170](#)
- LACP state might get stuck in Attached state after disabling peer active members. [PR1439268](#)
- Packet drop might be seen when chained composite next hop is enabled for Layer 3 VPN. [PR1439317](#)
- The **vlan all interface all** combination not work as expected under VSTP. [PR1439583](#)
- When group is applied at nonroot level, updating commands inside the group does not update the hierarchies where it is applied. [PR1439805](#)
- The bbe-smgd core files are seen after restart. [PR1439905](#)
- PRPD flexible tunnel profile queries do not return DMAC when set to all zeros by client. [PR1439940](#)
- CoS-related errors are seen and subscribers could not get service. [PR1440381](#)
- CPU might hang or interface might be stuck down on a particular 100-Gigabit Ethernet port. [PR1440526](#)
- FPC might stuck 100 percent CPU GRES and multiple daemons continuous restart on MX platforms. [PR1440676](#)
- DHCP offer packets toward IRB over LT interface getting dropped in DHCP relay environment. [PR1440696](#)
- The Layer 2 dynamic VLAN misses when an interface is added to or removed from an aggregated Ethernet interface. [PR1440872](#)
- The EX ports might stay in up state even if the EX4600 line of devices or the QFX5100 line of devices are rebooted. [PR1441035](#)
- For a route received through EBGp, the AIGP value might not be considered as expected. [PR1441438](#)
- The rpd might crash or consume 100 percent of CPU after flapping the routes. [PR1441550](#)

- New OID is added that calculates the buffer utilization where inactive memory is not considered as free memory. [PR1441680](#)
- The outgoing aggregated Ethernet traffic might be dropped after changes are made to the aggregated Ethernet interface. [PR1441772](#)
- Privilege escalation vulnerability in dual Routing Engine, Virtual Chassis, or HA cluster might allow unauthorized configuration change. [PR1441795](#)
- SNMP trap comes twice for FRU removal in MX10000—one trap with FRU name as FPC: JNP10K-LC2101 and second with FRU name as FPC @ 1/*/* . [PR1441857](#)
- The packets originating from the IRB interface might be dropped in a VPLS scenario. [PR1442121](#)
- The chassisd is unable to power off a faulty FPC after Routing Engine switchover, which leads to chassisd restart loop. [PR1442138](#)
- In enhanced-ip or enhanced-ethernet mode with destination-class-usage (DCU) 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. Now the severity of those logs is corrected as INFO. [PR1442598](#)
- The interface might go into the down state after FPC restarts with the PTP configuration enabled. [PR1442665](#)
- DHCPv6 client might fail to get an IP address. [PR1442867](#)
- On MX Series platforms, the bbe-smgd might crash. [PR1443109](#)
- Improper handling of specific IPv6 packets sent by clients eventually leads to kernel crash (vmcore) on the device. [PR1443576](#)
- 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](#)
- GRE packets that are larger than MTU get dropped on MX204 platforms 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](#)
- Access route might be stuck in bbe-smgd and rpd might not be cleared. [PR1445155](#)
- The CPCDD process might crash continuously if the **captive-portal-content-delivery** service is activated for dual-stack PPPoE/DHCPv6 subscriber. [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-Gigabit Ethernet interface on MX204 might stay down after the device is rebooted. [PR1445508](#)
- Irregular traffic drop might be seen when traffic is ingress from MPC3E and egress to MPC10E. [PR1445649](#)
- In Junos OS Release 19.2 the group level use of wildcard <*> is not an available option. [PR1445651](#)
- The l2ald 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](#)
- The jdncpd process might crash after you issue 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](#)
- Upgrade of jfirmware might fail on MX chassis with MX-SCBE3 installed. [PR1446205](#)
- The static route for NAT might never come up if switchover happens with MX Series route service interface that has NAT and GR configuration. [PR1446267](#)
- The rpd process might crash when it is terminated immediately after it has been started. [PR1446320](#)
- Accurate statistics might not include packets forwarded during the last two seconds before subscriber termination. [PR1446546](#)
- On MX Series routers with MPC10 or MPC11 line cards, the incoming packets might get dropped. [PR1446736](#)
- NAT service set in a certain scale might fail to get programmed. [PR1446931](#)
- MX Series-based MPC might crash and restart during unified ISSU with large-scale logical interfaces. [PR1446993](#)
- The J-Flow version 5 stops working after changing input rate value. [PR1446996](#)
- The bbe-smgd core file **bbe_ifd_add_vlan (ifid=0x8c3e835, ifl=0xcaf59f18)** is generated on the backup Routing Engine at `../../../../src/junos/usr/sbin/bbe-svcs/smd/infra/bbe_ifd.c:6374`. [PR1447493](#)
- Traffic silently drops when using ps- interface over RLT in Layer 2 circuit with **no-control-word** enabled. [PR1447917](#)
- The rpd process might crash if BGP is activated or deactivated multiple times. [PR1448325](#)
- The vehosd process might crash without generating core files and automatic restart of vehosd might fail. [PR1448413](#)
- Interface attributes might cause high CPU usage of dcd. [PR1448858](#)
- FPC might reboot when PIC 0 is taken offline on MX204 platforms. [PR1449067](#)

- The DHCP relay feature might not work as expected with **helpers bootp** configured. [PR1449201](#)
- Increase in the maximum value of **delegation-cleanup-timeout**. [PR1449468](#)
- Currently IS-IS is sending system host-name instead of system-id in OC paths in Isdb or adjacency xpaths in periodic streaming and on-change notification. [PR1449837](#)
- **No localhost ifl for rtt 65535** can be seen on MX Series routers running Junos OS enhanced subscriber management feature. [PR1450057](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)
- The **Mixed Master and Backup RE types** alarm is observed when MX2008 with RE-MX2008-X8-128G detects backup Routing Engine as RE-MX2008-X8-64G. [PR1450424](#)
- VLAN configuration change with l2ald restart might cause kernel synchronization issues and impact forwarding. [PR1450832](#)
- JNP10K-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](#)
- Main chassisd thread at a JNS GNF might stall upon the GNF SNMP polling for hardware-related OIDs. [PR1451215](#)
- SNMP query for IPsec decrypted/encrypted packets does not fetch correct values and **KMD_SNMP_FATAL_ERROR** error is observed. [PR1451324](#)
- DHCP snooping statistics binding does not take effect after deleting and readding the entries. [PR1451688](#)
- RMPC core file is found after configuration changes are done on the network for PTP/clock synchronization. [PR1451950](#)
- Firmware upgrade for PSU (JNP10K-AC2 and JNP10K-DC2) on MX10000 and PTX10000 systems with Routing Engine redundancy configuration enabled might fail due to lcmd being disabled by the firmware upgrade command. [PR1452324](#)
- PLL errors might be seen after FPC reboot or restart. [PR1452604](#)
- Framing errors and packet loss might be seen when high throughput traffic passes through MACsec device. [PR1452851](#)
- Sensord core file might be seen when the script runs on the MPC10E line card. [PR1452976](#)
- The values displayed in the output of **show snmp mib walk jnxTimingNotfnsMIB.3** are not correct. [PR1453436](#)
- PTP goes out of synchronization when HWDB is not accessible during initialization. [PR1453531](#)
- Alarm was not sent to syslog on MX10003 platforms. [PR1453533](#)
- The FPC might crash due to the memory corruption in JNH pool. [PR1453575](#)
- The ANCP interface-set QoS adjusts might not be processed. [PR1453826](#)

- The FPC might crash when the severity of error is modified. [PR1453871](#)
- RADIUS interim accounting statistics are not populated on the MX204. [PR1454541](#)
- The 100-Gigabit Ethernet interfaces might not come up again after going down on MPC3E-NG. [PR1454595](#)
- The access request for L2BSA port up might not be retransmitted if the RADIUS server used is unreachable. [PR1454975](#)
- CRAFTD syslog fatal errors along with junk characters are seen upon its startup and exits after four startup attempts. [PR1454985](#)
- JET/JSD RPC tag handling bug. [PR1455426](#)
- SmiHelperd process is not initialized in Junos OS running on PPC-based platforms. [PR1455667](#)
- Device chooses incorrect source address for locally originated IPv6 packets in routing-instance when destination address is reachable through static route with **next-table** statement. [PR1455893](#)
- BgpRouteInitialize API exits with error code 2. [PR1455967](#)
- Queue data might be missing from the following path: **/interfaces/interface/state**. [PR1456275](#)
- 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 a specified service-set name. [PR1456749](#)
- The bbe-statsd process continuously crashes 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 subscriber routes are not cleared from backup Routing Engine when the session is aborted. [PR1458369](#)
- Subscribers are unable to log in after more than 2 million multicast subscribers are being activated. [PR1458419](#)
- The correct VoIP VLAN information in LLDP-MED packets might not be sent after commit if dynamic VoIP VLAN assignment is used. [PR1458559](#)
- The FPC X major errors alarm might be raised after committing the PTP configuration change. [PR1458581](#)
- 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 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](#)
- Traffic is silently dropped and discarded upon interface flap after DRD autorecovery. [PR1459698](#)
- CPCDD core file is found at **ServicesManager::cpcddSmdInterface::processServiceNotifyMsg ,SmdInterface::cbStateSyncServiceNotifyMsgHandler ,statesync_consumer_poll_new_state_cb**. [PR1459904](#)
- Subscriber statistics might be broken after unified ISSU. [PR1459961](#)
- The PPTP does not work with destination NAT. [PR1460027](#)
- Multiple leaf devices and prefixes are missing when LLDP neighbor is added after a streaming starts at global level. [PR1460347](#)
- If **vlan-offload** is configured on the VMX platform, **input-vlan-map** might not work. [PR1460544](#)
- Support of **del_path** for the LLDP neighbor changes at various levels. [PR1460621](#)
- When you receive IPv6 over IPv4 IBGP session, the IPv6 prefix is hidden. [PR1460786](#)
- The PTP function might consume the kernel CPU for a long time. [PR1461031](#)
- Explicit deletion notification (del_path) is not received when LLDP neighbor is lost as a result of disabling local interface on the DuT using CLI (gNMI). [PR1461236](#)
- A bbe-smgd core file is generated when all RADIUS servers are unreachable. [PR1461340](#)
- Traffic might be impacted because the fabric hardening is stuck. [PR1461356](#)
- On the MPC10E line card, more output packets are seen than expected when the ping function is performed. [PR1461593](#)
- In EVPN scenario, memory leak might be observed when **proxy-macip-advertisement** is configured. [PR1461677](#)
- The rpd core files are generated during system startup. [PR1461796](#)
- Memory leak causes bbe-statsd and bbe-smgd crash. [PR1461821](#)
- On MPC11E line card, the PPS information on the physical interface is inaccurate and varies. [PR1461872](#)
- The rpd might crash after committing the **dynamic-tunnel-anchor-pfe** command. [PR1461980](#)

- The rpd process might crash if the **show v4ov6-tunnels information anti-spoof-ip** command is executed. [PR1462047](#)
- The **CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply failed** message appears when both DIP switches and power switch are turned off. [PR1462065](#)
- On MX204, RADIUS interim accounting statistics are not populated. [PR1462325](#)
- The EA WAN SerDes gets into the **Stuck** state that leads to continuous DFE **tuning timeout** errors and causes the link to stay down. [PR1463015](#)
- The Routing Engine switchover might not be triggered when the master CB clock fails. [PR1463169](#)
- MVPN traffic might be dropped after performing switchover. [PR1463302](#)
- 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](#)
- The bbe-smgd process generates core files on the backup routing engine. [PR1466118](#)
- The mspmand process might crash when stateful firewall and RPC ALG are used on MX Series platforms with MS-MIC or MS-MPC. [PR1464020](#)
- The IPoE subscriber route installation might fail. [PR1464344](#)
- 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](#)
- The PPP IPv6CP might fail if the **routing-services** command is enabled. [PR1464415](#)
- The CPU utilization on mgd daemon might get stuck at 100 percent after the NETCONF session is interrupted by a flapping interface. [PR1464439](#)
- The MS-MIC might not work when it is used on a specific MPC. [PR1464477](#)
- The **show task memory detail** command shows incorrect cookie information. [PR1464659](#)
- The PPPoE session goes in to the **Terminated** state and the accounting stops for the delayed session. [PR1464804](#)
- MPC5E or MPC6E might crash due to internal thread hogging the CPU. [PR1464820](#)
- MPC10E might crash due to inconsistencies during firewall filter add or delete operations. [PR1465153](#)
- The jdncpd might consume high CPU and no further subscribers can be brought up if there are more than 4000 dhcp-relay clients in the MAC-MOVE scenario. [PR1465277](#)
- The physical interface of aggregated Ethernet might take time to come up after disabling or enabling the interface. [PR1465302](#)

- The internal ixlv1 interface might not be created after PXE or network installation. [PR1465547](#)
- ICMP error messages do not appear even after enabling the **enable-asymmetric-traffic-processing** statement. [PR1466135](#)
- The PPPoE subscribers get stuck due to the PPPoE inline keepalives that do not work properly. [PR1467125](#)
- A few DHCP inform packets specific to particular VLAN might be taking the wrong resolve queue. [PR1467182](#)
- Layer 2 wholesale does not forward all the client requests with stacked VLAN. [PR1467468](#)
- Hot-swapping between MPC11E and legacy MPC9, MPC8, or MPC6 is not supported. [PR1467725](#)
- The rpd process might stop after several changes to the flow-spec routes. [PR1467838](#)
- Crypto code might cause high CPU utilization. [PR1467874](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- The tcp-log connections fail to reconnect and get stuck in the **Reconnect-In-Progress** state. [PR1469575](#)
- Memory leak on Layer 2 cpd process causes Layer 2 cpd to crash. [PR1469635](#)
- A **hierarchical-scheduler** should not be configured on a ps- interface. [PR1470049](#)
- The SNMP interface-mib stops working for the PPPoE clients. [PR1470664](#)
- Multiple FreeBSD vulnerabilities fixed in Junos OS. [PR1470693](#)
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card might cause major alarms on other FPCs in the system. [PR1471372](#)
- The clksyncd crash might be seen when PTP over an aggregated Ethernet interface is configured on the MX104 platform. [PR1471466](#)
- Phase or frequency synchronization might not work correctly when PTP is configured in the hybrid mode. [PR1471502](#)
- On the MX10008 and MX10016 line cards, the ARP suppression (default enabled) in EVPN does not work. [PR1471679](#)
- The pccd core file and PCEP session flaps might be seen in PCE-initiated or PCE-delegated LSP scenario. [PR1472051](#)
- Chassis alarm on BSYS might be observed: **RE0 to one or many FPCs is via em1: Backup RE**. [PR1472313](#)
- Service accounting statistics do not get updated after changes are made to the firewall filters. [PR1472334](#)
- The kernel might crash and VM core file might be generated after the configuration change is committed. [PR1472519](#)
- Performing back-to-back rpd restarts might cause rpd to crash. [PR1472643](#)
- Active errors counts do not increase for I2C in the synchronization cards. [PR1472660](#)

- SDB goes down very frequently if the **reauthenticate lease-renewal** statement is enabled for DHCP. [PR1473063](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- Ingress multicast replication does not work with the GRES configuration. [PR1474094](#)
- On MX150 routers, unable to see generated core files for the **show system core-dumps** command. [PR1474118](#)
- An MPC11 crash might occur on MX2000 platform using multidimensional advanced scale configuration that has inline keepalive sessions. [PR1474160](#)
- MX10000 QSA adapter lane 0 port goes in the down state when you disable one of the other lanes. [PR1474231](#)
- A newly added LAG member interface might forward traffic even though its micro BFD session is down. [PR1474300](#)
- The clksyncd generates core files after GRES. [PR1474987](#)
- SFW rule configuration deletion might lead to memory leakage. [PR1475220](#)
- The RADIUS accounting updates of the service session have incorrect statistic data. [PR1475729](#)
- Dark window size is more than expected and 31.0872721524375 seconds of traffic loss is observed. [PR1476505](#)
- The bbe-mibd might crash on the MX Series platform in subscriber environment. [PR1476596](#)
- Traffic loss might be observed to the LNS subscribers in case the **routing-service** statement is enabled under the dynamic-profile. [PR1476786](#)
- In a NAT-T scenario, IKE version 2 tunnel flaps if the tunnel initiator is not behind NAT. [PR1477483](#)
- The Packet Forwarding Engine might be disabled due to major error on MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9 line cards. [PR1478028](#)
- FPC memory leak might happen after you execute the **show pfe route** command. [PR1478279](#)
- Output chain filter counters are not proper. [PR1478358](#)
- The protocol MTU might not be changed on lt- interface from the default value. [PR1478822](#)
- The TCP-log sessions might be in **Established** state, but no logs get sent out to the syslog server. [PR1478972](#)
- The SCBE3 fabric plane gets into check state in MX Virtual Chassis. [PR1479363](#)
- After kmd restart IPsec SA comes up but traffic fails for some time. [PR1480692](#)
- The rpd process might crash when executing the **show route protocol l2-learned-host-routing** or **show route protocol rift** command on a router. [PR1481953](#)
- The MX204 router reboots when the PPPoE client starts to log in and no core files are generated. [PR1482431](#)

- Packet loss might be observed after device reboots or l2ald restarts in an EVPN-MPLS scenario. [PR1484468](#)
- ARP entry might not be created in the EVPN-MPLS environment. [PR1484721](#)
- The logical tunnel interface might not work on MPC10 line cards. [PR1484751](#)
- MPC9E line cards show "Bad Voltage" state when you power on by deleting **set chassis fpc X power off** setting. [PR1485216](#)
- Interface input error counters are not increasing on MX150 routers. [PR1485706](#)
- When rpd starts or restarts, krt-nexthop-ack-timeout might not pick up automatically. [PR1485800](#)
- The unified ISSU is not supported on NG-MPC cards. [PR1491337](#)
- On the MX240, MX480, and MX960 router with SCB3E, FPCs might stay down or restart when you swap MPC7, MPC8, or MPC9 with MPC10, MPC11, or vice versa in the same slot. [PR1491968](#)
- DHCP subscribers do not come up as expected after deactivating Virtual Chassis port. [PR1493699](#)
- The ptp-clock-global-freq-tracable leaf value is not changing to True. [PR1493743](#)
- UID might not release properly in some scenarios after service session deactivation. [PR1188434](#)
- Need to change the default parameters for resource-monitor rtt-parameters. [PR1407021](#)
- The **show system subscriber-management summary** command needs to include failure reason for standby disconnect when primary and backup Routing Engine memories do not match. [PR1422976](#)
- The **show subscriber extensive** command incorrectly displays the DNS address provided to DHCP clients. [PR1457949](#)
- DHCP relay with forward-only fails to send OFFER when the client is terminated on the logical tunnel interface. [PR1471161](#)
- Dynamic-profile for VPLS-PW pseudowire incorrectly reports the Dynamic Static Subscriber Base Feature license alarm. [PR1473412](#)
- DHCP-server: RADIUS given mask is being reversed. [PR1474097](#)

Infrastructure

- The duplex status of the management interface might not be updated in the output of the **show** command. [PR1427233](#)
- The operations on the console might not work if the **system ports console log-out-on-disconnect** command is configured. [PR1433224](#)
- On all Junos OS VM-based platforms, FPC might reboot if jlock hog occurs. [PR1439906](#)
- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later releases. [PR1462986](#)
- The scheduled tasks might not be executed if cron daemon goes down without restarting automatically. [PR1463802](#)

Interfaces and Chassis

- Restarting chassisd with GRES disabled might cause FPC to restart and some demux interfaces to be deleted. [PR1337069](#)
- Unrelated aggregated Ethernet interfaces might go down if you commit configuration changes. [PR1409535](#)
- MX Series Virtual Chassis unified ISSU is not supported when Redundant LT (RLT) is configured. [PR1411729](#)
- The demultiplexer interfaces will be down after the MTU of the underlying et- interface is changed. [PR1424770](#)
- Upgrade from releases before Junos OS Release 17.4R1 results in the generation of cfmd core files. [PR1425804](#)
- The NCP session might be brought down after IPCP configure-reject is sent. [PR1431038](#)
- VRRP mastership might flap when the tracked route is deleted or the tracked interface goes down. [PR1432361](#)
- In MX Series platform where PPPoE is used, the router might not send LCP termination-request or LCP terminate-ack. [PR1433489](#)
- The output of the **show interfaces <>** command for AFT card might be different from legacy card. [PR1435416](#)
- Mixed link-speed ae- bundle could not add new subinterface 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](#)
- The cfmd process might crash after a restart on Junos OS Release 17.1R1 and later. [PR1443353](#)
- Enhancement of add or delete a single VLAN in **vlan-id-list** under interface family bridge. [PR1443536](#)
- When the logical interface is associated to a routing-instance inside a LR, the logical interface is removed from routing-instance and the logical interface is not added to the default routing instance. [PR1444131](#)
- Unified 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](#)
- Continuous VRRP state transition (VRRP master or backup flaps) is observed when one device drops the VRRP packets. [PR1446390](#)
- Unable to connect to newly installed Routing Engine from other Routing Engine in MX Series Virtual Chassis. [PR1446418](#)
- VRRP dual-master status is seen after Routing Engine switchover on the backup router. [PR1447028](#)
- The l2ald might fail to update composite next hop. [PR1447693](#)
- The ifinfo daemon might crash on the execution of the **show interface extensive** command. [PR1448090](#)

- Interface descriptions might be missing under logical systems CLI. [PR1449673](#)
- Dual VRRP mastership might be seen after Routing Engine switchover occurs ungracefully. [PR1450652](#)
- LACP daemon crashes continuously. [PR1450978](#)
- The severity level log might be flooded when the QSFP-100GE-DWDM2 is inserted. [PR1453919](#)
- CFM UP MEP session does not come up in scaled scenario over L2VPN circuits on LAG interfaces. [PR1454187](#)
- The VRRP traffic loss is longer than one second for some backup groups after performing a GRES. [PR1454895](#)
- Mismatched MTU value causes the RLT interface to flap. [PR1457460](#)
- The EOAM CFM primary-vid functionality does not work if the **enhanced-cfm-mode** is enabled. [PR1465608](#)
- The vrrpv3mibs does work to poll the VRRPv6-related objects. [PR1467649](#)
- The voltage high alarm might not be cleared when the voltage level comes back to normal for MIC on MPC5. [PR1467712](#)
- When you configure ESI on a physical interface, the traffic drops when you disable the logical interface under the physical interface. [PR1467855](#)
- Executing commit might hang up due to stuck dcd process. [PR1470622](#)
- Commit error was not thrown when member link was added to multiple aggregation groups with different interface specific options. [PR1475634](#)
- The interface on MIC3-100G-DWDM might go down after performing an interface flap. [PR1475777](#)
- Multichassis aggregated Ethernet interface might be shown as an unknown status when you add the subinterface as part of the VLAN on the peer multichassis aggregated Ethernet node. [PR1479012](#)

J-Web

- Session fixation vulnerability in J-Web. [PR1410401](#)
- Cross-site scripting (XSS) in J-Web. [PR1434553](#)
- Some error messages might be seen when using J-Web. [PR1446081](#)
- Security vulnerability in J-Web and web based (HTTP/HTTPS) services. [PR1499280](#)

Junos Fusion for Enterprise

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

Layer 2 Ethernet Services

- LACP PDU might be looped toward peer MC-LAG nodes. [PR1379022](#)
- Error messages might be seen when you add a logical interface for physical interfaces. [PR1424106](#)
- The DHCP DECLINE packets are not forwarded to the DHCP server when **forward-only** is set within **dhcp-reply**. [PR1429456](#)
- DHCP request might get dropped in a DHCP relay scenario. [PR1435039](#)
- The jdhcpd process might go into an infinite loop and cause full utilization of CPU. [PR1442222](#)
- The ppman and aft process might crash if ppm control logs are enabled. [PR1443410](#)
- The **dhcp-relay** statement might not work on MX10008 and MX10016 platforms. [PR1447323](#)
- DHCPv6 authentication via RADIUS server might fail as a result of the missing VSA option 26-207. [PR1448100](#)
- Multiple vulnerabilities in JDHCPD allow for OS command injection and code execution of JDHCPD. [PR1449353](#)
- PPPoE holding DHCPv6 prefix causes DHCPv6 binding failure due to a duplicate prefix. [PR1453464](#)
- DHCP subscriber might not come online after rebooting the router. [PR1458150](#)
- DHCP packet might not be processed correctly if DHCP option 82 is configured. [PR1459925](#)
- The metric does not change when configured under DHCP. [PR1461571](#)
- The ISSU might fail during the subscriber in-flight login. [PR1465964](#)
- Telemetry data for **relay/bindings/binding-state-v4relay-binding** and **relay/bindings/binding-state-v4relay-bound** are not correct. [PR1475248](#)

Layer 2 Features

- LSI interface might not be created, which prevents MAC addresses from being learned. The following error is seen: **RPD_KRT_Q_RETRIES: ifl iff add: Device busy**. [PR1295664](#)
- VPLS neighbors might stay in the down state after changing the configuration in **vlan-id**. [PR1428862](#)
- After disabling and enabling the aggregate interface, the next hop of CE-facing aggregate interface might be in a wrong state. [PR1436714](#)
- The operational status of the interface in hardware and software might be out of synchronization in an EVPN setup with the proxy ARP feature enabled. [PR1442310](#)
- In Virtual Chassis scenario, traffic drop might be seen when one Virtual Chassis member reboots and rejoins the Virtual Chassis. [PR1453430](#)
- Connectivity is broken through LAG due to members configured with **hold-time** and **force-up**. [PR1481031](#)

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](#)
- The rpd might restart after an MPLS LSP flap if **no-cspf** and **fast-reroute** are configured in an LSR ingress router. [PR1368177](#)
- RSVP LSP might get stuck in down state in an OSPF multiarea topology. [PR1417931](#)
- MPLS LSP autobandwidth statistics miscalculations might lead to high bandwidth reservation. [PR1427414](#)
- Continuous rpd core files at **l2ckt_alloc_label** , **l2ckt_standby_assign_label** , **l2ckt_intf_change_process** in new backup during GRES in MX2010 box. [PR1427539](#)
- The LDP might withdraw a label for an FEC after the IGP route is inactive in inet.0. [PR1428843](#)
- MPLS ingress LSPs might not come up after MLPS is disabled or enabled. [PR1432138](#)
- SRLG entry shows unknown after removing it from configuration in the **show mpls lsp extensive** or **show mpls srlg** output. [PR1433287](#)
- Restart routing might result in RPD core files while GRES and NSR are enabled. [PR1433857](#)
- Traffic loss might occur if p2mp with NSR is enabled. [PR1434522](#)
- The P2MP LSP branch traffic might be dropped for a while when the sender PE device switchover occurs. [PR1435014](#)
- Traffic loss might be seen after the LDP session flaps rapidly. [PR1436119](#)
- The rpd might crash after executing the **ping mpls ldp** command. [PR1436373](#)
- The flow label is not pushed when **chained-composite-next-hop ingress l2ckt/l2vpn** is enabled. [PR1439453](#)
- The LDP route and LDP output label are not displayed in the inet.3 table and LDP database respectively if **OSPF rib-group** is enabled. [PR1442135](#)
- The active path of a no-cspf LSP might keep flapping when one or more transit nodes are shared by primary path and secondary path. [PR1442495](#)
- The backup LSP path messages are rejected if the bypass tunnel path is an interarea LSP. [PR1442789](#)
- RSVP path message with long refresh interval is dropped between nodes running Junos OS releases earlier than and later than Release 16.1. [PR1443811](#)
- P2MP LSP might get stuck in the down state after link flaps. [PR1444111](#)
- The rpd memory leak might be seen when the interdomain RSVP LSP is in the down state. [PR1445024](#)
- Traffic might be silently dropped or discarded if two consecutive PLRs along the LSP perform local repair simultaneously under certain misconfigured conditions. [PR1445994](#)
- The transit packets might be dropped if an LSP is added or changed on MX Series device. [PR1447170](#)

- Traffic drop might be seen after traceoption configuration is committed in RSVP P2MP scenario. [PR1447480](#)
- The LDP route timer resets when you commit unrelated configuration changes. [PR1451157](#)
- The traffic might be silently discarded after the LACP timeouts. [PR1452866](#)
- The rpd crash might be observed with traceoption enabled in MPLS. [PR1457681](#)
- All LDP adjacencies flap after changing LDP preference. [PR1459301](#)
- Previously configured credibility preference is not considered by CSPF despite that the configuration is 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](#)
- The rpdtdm process might crash while SNMP polls the statistics of the lpd interface. [PR1465729](#)
- The device might use the locally computed path for the PCE-controlled LSPs after link or node fails. [PR1465902](#)
- The rpd process might crash during shutdown. [PR1471191](#)
- The rpd crash might be seen after some commit operations, which might affect the RSVP ingress routes. [PR1471281](#)
- The following error messages continuously flood the backup Routing Engine:
(JTASK_IO_CONNECT_FAILED: RPDTM./var/run/rpdtdm_control: Connecting to 128.0,255.255,255.255,0.0.0,0.0.0, failed: No such file or directory). [PR1473846](#)
- RSVP LSPs might not come up in scaled network with a very high number of LSPs if NSR is used on the transit router. [PR1476773](#)
- Kernel crashes and device might restart. [PR1478806](#)
- RPD 100 percent CPU load and RPD core files are generated on the backup Routing Engine. [PR1479249](#)
- The rpd core files are generated during unified ISSU. [PR1493969](#)

Network Address Translation (NAT)

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

Network Management and Monitoring

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

Platform and Infrastructure

- The jcrypto syslog help package and events are not packaged even when errmsg is compiled. [PR1290089](#)
- LACP DDoS policer is incorrectly triggered by other protocol's traffic. [PR1409626](#)

- Error logs might be observed after performing unified ISSU. [PR1412463](#)
- The slax scripts triggered by event options might be stuck forever. [PR1422939](#)
- Packet drops, replication failure, or ksyncd crashes might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- With CNH for 6PE, MPLS EXP rewrite rule for non-VPN IPv4 over MPLS traffic might not work. [PR1430878](#)
- The FPC might crash when the firewalls filter manager deals with the firewall filters. [PR1433034](#)
- TWAMP session remains in pending state when cooperating with a non-Juniper device. [PR1434740](#)
- Traffic from the same physical interface cannot be forwarded. [PR1434933](#)
- The device might not be accessible after the upgrade. [PR1435173](#)
- BR for MAP-E does not return ICMP Type=3/Code=4 when over MTU sized packet comes with DF bit. [PR1435362](#)
- MAP-E encapsulation or de-encapsulation with specific parameter might work incorrectly. [PR1435697](#)
- The RPM http-get probe always returns HTTP 400 error. [PR1436338](#)
- The `/var/db/scripts` directory might be deleted after you execute the **request system zeroize** command. [PR1436773](#)
- With CNH enabled, the MPLS CoS rewrite does not work for 6PE traffic. [PR1436872](#)
- The BGP session might flap after you perform Routing Engine switchover simultaneously on both end of BGP peers. [PR1437257](#)
- The next-hop MAC address in the output for the **show route forwarding-table** command might be incorrect. [PR1437302](#)
- A certain combination of **allow-commands/deny-commands** does not work properly after Junos OS Release 18.4R1. [PR1438269](#)
- The inner IPv4 packet might get fragmented using the same size as the configured `mtu-v6`, which is used for the MAP-E software tunnel in the MAP-E configuration. [PR1440286](#)
- The RPM udp-ping probe does not work in a multiple routing instance scenario. [PR1442157](#)
- ARP resolution might fail after ARP HOLD next hops are added and deleted continuously. [PR1442815](#)
- When a host-bound packet is received in a 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 interfaces are unreachable after the router or FPC restarts. [PR1449615](#)

- FPC might reboot with vmcore due to memory leak. [PR1449664](#)
- REST API process becomes nonresponsive when a number of requests come in at a high rate. [PR1449987](#)
- In an 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](#)
- The MPC might drop packets after you enable the firewall fast lookup filter. [PR1454257](#)
- The DDoS protection does not stop logging when remote tracing is enabled. [PR1459605](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- CLI configuration flag **version-03** must be optional. [PR1462186](#)
- On the MX204 platform, Packet Forwarding Engine errors occur when the incoming GRE tunnel fragments get sampled and undergo inline reassembly. [PR1463718](#)
- EVPN-VXLAN T-5 tunnel does not work properly. [PR1466602](#)
- On the MX150 devices, the default subscriber management license does not include Layer 2 TP. [PR1467368](#)
- The Layer 2 traffic over ae- interfaces sent from one member to another member is corrupted on MX Series Virtual Chassis. [PR1467764](#)
- The JNH memory leaks after CFM session flap for LSI and VT interfaces. [PR1468663](#)
- The switch might not be able to learn MAC address with **dot1x** and **interface-mac-limit** configured. [PR1470424](#)
- SSH login might hang, and the TACAS plus server closes the connection without sending any authentication failure response. [PR1478959](#)
- The time convergence for the MVPN fast upstream failover might be more than 50 ms. [PR1478981](#)
- **Show system buffer** command displays all zeros in the MX104 chassis. [PR1484689](#)
- MAC malformation might occur in rare scenarios under MX Series Virtual Chassis. [PR1491091](#)
- A specific IPv4 packet might lead to FPC restart. [PR1493176](#)
- Routing Engine crash might be seen when a large number of next hops are quickly deleted and added again in a large ARP or ND scale scenario. [PR1496429](#)
- Traffic to VRRP VMAC/VIP drops as "my-mac check failed" when ingress queueing enabled on an aggregated Ethernet interface. [PR1501014](#)
- Python or Slax script might not be executed. [PR1501746](#)

Routing Policy and Firewall Filters

- The **route-filter-list** configuration with noncontinuous match might not work as expected after being updated. [PR1419731](#)
- Policy matching RD changes next hop of the routes that do not carry the RD. [PR1433615](#)
- The rib-group might not process the exported route correctly. [PR1450123](#)
- The rpd might crash after Routing Engine switchover when **prefix-list** is configured [PR1451025](#)
- Routes resolution might be inconsistent if any route resolves over the multipath route. [PR1453439](#)

Routing Protocols

- The rpd crashes in Junos OS Release 16.1 or later during BGP convergence. [PR1351639](#)
- Routing Engine-based micro-BFD packets do not go out with configured source IP address when the interface is in logical system. [PR1370463](#)
- The rpd might crash under a rare condition if GR helper mode is triggered. [PR1382892](#)
- Processing a large scale AS-path regex causes the flap of the route protocols to flap. [PR1396344](#)
- BFD link-failure detection of the broken path is delayed when IGP link-state update is received from the same peer through an alternative path. [PR1410021](#)
- BGP might get stuck in the idle state when the peer triggers a GR restart event. [PR1412538](#)
- The Layer 3 VPN link protection does not work after flapping the CE-facing interface. [PR1412667](#)
- Transit traffic does not forward under TI-LFA and IS-IS overLoad bit setting scenario. [PR1412923](#)
- Multicast traffic might be lost for around 30 seconds during Routing Engine switchover. [PR1427720](#)
- The next hop of IPv6 route remains empty when a new IS-IS link comes up. [PR1430581](#)
- The BGP configuration statement **multipath multiple-as** does not work in specific scenario. [PR1430899](#)
- IPv6 aggregate routes are hidden. [PR1431227](#)
- The rpd process might crash continuously if **egress-te** is configured under the EBGp VRF routing instance. [PR1431536](#)
- The **show isis adjacency extensive** output does not contain the state transition details. [PR1432398](#)
- In BFD and GR enabled scenario, BFD DOWN packets are not being sent immediately after BFD failure. [PR1432440](#)
- Per-prefix LFA might not work as expected where the last hop needs to be protected on the penultimate node. [PR1432615](#)
- PIM-SM join message might be delayed when MSDP is enabled. [PR1433625](#)
- The rpd crashes after removing MVPN configuration from a VRF instance. [PR1434347](#)
- With SR enabled, 6PE next hop is not installed. [PR1435298](#)

- The rpd might crash during the best-path changes in BGP-L3VPN with multipath and **no-vrf-propagate-ttl** enabled. [PR1436465](#)
- BGP route next hop can be incorrect in some scenarios with PIC edge configuration. [PR1437108](#)
- Removing SSH protocol version 1 from configuration. [PR1440476](#)
- RIP routes are discarded by Juniper Networks device over a /31 subnet interface. [PR1441452](#)
- The rpd process might crash in inter-AS option B Layer 3 VPN scenario if CNH is used. [PR1442291](#)
- The CPU utilization on rpd spins at 100 percent once the same external BGP route is learned on different VRF tables. [PR1442902](#)
- 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 rpd process crashes if the multicast scope with an invalid prefix is configured and committed. [PR1445746](#)
- BRP: RPC call is missing 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 a VRF instance. [PR1446437](#)
- The rpd might crash when the policy applied to the MoFRR is deleted. [PR1446472](#)
- 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 Series of devices, Layer 3 VPN PE-CE link protection exhibits unexpected behavior. [PR1447601](#)
- Junos OS BFD sessions with authentication flaps occur after some time. [PR1448649](#)
- The connection between ppmr (Routing Engine) and ppmn (FPC) might get lost due to session timeout. [PR1448670](#)
- The BGP routes might fail to be installed in routing instance if the **from next-hop** policy match condition is used in the VRF import policy. [PR1449458](#)
- The TI-LFA backup path for adj-sid is broken in OSPF. [PR1452118](#)
- SPRING-LDP interoperability issues are observed with colocated SRMS+SR-client+LDP-stitching. [PR1452956](#)
- The SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. [PR1454177](#)

- The rpd scheduler slip for BGP GR might be up to 120 seconds after the peer goes down. [PR1454198](#)
- MoFRR with MLDP inband signaling is not working. [PR1454199](#)
- The rpd memory might leak in a certain MSDP scenario. [PR1454244](#)
- Invalid BGP update sent to peer device might cause BGP session to terminate. [PR1454677](#)
- The rpd might crash when BGP features ORR and IS-IS are configured. [PR1454803](#)
- The rpd process 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](#)
- Packets 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](#)
- The rpd might crash when OSPF router-id gets changed for NSSA with the **area-range** configured. [PR1459080](#)
- The rpd memory leak might be observed on backup Routing Engine due to BGP flap. [PR1459384](#)
- The **other querier present interval** timer cannot be changed in an IGMP or MLD snooping scenario. [PR1461590](#)
- Rpd scheduler slips might be seen on RPKI route validation-enabled BGP peering router in a scaled setup. [PR1461602](#)
- Need to install all possible next hops for OSPF network LSAs. [PR1463535](#)
- IS-IS IPv6 multitopology 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](#)
- The rpd might crash if the IPv4 routes are programmed with IPv6 next hop through JET APIs. [PR1465190](#)
- The BGP peers might flap if the **hold-time** parameter is set as small. [PR1466709](#)
- The configured BGP damping policy might not take effect after BGP is disabled and then enabled followed by commit. [PR1466734](#)
- BGP multipath does not work for MT on cRPD. [PR1467091](#)
- The rpd might crash after configuring **independent-domain** under the master routing instance. [PR1469317](#)
- The mcsnoopd might crash when the STP moves the mrouter port to the blocked state. [PR1470183](#)
- The BFD client session might flap when removing the BFD configuration from the peer end (from other vendor) of the BFD session. [PR1470603](#)
- The rpd might stop when both the **instance-import** and **instance-export** policies contain the **as-path-prepend** action. [PR1471968](#)
- The rpd process might crash with BGP multipath and damping configured. [PR1472671](#)

- Removal of the cluster from BGP group might cause prolonged convergence time. [PR1473351](#)
- The rpd process might crash with BGP multipath and route withdrawal occasionally. [PR1481589](#)
- The rpd crashes if the same neighbor is configured in different RIP groups. [PR1485009](#)
- The BGP-LU routes do not have the label when BGP sharding is used. [PR1485422](#)
- The rpd might crash when you perform GRES with MSDP configured. [PR1487636](#)
- High CPU utilization might be observed when the outgoing BGP updates are sent slowly. [PR1487691](#)
- BGP RIB sharding feature cannot be run on a system with a single CPU. [PR1488357](#)
- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

Services Applications

- The kmd process might crash when DPD timeouts for some IKEv2 SAs occur. [PR1434521](#)
- Traffic might be dropped in an IPsec VPN scenario when the VPN peer is behind a NAT device. [PR1435182](#)
- The output of the **show subscriber user-name** on LTS shows only one session instead of two. [PR1446572](#)
- The jl2tpd process might crash during the restart procedure. [PR1461335](#)
- On an MX Series router, L2tp LTS fails to forward the agentCircuitId and agentRemoteld AVP toward the LNS. [PR1472775](#)
- The kmd might crash due to the incorrect IKE SA establishment after the remote peer's NAT mapping address is changed. [PR1477181](#)

Subscriber Access Management

- Subscriber filtering for general authentication services traceoptions could report debug messages for other users. [PR1431614](#)
- Incorrect Acct-Session-Time and no LCP Termination-Ack by MX Series BNG. [PR1433251](#)
- Subscriber deactivation might get stuck in terminated state. [PR1437042](#)
- Missing **<radius-server-data>** tags on test ppp aaa display XML output. [PR1444438](#)
- On MX Series platforms, there might be a false error for SAE policy activation or deactivation failure. [PR1447632](#)
- Subscriber's login fails when PCRF server is unreachable. [PR1449064](#)
- DHCPv6 subscribers might be stuck in a state after the authd process crashes. [PR1460578](#)
- The subscriber address allocation might fail after deleting the pool link in the middle of the chain. [PR1465253](#)

- The volume statistics attributes are missing in the accounting-stop for the Configuration Activated Services and CLI Activated Services. [PR1470434](#)
- The subinterfaces might be missing in the NAS port ID. [PR1472045](#)
- The authd process might crash after the unified ISSU setup from Junos OS Release 18.3 and earlier to Junos OS Release 18.4 and later. [PR1473159](#)
- Some address-relevant fields are missing when executing the **test aaa ppp** command. [PR1474180](#)
- The CoA request might not be processed if it includes the **proxy-state** attribute. [PR1479697](#)
- The **mac-address** CLI option is hidden under the **access profile *profile-name* radius options calling-station-id-format** statement. [PR1480119](#)

User Interface and Configuration

- The **show chassis hardware satellite** command is not available in Junos OS Release 17.3. [PR1388252](#)
- On an MX Series device, a J-Web page might not get redirected to login once the session expires with an idle timeout. [PR1459888](#)

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](#)
- MPLS LSP ping over Layer 2 circuit might not work when **flow-label** is enabled. [PR1421609](#)
- MVPN using PIM dense mode does not prune the OIF when PIM prune is received. [PR1425876](#)
- The P1 configuration delete message is not sent on loading baseline configuration if there is a prior change in VPN configuration. [PR1432434](#)
- The resumed multicast traffic for certain groups might be stopped in overlapping MVPN scenario. [PR1441099](#)
- 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 l2circuit displays MM status, which may cause traffic loss. [PR1462583](#)
- The Layer 2 circuit connections might become stuck in OL state after changing the Layer 2 circuit community and flapping the primary LSP path. [PR1464194](#)
- The rpd might crash when the **link-protection** is added to or deleted from LSP for the MVPN ingress replication selective provider tunnel. [PR1469028](#)
- Layer 2 circuit stuck in RD state at one end. [PR1498040](#)

Resolved Issues: 19.2R1

Application Layer Gateways (ALGs)

- DNS requests with the EDNS option might be dropped by the DNS ALG. [PR1379433](#)

Authentication and Access Control

- The dot1xd might crash when dot1xd receives incorrect reply length from the authd. [PR1372421](#)
- Push-to-JIMS now supports push auth entry to all online jims servers. [PR1407371](#)

Class of Service (CoS)

- Traffic drop occurs when deleting MPLS family or disabling the interface that has non-default EXP rewrite-rules. [PR1408817](#)

EVPN

- The rpd process would crash if deactivating the Autonomous-System (AS) in an EVPN scenario. [PR1381940](#)
- The RA packets might be sent out without using the configured virtual gateway address. [PR1384574](#)
- [EVPN/VXLAN] VTEP tunnel does not get deleted when EVPN peer goes down. [PR1390965](#)
- On EVPN setups, incorrect destination MAC addresses starting with 45 might show up when using the **show arp hostname** command. [PR1392575](#)
- The rpd process might crash with EVPN type-3 route churn. [PR1394803](#)
- The rpd process generates core files upon Routing Engine switchover with scaled EVPN configuration. [PR1401669](#)
- The rpd crashes due to memory corruption in EVPN. [PR1404351](#)
- EVPN database and bridge MAC-table are out of sync due to the interface's flap. [PR1404857](#)
- EVPN routes might show **Route Label: 0** in addition to the real label. [PR1405695](#)
- The rpd might crash after NSR switchover in an EVPN scenario. [PR1408749](#)
- Local L2ALD proxy MAC+IP advertisements accidentally delete MAC+IP EVPN database state from remotely learned type 2 routes. [PR1415277](#)
- The rpd process crash on backup Routing Engine after enabling nonstop-routing with EVPN. [PR1425687](#)
- The device might proxy the ARP probe packets in an EVPN environment. [PR1427109](#)
- IP is missing in mac-ip-table of EVPN database but is present in the EVPN database when CE interface has two primary IP address. [PR1428581](#)
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs. [PR1429821](#)
- Stale MAC addresses are present in the bridge MAC-table in a EVPN/MPLS scenario. [PR1432702](#)

- Configuring ESI on a single-homed 25G port might not work. [PR1438227](#)
- The RPD process might crash after you commit the changes. [PR1439537](#)

Flow-based and Packet-based Processing

- Fragmentation and ALG support for Power Mode IPsec. [PR1397742](#)

Forwarding and Sampling

- The LSI binding for the IPv6 neighbor is missing. [PR1388454](#)
- Firewall flexible match syntax clarification. [PR1389103](#)
- In Junos OS Release 13.3R9.13, the firewall filter action, "decapsulate gre", decapsulates gre, ip-over-ip, and ipv6-over-ip, but in 17.3R3.9, it only decapsulates gre. [PR1398888](#)

General Routing

- In a BGP/MPLS scenario, if the next hop type of label route is indirect, disabling and enabling the **family mpls** of the next hop interface might cause the route to go into a dead state. [PR1242589](#)
- Large-scale user's log in and log out might cause mgd memory leak. [PR1352504](#)
- Packet Forwarding Engine selector get stuck in rerouted state on unilist NH after primary aggregated Ethernet interface is link deactivated and activated. [PR1354786](#)
- The voltage high alarm might not be cleared when the voltage level comes back to normal for a MIC on an MPC5E. [PR1370337](#)
- The filter service might fail to get installed for the subscriber in a scaled BBE scenario. [PR1374248](#)
- In a subscriber scenario, FPC errors might be seen. [PR1380566](#)
- The routes learned over an interface will be marked as "dead" next hop after changing the prefix-length of IPv6 address on that interface. [PR1380600](#)
- Traffic is silently discarded that is caused by FPC offline in a MC-LAG scenario. [PR1381446](#)
- High cpu utilization for chassisd on bsys, approximately 20 percent at steady state. [PR1383335](#)
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent Major Alarm. [PR1384435](#)
- Subscriber connection setup is 30 percent lower than expected. [PR1384722](#)
- The rpd might crash when switchover is performed along with configuration changes being committed. [PR1385005](#)
- Incorrect log message for chip errors (extra dash "-"). [PR1385066](#)
- The MPC10E line card interface filter statistics are not showing the input packet count or rejects. The **show pfe statistics traffic** statement does not report for any normal discard. [PR1383579](#)
- The rpd and KRT queue might get stuck in a VRF scenario. [PR1386475](#)

- Behavior of the **set interfaces ams0 service-options session-limit rate <integer value>** has changed. [PR1386956](#)
- Migrate from syslog API to errmsg API - VMhost messages on Junos OS. [PR1387099](#)
- Some SFBs might go down when one of the PSMs in the chassis generates a bad output voltage which is out-of-range. [PR1387737](#)
- IPsec IKE keys are not cleared when delete/clear notification is received. [PR1388290](#)
- BBE SMGD core files are generated if MTU is changed while subscribers are logged in on the physical interface. [PR1389611](#)
- The jnxFruState might show incorrect PIC state after replacing a MPC with another MPC having less PICs. [PR1390016](#)
- Traffic destined to VRRP VIP gets dropped as filter is not updated to related logical interface. [PR1390367](#)
- **Delete chassis redundancy** will not give commit warning. [PR1390575](#)
- The BNG might not respond with PADO and create any Demux interface when PPPoE PADI packet is received. [PR1390989](#)
- The Packet Forwarding Engine might not respond with ICMP time exceeded error when packet arrives from the subscriber. [PR1391932](#)
- FPC might reboot on vMX in a subscriber scenario. [PR1393660](#)
- Junos OS enhancement configuration statement to modify mcontrol watchdog timeout. [PR1393716](#)
- The FPC cards might not come up while performing unified ISSU on MX10003. [PR1393940](#)
- IDS aggregate configuration statement should not be considered for the installation of the IDS dynamic filter [PR1395316](#)
- L3 gateway did not update ARP entries if IP or MAC quickly move from one router to another router in EVPN-VXLAN environment. [PR1395685](#)
- The MPC, and Forwarding Engine Board (AFEB or TFEB) with channelized OC MIC might crash with the generation of core files. [PR1396538](#)
- Adding IRB to bridge-domain with PS interface causes kernel crash. [PR1396772](#)
- Subscriber flapping might cause SMID resident memory leak. [PR1396886](#)
- The routing protocol process (rpd) has facilities to attempt to trap certain classes of nonfatal bugs by continuing to run, but it generates a "soft" core file. [PR1396935](#)
- Seeing **VMHost RE 0 Secure BIOS Version Mismatch** and **VMHost RE 1 Secure Boot Disabled** alarms. [PR1397030](#)
- The service PIC might crash while changing CGNAT mode. [PR1397294](#)
- The **show system firmware** command might provide unexpected output on some MX Series routers such as MX104. [PR1398022](#)

- Wrong transmit clock quality is observed when router is in holdover. [PR1398129](#)
- MPLSoUDP/MPLSoGRE tunnel might not come up on the interface route. [PR1398362](#)
- JET/PRPD incompatibility for the rib_service.proto field RouteGateway.weight from Junos OS Release 18.4R1 to Release 18.4R2 onward. [PR1400563](#)
- The mgd-api might crash due to a memory leak. [PR1400597](#)
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- The **show | compare** command output on global group changes lose the diff context after a rollback or 'load update' is performed. [PR1401505](#)
- The TCP connection between ppmd and ppmn might be dropped due to a kernel issue. [PR1401507](#)
- The FPC generates core files due to a corner case scenario (race condition between RPF, IP flow). [PR1401808](#)
- Traffic loss is seen in IGMP subscribers after GRES. [PR1402342](#)
- The MPC might crash due to the CPU hogging by dfw thread. [PR1402345](#)
- DHCP subscriber cannot reconnect over dynamic VLAN demux interfaces due to RPF check failure. [PR1402674](#)
- Observed rpd core files when few colored LSPs changed to uncolored LSPs. [PR1403208](#)
- The sync_response received earlier for interface sensor subscribed in on-change mode. [PR1403672](#)
- Continuous kernel crashes might be observed in the backup Routing Engine or VC-BM. [PR1404038](#)
- With MS-MPC and MS-MIC service cards, Syslog messages for port block interim might show 0.0.0.0 for the private-IP and PBA release messages might show the NAT'd IP as the private IP. [PR1404089](#)
- The FPC might crash in a CoS scenario. [PR1404325](#)
- The repd continues to generate core files on VC-Bm when there are too many IPv6 addresses on one session. [PR1404358](#)
- Incorrect output of the assigned prefixes to the subscriber in the output of the **show interface < dynamic demux interface>** command. [PR1404369](#)
- On an MX10003 and an MX10008, its i2c bus might fail a read operation. [PR1405787](#)
- MPC might generate core files after restarting the FPC that belongs to targeting aggregate Ethernet and host subscribers. [PR1405876](#)
- NAT64 translation issues of **ICMPv6 Packet Too Big** message with MS-MPC/MS-PIC. [PR1405882](#)
- The FPC crash might be observed in MS-MPC HA environment. [PR1405917](#)
- Fabric performance drops on MPC7, MPC8, and MPC9E and SFB2 based MX2000 routers. [PR1406030](#)
- A rpd crash is seen post configuration commit and bt has pointers on receiving SNMP packet. [PR1406357](#)

- Traffic impact might be seen if auto-bandwidth is configured for RSVP LSPs. [PR1406822](#)
- New CLI option to display DF and MLR in split format. [PR1406884](#)
- MX10003 gives a cosmetic error message **ALARMD_CONNECTION_FAILURE: after 60 attempts crafted connect returned error: Connection refused**. [PR1406952](#)
- Layer 2 VPN might flap repeatedly after the link up between PE and CE devices. [PR1407345](#)
- The rpd might crash when a commit check is executed on LDP trace options filtering. [PR1407367](#)
- NPC core file is generated after daemon restart in `jnh_get_oif_nh ()` routine. [PR1407765](#)
- Ephemeral database might get stuck during commit. [PR1407924](#)
- Traffic forwarding fails when crossing VCF members. [PR1408058](#)
- **openconfig-network-instance:network-instances** support for IS-IS must be hidden unless supported. [PR1408151](#)
- Group VPN (GVPN): ToS/DSCP byte is not copied into the outer IPSec header during IP header preservation. [PR1408168](#)
- Alarm mismatch in total memory is detected after **reboot vmhost both**. [PR1408480](#)
- The MPC line cards might crash when performing unified ISSU to Junos OS Release 19.1R1 or above. [PR1408558](#)
- Python script might stop working due to **Too many open files** error. [PR1408936](#)
- MX-service templates are not cleaned up. [PR1409398](#)
- MX-MPC2-3D-EQ and MPC-3D-16XGE-SFPP will now show the **Exhaust A** temperature, rather than the Intake temperature. [PR1409406](#)
- MIC-MACSEC-20GE supports Extended Packet Numbering (XPN) mode on 1-Gigabit or 10-Gigabit Ethernet interfaces [PR1409457](#)
- Telemetry: **interface-set meta-data** needs to include the CoS TCP names in order to aid collector reconciliation with queue-stats data. [PR1409625](#)
- The non-existent subscribers might appear at **show system resource-monitor subscribers-limit chassis extensive** output. [PR1409767](#)
- FPC might crash during next hop change when using MPLS inline-jflow. [PR1409807](#)
- MX80 drops DNS responses which contain an underscore. [PR1410062](#)
- When using SFP+, the interface optic output might be non-zero even though the interface has been disabled. [PR1410465](#)
- Traffic loss might be seen on MPC8E or MPC9E after request one of the SFB2s offline/online. [PR1410813](#)
- Kernel replication failure might be seen if an IPv6 route next hop points to an ether-over-atm-llc ATM interface. [PR1411376](#)

- Packet Forwarding Engine heap memory leak might happen during frequent flapping of PPPoE subscribers connected over aggregated Ethernet interface. [PR1411389](#)
- Virtual Route Reflector might report **DAEMON-3-JTASK_SCHED_SLIP_KEVENT** error on some hypervisor or host machine because of NTP sync. Routing protocol might be impacted. [PR1411679](#)
- If GRE over GRE tunnel is used for sending Routing Engine-originating traffic, the traffic cannot be encapsulated properly although the GRE over GRE tunnel works for transit traffic. [PR1411874](#)
- The **file copy** command might not work if the routing-instance option is not specified. [PR1412033](#)
- On MX10003 router, the rpd process crash with switchover-on-routing-crash does not trigger the Routing Engine switchover and the rpd process on the master Routing Engine goes into STOP state. [PR1412322](#)
- Junos OS PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. [PR1412659](#)
- PPPoE subscribers might not be able to login after unified ISSU. [PR1413004](#)
- The rpd memory leak might be seen due to a wrong processing of a transient event. [PR1413224](#)
- During unified ISSU from Junos OS Release 16.1R4-S11.1 to Junos OS Release 18.2R2-S1.2, CoS GENCFG write failures are observed. [PR1413297](#)
- The support of inet6 filter attribute for ATM interface is broken in the Junos OS Release 17.2R1 and onwards [PR1413663](#)
- DHCP subscribers over HAG might cause core file generation. [PR1413862](#)
- The services load balance might not be effective for AMS if the hash-key under the **forwarding-options** hierarchy is configured. [PR1414109](#)
- FPC crash might be observed if it reaches the heap utilization limit. [PR1414145](#)
- Firewall filters are not getting programmed into Packet Forwarding Engine. [PR1414706](#)
- The user might not enter the configure mode due to mgd is in lockf status. [PR1415042](#)
- PMTU issue IPv4/IPv6 MX does not respond when MTU exceeded for clients terminated on tunnel type interfaces. [PR1415130](#)
- Port speed change and scaled aggregate Ethernet configuration can lead to MQSS errors and subsequent card crash. [PR1415183](#)
- PCE-initiated LSPs get deleted from the PCC if the PCEP session goes down and gets re-established within the configured **delegation-cleanup-timeout** period. [PR1415224](#)
- The bbe-smgd process might have memory leak while running the **show system subscriber-management route route-type <> routing-instance <>** command. [PR1415922](#)
- jdhcpd core file is observed after deletion of the active lease-query configurations. [PR1415990](#)
- BMP type 1 message with extra 24 bytes at end of the message. [PR1416301](#)

- After a GRES on a MX104 some tunnels will fail to pass traffic after a re-key. [PR1417170](#)
- The ECMP fast reroute protection feature might not work on MX5, MX10, MX40, MX80, and MX104 routers. [PR1417186](#)
- An IPv4 packet with a zero checksum might not be translated to IPv6 packet properly under NAT64 scenario. [PR1417215](#)
- With NETCONF the **xmlns** attribute is displayed twice when the RPC **get-arp-table-information** is sent to the router. [PR1417269](#)
- Some subscribers might be offline when doing GRES or daemon restart. [PR1417574](#)
- Observed zero tunnel statistics on the soft-gre tunnel. [PR1417666](#)
- The BGP session might flap after Routing Engine switchover. [PR1417966](#)
- CGNAT with MS-MPC card does not account for AP-P out of port errors or generate a syslog message when this condition is met. [PR1418128](#)
- There is no SNMP trap message generated for jnxHardDiskMissing/jnxHardDiskFailed on MX10003 routers. [PR1418461](#)
- **Clear PRBS statistics** is ineffective on latest build. [PR1418495](#)
- lsp-cleanup-timer is not being honored when lsp-cleanup-timer is configured to be greater than 2147483647. [PR1418937](#)
- PPPoE compliance issue with RFC2516, the MX allows PPPoE session-id 65535. [PR1418960](#)
- A PPP session under negotiation might be terminated if another PPPoE client bearing the same session ID. [PR1419500](#)
- CPU usage on Service PIC might spike while forming an IPsec tunnel under DEP/NAT-T scenario. [PR1419541](#)
- A new tunnel could not be established after changing the NAT mapping IP address until the IPEC SA Clear command is run. [PR1419542](#)
- rtsock_peer_unconsumed_obj_free_int: unable to remove node from list logged extensively. [PR1419647](#)
- A bbe-mibd memory leak is causing daemon crash when having live subscribers and SNMP OIDs query. [PR1419756](#)
- In the scenario where the MX Series router and the peer device both try to bring an IPsec tunnel up, so both sides are acting as an initiator, if the peer side does not answer the MX ISAKMP requests the MX can bring the peer initiated tunnel down. [PR1420293](#)
- On MX Series routers, the PTP phase is aligned but TE/cTE not good. [PR1420809](#)
- The FPC CPU might be hogged if channelized interfaces are configured. [PR1420983](#)
- Failed to reload keyadmin database for `/var/etc/keyadmin.conf`. [PR1421539](#)
- **bbemg_smgd_lock_cli_instance_db** should not log as error messages. [PR1421589](#)

- MX-VC: VCP port reports MTU value 9152 in the ICMP MTU exceeded message while the VCP port MTU is set to 9148. [PR1421629](#)
- RPM syslogs are not getting generated after deactivating the aggregate Ethernet interface. [PR1421934](#)
- Remote gateway address change is not effective on MX150 router when its an initiator. [PR1421977](#)
- The CoS IEEE-802.1 classifier might not get applied when it is configured with service activation on underlying interface. [PR1422542](#)
- On the MX204 router, the number of PICs per FPC is incorrectly used as 8, that causes MAC allocation failure on the physical interfaces. [PR1422679](#)
- Added support for SFP-T with QSA adapter in MX10003. [PR1422808](#)
- Incorrect PIC mode on MX10003 MX1RU when pic mode is changed to default mode. [PR1423215](#)
- While committing huge configuration customer is seeing the **error: mustd trace init failed** error. [PR1423229](#)
- MX10003: **enhanced-hash-key symmetric** is not effective and not shown on FPC. [PR1423288](#)
- Traffic is dropped after FPC reboot with aggregated Ethernet member links deactivated by the remote device. [PR1423707](#)
- The MPC10 line card crash is seen on Ktree alloc (jnh_dfw_instance_add (filter_index=< optimized out>)) at `../../../../src/pfe/common/applications/dfw/dfw_iff.c:1030 with inline + scale prefix filter`. [PR1423709](#)
- On MX204 optics, "SFP-1GE-FE-E-T" I2C read errors are seen when an SFP-T is inserted into a disabled state port. [PR1423858](#)
- The bbe-smgd process might crash after executing the **show system subscriber-management route prefix <>** command. [PR1424054](#)
- MX10000 port configured for 1-Gigabit flaps after a Routing Engine switchover. [PR1424120](#)
- The interface configured with 1-Gigabit speed on JNP10K-LC2101 cannot come up. [PR1424125](#)
- mgd-api core file is seen while running the gNMI set operation. [PR1424128](#)
- Continuous MAC change might cause CPU hogs and FPC reboot. [PR1424653](#)
- [vMX]Continuous disk error logs on vCP Console (Requesting switchover due to disk failure on ada1). [PR1424771](#)
- The jdhcpd might consume 100 percent CPU and then crash if **dhcp-security** is configured. [PR1425206](#)
- The rpd might crash continuously when MD5 authentication on any protocols is used along with master password. [PR1425231](#)
- Soft-gre tunnel route is lost after reboot or GRES or upgrade in WAG scenario. [PR1425237](#)
- Log messages are seen continuously on MX204 router **fru_is_present: out of range slot 0 for**. [PR1425411](#)
- All interfaces creation fails after NSSU. [PR1425716](#)
- Sometimes, the interface is down after rebooting. [PR1426349](#)

- Traffic loss might be seen when multiple IPsec tunnels are established with the remote peer. [PR1426975](#)
- Traffic is not flowing through MACsec interfaces when configured with an unknown cipher algorithm and change back. [PR1427294](#)
- Execution of the **clear-session re-cli** command should not be allowed from Standby DUT. [PR1428353](#)
- The subscriber IP route might get stuck in bbe-smgd if the subscriber IP address is the same with local IP address. [PR1428428](#)
- Incorrect normalization on routing instance where an interface includes a **vlan-id-range**. [PR1428623](#)
- PTSP subscriber is stuck in configured state. Auto-clear-timer does not work as well. [PR1428688](#)
- Incorrect IGMP statistics for dynamic PPP interfaces are observed. [PR1428822](#)
- L2TP subscriber and MPLS Pseudowire Subscriber volume accounting statistics value remains unchanged post unified ISSU. [PR1429692](#)
- The rpsd daemon is not getting killed on when unconfigured simulatenous to toggling rpd 'force-64-bit', rpsd core file is seen 10 minutes later. [PR1429770](#)
- **Cmerror Op set** log message is missing for **bringup jspec** command-based error simulation in EVO. [PR1430300](#)
- Configuration is prevented from being applied on MX Series routers in subscriber scenario. [PR1430360](#)
- Destination unreachable counter is counting up without receiving traffic. [PR1431384](#)
- The bbe-smgd process might crash if PPPoE subscribers are trying to log in when commit is in progress. [PR1431459](#)
- MX10003 - PEM not present alarm is raised when minimum required PEM exist in the system. [PR1431926](#)
- Error message for **show system resource-monitor** and **show system resource-cleanup** is **error: command is not valid on the qfx5220-32cd**. [PR1435136](#)
- A unified ISSU fails from Junos OS Release 19.1R1 legacy Junos OS release images. [PR1438144](#)

Infrastructure

- SNMP OID IFOutDiscards is not updated when drops increase. [PR1411303](#)
- Increase in Junos image size for Junos OS Release 19.1R1. [PR1423139](#)

Interfaces and Chassis

- LFM sessions might flap during unified ISSU. [PR1377761](#)
- Changing the value of **mac-table-size** to default might lead all FPC to reboot. [PR1386768](#)
- The dcd memory leak might be seen when committing configuration change on static route tag. [PR1391323](#)
- The dcd crash might be seen after deleting the sub interface from VPLS routing-instance and mesh-group. [PR1395620](#)

- NPC crashes at `rt_nh_install (rnh=0x618123d8, rnh_src=0x0, rt=< optimized out>, p_rtt=0x74f886c0)` at `.././../src/pfe/common/pfe-arch/trinity/applications/route/rt_nh.c:631`. [PR1396540](#)
- Static demux0 logical interfaces do not come up after a configuration change if the underlying interface is et. [PR1401026](#)
- Certain otn-options cause interface flapping during commit. [PR1402122](#)
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. [PR1402606](#)
- The subscriber might not be able to access the device due to the conflicted assigned address. [PR1405055](#)
- On MX Series routers, the EX-SFP-1FE-LX SFP transceiver does not initialize with MIC-3D-20GE-SFP-E(EH). [PR1405271](#)
- The `cfmd` might fail to start after it is restarted. [PR1406165](#)
- The **aaa-options** configuration statement for PPPoE subscribers does not work on the MX80 and MX104 routers. [PR1410079](#)
- OAM CFM MEP flaps might occur when hardware-assisted keep alives are enabled. [PR1417707](#)
- **Monitor Ethernet loss-measurement** command returns Invalid ETH-LM request for unsupported outgoing logical interface. [PR1420514](#)
- Incorrect value on speed will cause traffic destined to the IRB's VIP to be dropped. [PR1421857](#)
- The syslog message `/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex` is seen on executing LFM related configuration commit on the aggregated Ethernet interfaces. [PR1423586](#)
- [EVPN] Aggregate Ethernet interface flaps followed by commit. [PR1425339](#)
- **flexible-queuing-mode** is not working on MPC5E of Virtual Chassis member1. [PR1425414](#)
- PEMs lose DC output power load sharing after PEM switch off and on operation on MX routers. [PR1426350](#)
- CFM message flooding. [PR1427868](#)
- Vrrpd crashes during group mastership change if preemption is configured and logical interface was enabled/activated some time after disabling/deactivation. [PR1429906](#)

Layer 2 Features

- The unicast traffic from IRB interface towards LSI might be dropped due to Packet Forwarding Engine mismatch at egress processing. [PR1381580](#)
- Traffic loss might be seen over LDP-VPLS scenario. [PR1415522](#)
- The `rpd` crashes after `iw0` interface is configured under a VPLS instance. [PR1406472](#)
- In a Layer 2 domain, there might be unexpected flooding of unicast traffic at every 32-40 seconds interval towards all local CE-facing interface. [PR1406807](#)

- Broadcast traffics might be discarded in a VPLS local-switching scenario. [PR1416228](#)
- Commit error will be seen but the commit is processed if adding more than one site under **protocols vpls** in the VPLS routing-instances. [PR1420082](#)

Layer 2 Ethernet Services

- The SNMP query on LACP interface might lead to lacpd crash. [PR1391545](#)
- Log messages **dot1xd[]: task_connect: task ESP CLIENT:....: Connection refused** might be reported in Junos OS Release 17.4 or later. [PR1407775](#)
- DMAC problem of IRB interface for traffic over the Layer 2 circuit. [PR1410970](#)
- The IRB interface might flap after committing configuration change on any interface. [PR1415284](#)
- The IPv6 neighbor might become unreachable after the primary link goes down in a VPLS scenario. [PR1417209](#)
- The jdhcpd becomes aware about some of the existing configuration only after 'commit full' or jdhcpd restart. [PR1419437](#)
- Change the nd6 next hops to reject NH once Layer 2 interfaces gets disassociated with IPv6 entries. [PR1419809](#)
- The jdhcpd process might consistently run at 100 percent CPU and not provide service if **delay-offer** is configured for the DHCP local server. [PR1419816](#)
- JDI-RCT:BBE:DHCP subscribers on non-default routing instance went down after unified ISSU. [PR1420982](#)
- The jdhcpd daemon might crash during continuous stress test. [PR1421569](#)

MPLS

- Not found number of ingress, transit, and egress LSP's as expected. [PR1242558](#)
- Collecting LDP statistics do not work correctly and kernel memory leak is observed after configuring **ldp traffic-statistics**. [PR1258308](#)
- With an SR-TE path with "0" explicit NULL as the innermost label, SR-TE path does not get installed with label "0". [PR1287354](#)
- A RSVP-signaled LSP might stay in down state after a link in the path flaps. [PR1384929](#)
- The rpd process might crash when executing **traceroute mpls bgp**. [PR1399484](#)
- MPLS LSP traffic loss might be seen under rare conditions if CSPF is enabled. [PR1402382](#)
- Scaled MPLS labels might cause slow labels allocation and high CPU utilization. [PR1405033](#)
- The Layer 2 circuit information is not advertised over the LDP session if **ldp dual-transport inet-lsr-id** is different from the router-id. [PR1405359](#)
- Resources might be reserved for stale RSVP LSP when RSVP is disabled on the interface. [PR1410972](#)

- The rpd might crash in BGP-LU with egress-protection while committing configuration changes. [PR1412829](#)
- The rpd might crash if **longest-match** is configured for LDP. [PR1413231](#)
- LDP route is not present in inet6.3 if IPv6 interface address is not configured. [PR1414965](#)
- Rpd memory might leak when RSVP LSP is cleared/re-signaled. [PR1415774](#)
- RSVP signalled LSP takes 3 - 4 minutes before LSP switchover begins, causing long traffic to be silently discarded. [PR1416487](#)
- LDP route might be missing in inet.3 when enabling **sr-mapping-client** on LDP-SR stitching node. [PR1416516](#)
- Traffic might be dropped because of the LDP label corruption after Routing Engine switchover. [PR1420103](#)
- Bad length for Sub-TLV 34 (RFC 8287) in MPLS echo request. [PR1422093](#)
- LDP route metric might not match IGP route metric even with **ldp track-igp-metric** configured. [PR1422645](#)
- Bypass dynamic RSVP LSP tears down too soon when being used for protecting LDP LSP with **dynamic-rsvp-lsp** statement. [PR1425824](#)
- MPLS ping sweep stops working and gets CLI irresponsive. [PR1426016](#)
- When MBB for P2MP LSP fails, it is stuck in the old path. [PR1429114](#)

Network Management and Monitoring

- The chassisd might crash and restart after the AGENTX session timeout between master(snmpd) and sub-agent. [PR1396967](#)
- The snmp query might not get data in scaled L2 circuit environment. [PR1413352](#)
- Syslog filtering(match "regular-expression" statement) does not work if each line of **/etc/syslog.conf** is over 2048 bytes. [PR1418705](#)

Platform and Infrastructure

- The kernel and ksyncd core after dual cb flap at rt_nhfind_params: rt_nhfind() found an nh different from that onmaster 30326. [PR1372875](#)
- Traffic is being dropped when passing through MS-DPC to MPC. [PR1390541](#)
- All FPCs might restart after the Layer 3 VPN routes churn multiple times. [PR1398502](#)
- MAP-E some ICMP types cannot be encapsulated or decapsulated on the SI interface. [PR1404239](#)
- Abnormal queue-depth counters are seen in the **show interface queue** command output on interfaces that are associated to XM2 and 3. [PR1406848](#)
- IPv6 traffic might be dropped between VXLAN bridge-domain and IP/MPLS network. [PR1407200](#)
- CoS configuration changes might lead to traffic drop on cascade port in a Junos fusion setup. [PR1408159](#)

- Traffic is getting dropped when there is a combination of DPC/FPC card and MPC card on egress PE router in Layer 3 VPN. [PR1409523](#)
- The VLAN tag is wrongly inserted on the access interface if the packet is sent from an IRB interface. [PR1411456](#)
- The MPC might crash when one MIC is pulled out while the MIC is booting up. [PR1414816](#)
- Distributed multicast forwarding to the subscriber interface might not work. [PR1416415](#)
- The **op url** command cannot run a script with libs from **/config/scripts**. [PR1420976](#)
- arp request is not replied although **proxy-arp** is configured. [PR1422148](#)
- **show jnh trap-info** with incorrect LU instance crashes and generates a core file on FPC. [PR1423508](#)
- The native VLAN ID of packets might fail to be removed when leaving out. [PR1424174](#)
- The policer bandwidth might be incorrect for the aggregated Ethernet interface after activating the **shared-bandwidth-policer** statement. [PR1427936](#)
- Pre-fragmented ICMP IPv4 packets might fail to arrive at the destination. [PR1432506](#)
- Enable sensor **/junos/system/linecard/qmon/** causing continuous **ppe_error_interrupt** errors. [PR1434198](#)

Routing Policy and Firewall Filters

- The rpd process might crash when the policy configuration is being changed. [PR1357802](#)
- MX-Series: The CLI statement **as-path-expand last-as** causes commit failures. [PR1388159](#)
- The rpd process might crash when **routing-options flow** configuration is removed. [PR1409672](#)

Routing Protocols

- Dynamic NextHop template cache does not shrink when application frees the NextHop template. [PR1346984](#)
- Qualified next hop of static route might not be withdrawn when BFD is down. [PR1367424](#)
- The static route might persist even after its BFD session goes down. [PR1385380](#)
- BGP sessions might keep flapping on backup RE if proxy-macip-advertisement is configured on IRB interface for EVPN-VXLAN. [PR1387720](#)
- In rare cases, rpd process might crash after Routing Engine switchover when BGP multipath and L3VPN **vrf-table-label** are configured [PR1389337](#)
- BGP IPv6 routes with IPv4 next hop causes a rpd crash. [PR1389557](#)
- Multicast traffic might be interrupted in some H-VPLS scenario. [PR1394213](#)
- BGP DMZ LINK BANDWIDTH - not able to aggregate bandwidth, when applying the policy. [PR1398000](#)
- The process rpd might crash in a BGP setup with NSR enabled. [PR1398700](#)

- Unexpectedly high packet loss might be observed after an uplink failure when the MoFRR feature is used in a scaled environment. [PR1399457](#)
- There might be unexpected packets drop in MoFRR scenario if active RPF path is disabled. [PR1401802](#)
- The rpd might be stuck at 100 percent when **auto-export** and **BGP add-path** are configured. [PR1402140](#)
- BGP router on the same broadcast subnet with its neighbors might cause IPv6 routing issue on the neighbor from other vendors. [PR1402255](#)
- Some times when a new logical router is configured, the logical router core might be seen on the system. [PR1403087](#)
- Memory leaks when labeled-isis transit routes is created as a chain composite next hop. [PR1404134](#)
- Extended traffic loss might be seen after link recovery when **source-packet-routing** is used on OSPF P2P links. [PR1406440](#)
- SBFD failure is seen with a special IP address like 127.0.0.1 under interface lo0. [PR1406631](#)
- IGMP join through PPPOE sub is not propagated to upstream PIM. [PR1407202](#)
- The rpd crashes on static route configuration for multicast source. [PR1408443](#)
- On MX Series routers, mcsnoopd core file is generated immediately after the commit change related to EVPN-VXLAN configuration. [PR1408812](#)
- SID label operation might be performed incorrectly in an OSPF SPRING environment. [PR1413292](#)
- An unexpected AS prepending action for AS path might be seen after the **no-attrset** statement is configured or deleted with vrf-import/vrf-export configuration. [PR1413686](#)
- The CPU utilization of the rpd process is stuck at 100 percent if BGP multipath is configured. [PR1414021](#)
- Dynamic routing protocol flaps with vmhost Routing Engine switchover on Next Generation-Routing Engine. [PR1415077](#)
- The IS-IS SR route sent by the mapping server might be broken for ECMP. [PR1415599](#)
- Route info might be inconsistent between RIB and OSPF database when using the OSPF LFA feature. [PR1416720](#)
- A memory leak in rpd might be seen if source packet routing is enabled for the IS-IS protocol. [PR1419800](#)
- IPv6 IS-IS routes might be deleted and not be reinstalled when MTU is changed under the logical interface level for family inet6. [PR1420776](#)
- A timing issue is seen while closing a PIM task and an auto-RP at the same time that might sometimes result in an rpd core file generation. [PR1426711](#)
- The rpd might crash while handling the withdrawal of an imported VRF route. [PR1427147](#)
- The rpd process might crash with OSPF overload as external configuration. [PR1429765](#)
- The **request system core-dump routing** CLI is not supported in cRPD. [PR1433349](#)

Services Applications

- Hide HA information when the service set does not have HA configured. [PR1383898](#)
- The following log message is seen: **SPD_CONN_OPEN_FAILURE: spd_svc_set_summary_query: unable to open connection to si-0/0/0 (No route to host)**. [PR1397259](#)
- Inconsistent content might be observed to the access line information between ICRQ and PPPoE. message [PR1404259](#)
- The stale si- logical interface might be seen when L2TP subscribers with duplicated prefixes or framed-route log in. [PR1406179](#)
- The kmd process might crash on MX/ACX platforms when IKEv2 is used. [PR1408974](#)
- The ERA value does not match with configured values while verifying if the new ERA settings are reflected in messages log. [PR1410783](#)
- The jpppd generates core files on LNS. [PR1414092](#)
- L2TP LAC might fail to tunnel static pp0 subscriber to the desired LNS. [PR1416016](#)
- IPsec SA might not come up when the local gateway address is a VIP for a VRRP configured. interface. [PR1422171](#)
- In subscriber with L2TP scenario, subscribers are stuck in INIT state forever. [PR1425919](#)
- Some problems might be seen if client negotiates LCP with no ppp-options to LAC. [PR1426164](#)

Software Installation and Upgrade

- The configuration loss and traffic loss might be seen if the backup Routing Engine is zeroized and is then switched over to master within short time. [PR1389268](#)
- JSU might be deactivated from FPC in case of power cycle. [PR1429392](#)

Subscriber Access Management

- The DHCPv6-PD client connection might be terminated after commit when the RADIUS assigned address is not defined within the range of a local pool. [PR1401839](#)
- The authd crash might be seen due to a memory corruption issue. [PR1402012](#)
- Adding a firewall filter service through the **test aaa** command causes a crash in dfwd. [PR1402051](#)
- The authd re-uses address too quickly before jdhcpd completely cleans up the old subscriber that is causing the flooding error log **DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add x.x.x.x as it is already used by xxx**. [PR1402653](#)
- Continuous log message **authd[18454]: %DAEMON-3-LI: liPollTimerExpired returned 0**. [PR1407923](#)
- Authd telemetry: Linked pool head attribute is incorrect for single pools. [PR1413293](#)
- CoA-NACK is not sent when performing negative COA request tests by sending incorrect session-id. [PR1418144](#)

- Subscribers might not be able to re-login in Gx-plus provisioning scenario. [PR1418579](#)
- PPPoE session might be disconnected when LI attributes are received in access-accept with invalid data. [PR1418601](#)
- Address allocation issue with linked pools when using linked-pool-aggregation. [PR1426244](#)
- RADIUS authentication server might always be marked as DEAD. [PR1429528](#)

User Interface and Configuration

- The **show configuration** and **rollback compare** commands are causing high CPU usage. [PR1407848](#)
- Commit reject occurs for ae0.0 vlan-id-list and routing-instance vlan-id (but does not reject for vlan-range). [PR1427278](#)

VPNs

- The receivers belonging to a routing instance might not receive multicast traffic in an Extranet next-generation MVPN scenario. [PR1372613](#)
- Downstream interface is not removed from multicast route after getting PIM prune. [PR1398458](#)
- Routes with multiple communities are being rejected in an inter-AS next-generation MVPN scenario. [PR1405182](#)
- For rosen MVPN configuration with data-mdt, the **show pim mdt data-mdt-limit instance < instance name>** CLI command with family option causes high CPU usage of the rpd. [PR1405887](#)
- The rpd might crash in rosen MVPN scenario when the same provider tunnel source address is being used for both IPv4 and IPv6. [PR1416243](#)
- The deletion of (S,G) entry might be skipped after the PIM join timeout. [PR1417344](#)
- The rpd process might crash in rare conditions when Extranet next-generation MVPN is configured. [PR1419891](#)
- A permanent traffic loss is seen on next-generation MVPN selective tunnels after Routing Engine switchover (one-time). [PR1420006](#)
- The rpd process might crash and core file is generated during **mpls ping** command on L2 circuit. [PR1425828](#)

SEE ALSO

[What's New | 92](#)

[What's Changed | 113](#)

[Known Limitations | 123](#)

[Open Issues | 127](#)

Documentation Updates

IN THIS SECTION

- [Installation and Upgrade Guide | 191](#)
- [Subscriber Management Provisioning Guide | 191](#)

This section lists the errata and changes in Junos OS Release 19.2R3 documentation for the MX Series.

Installation and Upgrade Guide

- **Veriexec explained (MX Series)**—Verified Exec (also known as veriexec) is a file-signing and verification scheme that protects the Junos operating system (OS) against unauthorized software and activity that might compromise the integrity of your device. Originally developed for the NetBSD OS, veriexec was adapted for Junos OS and enabled by default from Junos OS Release 7.5 onward.

[See [Veriexec Overview](#).]

Subscriber Management Provisioning Guide

- The Broadband Subscriber Sessions User Guide published for Junos OS Release 19.2R1 erroneously reported support for a feature to manage certain PCRF server errors. Support for an extended session ID was also incorrectly reported. The incorrect information has been removed from the affected topics.
- The Broadband Subscriber Sessions User Guide published for Junos OS Release 19.2R1 reported that the **juniper-access-line-attributes** option is backward compatible. This option is not backward compatible with Junos OS Release 19.1 or earlier releases. This means that if you have configured **juniper-access-line-attributes** option in Junos OS Release 19.2 or higher releases, you must perform the following steps to downgrade to Junos OS Release 19.1 or earlier releases:
 1. Delete the **juniper-access-line-attributes** option from all access profiles that include it.
 2. Perform the software downgrade.
 3. Add the **juniper-dsl-attributes** option to the affected access profiles.

SEE ALSO

What's New 92
What's Changed 113
Known Limitations 123
Open Issues 127
Resolved Issues 139
Migration, Upgrade, and Downgrade Instructions 192

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.2 | 193](#)
- [Procedure to Upgrade to FreeBSD 11.x based Junos OS | 193](#)
- [Procedure to Upgrade to FreeBSD 6.x based Junos OS | 195](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 197](#)
- [Upgrading a Router with Redundant Routing Engines | 198](#)
- [Downgrading from Release 19.2 | 198](#)

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

Starting in Junos OS 18.3R1 release, FreeBSD 11.x is the underlying OS for all Junos OS platforms which were previously running on FreeBSD 10.x based Junos OS. FreeBSD 11.x does not introduce any new Junos OS related modifications or features but is the latest version of FreeBSD.

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

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

Basic Procedure for Upgrading to Release 19.2

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-x86-32-19.2R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.2R3.9-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-x86-32-19.2R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-19.2R3.9-limited.tgz
```

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

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

- `http://hostname/pathname`
- `scp://hostname/pathname`

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 19.2 **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-19.2R3.9-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-19.2R3.9-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 19.2 **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.

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 19.2

To downgrade from Release 19.2 to another supported release, follow the procedure for upgrading, but replace the 19.2 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 92
What's Changed 113
Known Limitations 123
Open Issues 127
Resolved Issues 139
Documentation Updates 191

Junos OS Release Notes for NFX Series

IN THIS SECTION

- What's New | 199
- What's Changed | 201
- Known Limitations | 202
- Open Issues | 204
- Resolved Issues | 206
- Documentation Updates | 210
- Migration, Upgrade, and Downgrade Instructions | 210

These release notes accompany Junos OS Release 19.2R3 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- What's New in Release 19.2R3 | 200
- What's New in Release 19.2R2 | 200
- What's New in Release 19.2R1 | 200

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

What's New in Release 19.2R3

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 19.2R3.

What's New in Release 19.2R2

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 19.2R2.

What's New in Release 19.2R1

Architecture

- **Open vSwitch (OVS) integrated with Data Plane Development Kit (DPDK)**—Starting in Junos OS Release 19.2R1, NFX150-S1 and NFX150-S1E devices support OVS integrated with DPDK that offers better network packet throughput and lower latencies.

[See [Benefits and Uses of NFX150](#).]

Application Security

- **Application-based multipath support (NFX Series)**—Starting in Junos OS Release 19.2R1, application-based multipath routing is supported on NFX150 devices.

Multipath routing allows the sending device to create copies of packets, send each copy through two or more WAN links. On the other end, multipath calculates the jitter and packet loss for the combined links and estimates the jitter and packet loss for the same traffic on individual links. You can compare the reduction in packet loss when combined links instead of individual links are used. Sending multiple copies of traffic ensures timely delivery of the sensitive application traffic.

Multipath support in SD-WAN uses case enhances application experience.

[See [Application Quality of Experience](#).]

- **Application-level logging for AppQoE (NFX Series)**—Starting in Junos OS Release 19.2R1, NFX series devices support application-level logging for AppQoE. This feature reduces the impact on the CSO or log collector device while processing a large number of system log messages generated at the session-level. The device maintains session-level information and provides system log messages for the session level. Replacing session-level logging with application-level logging decreases the overhead on the device and increases AppQoE throughput.

[See [AppQoE](#).]

Virtual Network Functions

- **Disable VNF interfaces (NFX150 and NFX250 NextGen)**—Starting in Junos OS Release 19.2R1, you can manually disable the VNF interfaces (eth0 through eth9) on the OVS or custom bridge on NFX150 and NFX250 NextGen devices.

[See [Configuring VNF Interfaces and VLANs.](#)]

- **MAC flooding on VNF interfaces (NFX150 and NFX250 NextGen)**—Starting in Junos OS Release 19.2R1, changes to the default MAC flooding behavior of the virtualized network function (VNF) interfaces improve the performance of multicast traffic. If a VNF interface is not attached to a VLAN, drop flow is not configured. The interface functions as a trunk port that can receive and forward the VLAN traffic.

In earlier releases, if a VNF interface is not attached to a VLAN, drop flow is configured and the VNF interface drops the outgoing traffic.

[See [Configuring VNF Interfaces and VLANs.](#)]

- **Bootstrap configuration of a VNF using a config-drive (NFX150 and NFX250 NextGen)**—Starting in Junos OS Release 19.2R1, you can bootstrap a VNF using an attached config drive that contains a bootstrap-config ISO file on NFX150 and NFX250 NextGen devices. The config drive is a virtual drive, which can be a CD-ROM, USB drive or Disk drive associated to a VNF with the configuration data. Configuration data can be files or folders, which are bundled in the ISO file that makes a virtual CD-ROM, USB drive, or Disk drive.

[See [Preparing the Bootstrap Configuration on NFX150 Devices.](#)]

[See [Preparing the Bootstrap Configuration on NFX250 NextGen Devices.](#)]

SEE ALSO

[What's Changed | 201](#)

[Known Limitations | 202](#)

[Open Issues | 204](#)

[Resolved Issues | 206](#)

[Documentation Updates | 210](#)

[Migration, Upgrade, and Downgrade Instructions | 210](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.2R3 | 202](#)
- [What's Changed in Release 19.2R2 | 202](#)
- [What's Changed in Release 19.2R1 | 202](#)

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

What's Changed in Release 19.2R3

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 19.2R3 for NFX Series devices.

What's Changed in Release 19.2R2

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 19.2R2 for NFX Series devices.

What's Changed in Release 19.2R1

Factory-default Configuration

- **Plug-and-play configuration (NFX150 and NFX250 NextGen devices)**—Starting in Junos OS Release 19.2R1, the factory default configuration is modified to include the secure router plug-and-play configuration.

SEE ALSO

What's New 199
Known Limitations 202
Open Issues 204
Resolved Issues 206
Documentation Updates 210
Migration, Upgrade, and Downgrade Instructions 210

Known Limitations

IN THIS SECTION

- [Interfaces | 203](#)
- [Platform and Infrastructure | 203](#)
- [Virtual Network Functions \(VNFs\) | 204](#)

Learn about known limitations in this release for NFX Series devices.

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

Interfaces

- On NFX150 devices, the TCP and ICMP RPM probes take the best-effort queue of the outgoing interface, instead of the network control queue. As a workaround, configure a DSCP value such as nc1 for the RPM probes to take the network control queue. [PR1329643](#)
- On NFX150 devices, the PPPoE session does not come up on the interface due to the hardware limitation for both tagged and untagged cases. As a workaround, enable the promiscuous mode on the interface. [PR1347830](#)
- On NFX150 devices, the link does not come up if a 1-GbE 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](#)

Platform and Infrastructure

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

As a workaround, interrupt the boot process to select the primary disk. [PR1344342](#)

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

Virtual Network Functions (VNFs)

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

SEE ALSO

[What's New | 199](#)

[What's Changed | 201](#)

[Open Issues | 204](#)

[Resolved Issues | 206](#)

[Documentation Updates | 210](#)

[Migration, Upgrade, and Downgrade Instructions | 210](#)

Open Issues

IN THIS SECTION

- [Interfaces | 205](#)
- [Platform and Infrastructure | 205](#)
- [Routing Protocols | 205](#)
- [Virtual Network Functions \(VNFs\) | 206](#)

Learn about open issues in this release for the 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 NFX150 and NFX250 NextGen devices, the **link disable** option puts the analyzer interface in inconsistent state with the link state as **down** and admin state as **up**. [PR1442224](#)
- On NFX150 and NFX250 NextGen devices, while configuring **vmhost vlans** using **vlan-id-list**, the system allows duplicate VLAN IDs in the VLAN ID list. [PR1438907](#)
- On NFX250 devices, if the IRB interface configuration and DHCP service configuration on JDM are removed and rolled back while retaining the VLAN mapping to the IRB interface, the DHCP service fails to assign IP addresses to the corresponding VNF interfaces and the service chaining fails. As a workaround, remove the VLAN mapping to the IRB interface along with IRB and DHCP service configuration on JDM. [PR1234055](#)
- On NFX150 devices, only the CFM cells that are configured for MEP levels are exchanged across xDSL MEPs. Other MEP-level CFM packets are dropped, whereas for Ethernet All MD level along with above level will be exchanged. [PR1409576](#)
- When you issue the **show interface** command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. [PR1306191](#)
- When a DHCP server assigns a conflicting IP address to the NFX Series device interfaces, the device does not send a **DHCP DECLINE** message in response. [PR1398935](#)
- 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](#)
- When the interface configuration has the encapsulation **flexible-ethernet-services** enabled on a 10-Gigabit Ethernet interface, traffic is dropped. [PR1425927](#)

Platform and Infrastructure

- On NFX250 devices, Virtual Port Peer (VPP) is not running on dual CPE and occasionally on single CPE. [PR1461238](#)
- On NFX150 devices, the following messages are seen during **FTP: ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. [PR1315605](#)
- Login access to JDM through TACACS failed after upgrade to Junos OS Release 18.4R3 As a workaround, log in as a local user. [PR1504915](#)

Routing Protocols

- When a static route and an OSPF route are active in the routing table for a specific destination network, a ping initiated to that destination network from the NFX Series device will fail. [PR1438443](#)
- On NFX150 devices, SNMP does not work for the following commands:

- `show snmp mib walk jnxIpSecTunMonOutEncryptedBytes`
- `show snmp mib walk jnxIpSecTunMonOutEncryptedPkts`
- `show snmp mib walk jnxIpSecTunMonInDecryptedBytes`
- `show snmp mib walk jnxIpSecTunMonInDecryptedPkts`
- `show snmp mib walk jnxIpSecTunMonLocalGwAddr`
- `show snmp mib walk jnxIpSecTunMonLocalGwAddrType`

[PR1386894](#)

Virtual Network Functions (VNFs)

- 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](#)
- On issuing the `show virtual-network-functions vnf-name` command on NFX250 devices, the system creates a defunct process due to a mismatch between the `popen()` and `pclose()` calls. [PR1415210](#)

SEE ALSO

What's New 199
What's Changed 201
Known Limitations 202
Resolved Issues 206
Documentation Updates 210
Migration, Upgrade, and Downgrade Instructions 210

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3 | 207](#)
- [Resolved Issues: 19.2R2 | 207](#)
- [Resolved Issues: 19.2R1 | 209](#)

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

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

Resolved Issues: 19.2R3

Application Security

- On NFX150 devices, 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](#)

Interfaces

- On NFX350 devices, deleting and adding SXE interfaces results in moving the SXE interface to STP BLK state and traffic drops on that interface [PR1475854](#)

Platform and Infrastructure

- In Junos OS Release 18.4R3.3, LTE module is not contained on platform NFX250-ATT-LS1. [PR1493711](#)
- Traffic tunnels are down in GWR, after jdm image upgrade from Junos OS Release D497.1 to 18.4R3-S2 porter2 image [PR1507165](#)

Resolved Issues: 19.2R2

Class of Service (CoS)

- Traffic is sent to the incorrect queue when you are configuring COS with forwarding-classes class versus queue. [PR1436408](#)
- The CoS rewrite rule does not work for the st0 interface. [PR1439401](#)

High Availability

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

Interfaces

- On NFX150 devices, if VLAN-tagged PPPoE is configured and you commit the configuration without using the **delete interfaces** command, then PPPoE does not come up and we will see malformed offer packet at CPE. [PR1409475](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- The temperature field is displayed as **Testing** in the **show chassis fpc** and **show chassis fpc details** CLI command output. [PR1433221](#)

- The limit on the maximum OVS interfaces is restored to the originally defined limit 25 for backward compatibility. [PR1439950](#)
- On NFX Series devices, ping does not work between the cross-connected interfaces configured with the **interface deny-forwarding** option. [PR1442173](#)
- On NFX150 and NFX250 NextGen devices, cross-connect stays down even if all linked interfaces are up. [PR1443465](#)
- On NFX250 NextGen devices, when you change the performance mode from throughput to compute, the FPC0 interface goes down after the reboot. [PR1448246](#)
- The heth-0-4 and heth-0-5 ports do not detect traffic when you try to activate the ports by plugging in or unplugging the cable. [PR1449278](#)
- On NFX Series devices, the static MAC address is replaced by a random MAC address. [PR1458554](#)
- After upgrading vSRX3 from Junos OS Release 18.4R1.8 to Junos OS Release 18.4R1-S4 to solve swap memory issues, core files are generated and traffic is dropped. [PR1465132](#)
- On NFX150 devices, GRE tunnel interface (gr-1/0/0) might not appear if the **clear-dont-fragment-bit** option is configured for the GRE interface. [PR1472029](#)

Layer 2 Ethernet Services

- DHCP request may be dropped in DHCP relay scenario. [PR1435039](#)

Platform and Infrastructure

- REST API process becomes nonresponsive when a number of requests come in at a high rate. [PR1449987](#)
- Packet drops, replication failure, or ksyncd crashes might be seen on the logical system of a device running Junos OS after Routing Engine switchover. [PR1427842](#)
- LACP state might remain in **Attached** state after peer active members are disabled. [PR1439268](#)
- On NFX150 devices, when you issue the **show security dynamic-address** command, the **show security dynamic-address does NOT work on port3** message is displayed. [PR1448594](#)
- On NFX150 devices, version compare in PHC might fail, making the PHC download the same image. [PR1453535](#)
- NFX250 devices do not allow jdm (case-insensitive) as a VNF name. You can use jdm as part of the name. For example, *jdm123*, *abcJDM*, and *abcJDM123* are valid VNF names, whereas, *jdm*, *JDM*, *Jdm*, *JDm* are not valid VNF names. [PR1463963](#)
- On NFX Series devices, after a power outage, JDMD might become unresponsive because the **/etc/hosts** file is corrupted. [PR1477151](#)

Routing Protocols

- The **other querier present interval** timer cannot be changed in an IGMP/MLD snooping scenario. [PR1461590](#)

- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

Virtual Network Functions (VNFs)

- Duplicate host entries are observed in `/etc/hosts` on JDM when VNF interfaces are moved from the default OVS bridge to a custom OVS bridge. [PR1434679](#)
- When you downgrade from Junos OS Release 19.2 to Junos OS Release 18.4, the **show virtual-network-functions vnf-name** command does not display the VNF information. [PR1437547](#)
- No error is displayed for **native-vlan-id** that is configured on an access VNF interface, although the commit fails. [PR1438854](#)
- Management ports are not disabled with the **link disable** command on NFX150-S1 devices. [PR1442064](#)

Resolved Issues: 19.2R1

Interfaces

- On NFX250 devices, an SFP-T interface does not become active when it is plugged into a ge-12/0/0 or a ge-13/0/0 interface. [PR1404756](#)
- On a NFX250 devices with xDSL SFP used on the fiber ports the status of the xDSL SFP was displayed with **Adsl Status** field under cli command **show interfaces int-name**. But whenever a user hot-swaps a xDSL SFP with another xDSL SFP on the same port, then the **Adsl Status** field was not displayed in the **show interfaces** command output. [PR1408597](#)
- On NFX150 devices, FPC0 may not be online after an upgrade and a device reboot is required. [PR1430803](#)

Platform and Infrastructure

- Software upgrade does not delete all images from a previous installation. This occupies about 1GB of storage per upgrade and leads to depletion of storage after several upgrades. [PR1408061](#)
- JDM depends on the libvirt daemon to manage the guest VM through cli. The libvirt daemon was stuck and vjunos VM start up failed, which resulted in in-band connectivity issues, the guest VM could not start, and the console was hung. [PR1314945](#)
- The **NFX3/ACX5448:LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** message is displayed when you commit the configuration on config prompt. As a workaround to exclude this from messages or syslogs, run the **set system syslog user * match "(LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified)"** and commit. [PR1376665](#)

SEE ALSO

What's New	 199
What's Changed	 201
Known Limitations	 202
Open Issues	 204
Documentation Updates	 210
Migration, Upgrade, and Downgrade Instructions	 210

Documentation Updates

There are no errata or changes in Junos OS Release 19.2R3 for documentation for NFX Series.

SEE ALSO

What's New	 199
What's Changed	 201
Known Limitations	 202
Open Issues	 204
Resolved Issues	 206
Migration, Upgrade, and Downgrade Instructions	 210

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 211
- [Basic Procedure for Upgrading to Release 19.2](#) | 211

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 19.2

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: 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 19.2R3:

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 | 199](#)

[What's Changed | 201](#)

[Known Limitations | 202](#)

[Open Issues | 204](#)

[Resolved Issues | 206](#)

[Documentation Updates | 210](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- [What's New | 213](#)
- [What's Changed | 221](#)

- Known Limitations | 227
- Open Issues | 228
- Resolved Issues | 230
- Documentation Updates | 238
- Migration, Upgrade, and Downgrade Instructions | 238

These release notes accompany Junos OS Release 19.2R3 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- What's New in 19.2R3 | 213
- What's New in 19.2R2 | 214
- New and Changed Features: 19.2R1-S4 | 214
- New and Changed Features: 19.2R1-S1 | 215
- New and Changed Features: 19.2R1 | 215

Learn about new features introduced in the main and maintenance releases for PTX Series routers.

What's New in 19.2R3

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 19.2R3.

What's New in 19.2R2

Network Management and Monitoring

- **Implement new MIBs using telemetry-based model (PTX Series)**—Starting in Junos OS Release 19.2R2, new MIBs `mplsMldpInterfaceStatsEntry` and `mplsMldpFecUpstreamSessTable` are introduced. The Routing Engine uses a telemetry-based approach to collect statistics to provide MIB data for these MIBs. A new statement, **sensor-based-stats** at the `[edit protocols ldp traffic-statistics]` hierarchy level, enables telemetry-based collection. Setting this statement is mandatory for `mplsMldpInterfaceStatsEntry` and `mplsMldpFecUpstreamSessTable`.

Routing Protocols

- **Option to pause BGP multipath computation during BGP peering churn (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R2, you can choose to defer multipath computation for all families during a BGP peering churn. In very large-scale network deployments during BGP peering churn there is a temporary spike in multipath computation, which takes a toll on the Packet Forwarding Engine resources. This feature allows you to pause the multipath computation and to resume after the peering churn settles down. Note that if there is no BGP peering churn, then multipath computation is not paused.

To enable the pause option for BGP multipath computation during BGP peering churn, include the **pause computation** statement at the `[edit protocols BGP multipath]` hierarchy level.

New and Changed Features: 19.2R1-S4

Interfaces and Chassis

- **Support for 1-Gbps speed on QFX-60S line card on PTX10008 and PTX10016 Routers**—Starting in Junos OS Release 19.2R1-S4, QFX10000-60S-6Q line card supports 1-Gbps speed on its ports (0 through 59). The QFX10000-60S-6Q line card contains 60 SFP+ ports that support 10-Gbps, two dual-speed QSFP28 ports that support either 40-Gbps or 100-Gbps, and four QSFP+ ports that support 40-Gbps. You can individually configure ports 0 to 59 for 10-Gbps or 1-Gbps port speed. Use the **set chassis fpc fps-slot-number pic pic-number port port-number speed 1G** command to change the mode of a port from 10-Gbps to 1-Gbps. The transceivers supported for 1-Gbps are QFX-SFP-1GE-LX, QFX-SFP-1GE-SX, and QFX-SFP-1GE-T.

[See [QFX10000 Line Cards](#) for details on the combination of modes supported on the ports.]

Services Applications

- **Support for Two-Way Active Measurement Protocol (TWAMP) and hardware timestamping of RPM probe messages (MX10000 and PTX10000 routers)**—Starting in Release 19.2R1-S4, Junos OS supports TWAMP and hardware timestamping of RPM probe messages on the MX10008, MX10016, PTX10008 and PTX10016 routers. You can use TWAMP to measure IP performance between two devices in a network. By enabling hardware timestamping of RPM you can account for the latency in the communication of probe messages and also generate more accurate timers in the Packet Forwarding Engine.

[See [Understanding Two-Way Active Measurement Protocol on Routers](#) and [Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX and MX Series Routers](#).]

New and Changed Features: 19.2R1-S1

MPLS

- **Color-based mapping of VPN services over SRTE (PTX Series)**—Starting in Junos OS Release 19.2R1-S1, 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](#).]

Routing Protocols

- **Decouple RSVP for IGP-TE (MX Series, PTX Series, ACX Series, QFX Series, SRX Series, and EX Series)**—Starting in Junos OS Release 19.2R1-S1, 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.

New and Changed Features: 19.2R1

Hardware

- **Advanced Cooling and Power Components (PTX10000 Routers)**—Starting in Junos OS Release 19.2R1, PTX10000 routers offer 5.5 KW power supplies, new high performance fan trays, and compatible fan tray controllers. The JNP10K-PWR-AC2 power supply supports AC, high-voltage alternating current (HVAC), DC, or high-voltage direct current (HVDC). The JNP10K-PWR-DC2 provides a 5.5 KW upgrade for DC users. The JNP10008-FAN2 and JNP10016-FAN2 offer increased air flow through the chassis. The JNP10008-FAN2 offers 1793 cubic feet per minute (CFM) per fan tray, while the JNP10016-FAN2 provides 3423 CFM per fan tray. Two new fan tray controllers support the new fan trays, the JNP10008-FTC2 and the JNP10016-FTC2.

[See [PTX10016 System Overview](#).]

Interfaces and Chassis

- **Domain Name System (DNS) is VRF aware (PTX Series)**—Starting in Junos OS Release 19.2R1, when the **management-instance** statement is configured at the **[edit system]** hierarchy level, you can use the non-default management routing instance **mgmt_junos** as the routing instance through which the DNS name server is reachable. To specify the routing instance **mgmt_junos**, configure our new configuration statement **routing-instance mgmt_junos**, at the **[edit system name-server server-ip]** hierarchy level.

[See [Management Interface in a Nondefault Instance, Configuring a DNS Name Server for Resolving a Hostname into Addresses](#), [name-server](#), and [show host](#).]

- **Support for health monitoring on the Routing Engine hard disk (PTX10008, PTX1000, PTX5000, and PTX10016)**—Starting in Junos OS Release 19.2R1, you can configure the device to perform certain health checks on the Routing Engine solid-state drive (SSD) and log a health event or raise an alarm in case a predefined health attribute threshold is breached. You can use the **set chassis routing-engine disk smart-check** command to instruct the system to raise an alarm when an SSD health attribute threshold is breached. You can view the alarm by using the command **show chassis alarms**.

[See [smart-check](#)]

Junos Telemetry Interface

- **Sensor- level statistics support on JTI (MX960, MX2008, MX2010, MX2020, PTX5000, PTX1000, and PTX10000 routers and QFX5100 and QFX5200 switches)**—Starting with Junos OS Release 19.2R1, you can issue the Junos operational mode command **show network-agent statistics** to provide more information on a per-sensor level for statistics being streamed to an outside collector by means of remote procedure calls (gRPC) and Junos telemetry interface (JTI). Only sensors exported with gRPC are supported. The command does not support UDP-based sensors.

[See [show network-agent statistics](#) and [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **ONCE mode supported using gNMI services and JTI (MX Series, QFX Series, PTX Series)**—Starting in Junos OS Release 19.2R1, you can include the ONCE mode with the **Subscribe** RPC when subscribing to gRPC Network Management Interface (gNMI) services to export statistics for telemetry monitoring and management using Junos telemetry interface (JTI). ONCE mode ensures that the collector is only streamed telemetry information one time.

The Subscribe RPC and subscription parameters are defined in the [gnmi.proto](#) file.

Streaming telemetry data through gNMI 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 JTI.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Packet Forwarding Engine statistics export using gNMI and JTI (PTX1000 and PTX10000 routers)**—Starting in Junos OS Release 19.2R1, you can stream Packet Forwarding Engine statistics to an outside collector using gRPC Management Interface (gNMI) version 0.7.0 and junos telemetry interface (JTI). Before this release, these statistics were exported using OpenConfig gRPC and UDP protocol buffer

(gpb) format. OpenConfig gRPC and gNMI are both protocols used to modify and retrieve configurations as well as export telemetry streams from a device in order to manage and monitor it

To provision Packet Forwarding Engine sensors to export data through gNMI, use the Subscribe RPC defined in the [gnmi.proto](#) to specify request parameters.

Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **Specify Routing Instance for JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R1, you can specify the routing instance to use for remote procedure call (gRPC) services. Include the **routing-instance** *instance-name* at the **[edit system services extension-service request-response grpc]** hierarchy level. The routing instance name specified should match the name of the existing routing instance, such as a name configured under the **[routing-instances]** hierarchy level or **mgmt_junos** if **system management-instance** is configured (the dedicated management routing instance).

Configuring the routing instance lets you choose the VRF for gRPC services. When the routing instance is not configured, the default behavior is that all gRPC-related services are available through the management **fxp0/em0** interface.

MPLS

- **Dynamic creation of segment routing LSPs using BGP protocol next hops (PTX Series)**—Starting in Junos OS Release 19.2R1, you can configure tunnel templates on colored and non-colored segment routing traffic-engineered (SR-TE) paths. These templates enable dynamic creation of segment routing tunnels using protocol next hops with BGP prefixes to resolve destination segment identifiers (SIDs).

With this feature, you can benefit from reduced configuration, especially when the network deployment requires connectivity from each provider edge (PE) device to every other PE device.

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

- **Support for IS-IS segment routing (PTX Series)**—Starting in Junos OS Release 19.2R1, you can use IS-IS segment routing through MPLS. Currently, label advertisements are supported for IS-IS only. IS-IS creates an adjacency segment per adjacency, per level, and per address family (one each for IPv4 and IPv6). Junos OS IS-IS implementation allocates node segment label blocks in accordance with the IS-IS protocol extensions for supporting segment routing node segments. It provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:
 - **source-packet-routing**—Enable the source packet routing feature.
 - **node-segment**—Enable source packet routing at all levels.

- **use-source-packet-routing**—Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and for primary IS-IS source packet routing node segments.
- **no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.

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

Network Management and Monitoring

- **Support for displaying valid user input in the CLI for command options and configuration statements in custom YANG data models (PTX Series)**—Starting in Junos OS Release 19.2R1, the CLI displays the set of possible values for a given command option or configuration statement in a custom YANG data model when you include the **action-expand** extension statement in the option or statement definition and reference a script that handles the logic. The **action-expand** statement must include the **script** child statement, which defines the Python action script that is invoked when a user requests context-sensitive help in the CLI for the value of that option or statement.

[See [Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules](#).]

- **Enhanced sFlow (PTX1000)**—Starting in Junos OS Release 19.2R1, MPLS and GRE traffic flows are added to sFlow functionality for the PTX1000 routers.

[See [Overview of sFlow Technology](#).]

Port Security

- **Fallback PSK for Media Access Control Security (MACsec) (PTX Series)**—Starting in Junos OS Release 19.2R1, fallback PSK for MACsec is supported on PTX Series routers that support MACsec. The fallback PSK provides functionality to establish a secure session in the event that the primary PSKs on each end of a MACsec-secured link do not match.

[See [Configuring Media Access Control Security \(MACsec\) on MX Series Routers](#).]

Routing Policy and Firewall Filters

- **Support for interface, forwarding-class, and loss priority match conditions on egress interfaces (PTX10008, PTX10016)**—Starting with Junos OS Release 19.2R1, you can apply the **interface**, **forwarding-class**, and **loss-priority** firewall filter match conditions in the egress direction on IPv4 and IPv6 interfaces. You configure the match conditions at the **[edit firewall]** hierarchy level. This feature was previously supported in an "X" release of Junos OS.

[See [Firewall Filter Match Conditions and Actions \(QFX10000\)](#).]

Routing Protocols

- **Support for export of BGP Local RIB through BGP Monitoring Protocol (BMP) (PTX Series)**—Starting in Junos OS Release 19.2R1, BMP is enhanced to support monitoring of local RIB (**loc-rib**) policy. The **loc-rib** policy is added to RIB types under the **bmp route-monitoring** statement.

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

Services Applications

- **Support for IPv6 BGP next-hop address in IPv6 and MPLS-IPv6 inline flow record templates (MX Series and PTX Series)**—Starting in Junos OS Release 19.2R1, a new element, IPv6 BGP NextHop Address, is available in the IPv6 inline flow record template and the MPLS-IPv6 inline flow record template to add support for the IPv6 BGP NextHop information element. The new element is supported on both version 9 and version 10 (IPFIX) export formats. The element ID is 63 and the element size is 16 bytes.

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

- **Two-Way Active Measurement Protocol (TWAMP) Support (PTX Series)**—Starting in Junos OS Release 19.2R1, PTX Series routers support the Two-Way Active Measurement Protocol (TWAMP). TWAMP defines a standard for measuring IPv4 performance between two devices in a network. You can use the TWAMP-Control protocol to set up performance measurement sessions between a TWAMP client and a TWAMP server, and use the TWAMP-Test protocol to send and receive performance measurement probes.

The destination interface **si-x/y/z** attribute, which is meant for enabling inline services is not supported on PTX Series routers for TWAMP client configurations.

[See [Understanding Two-Way Active Measurement Protocol on Routers](#).]

Software Defined Networking (SDN)

- **PCE-initiated bypass LSPs (PTX Series)**—Starting in Junos OS Release 19.2R1, the Path Computation Element Protocol (PCEP) functionality is extended to allow a stateful Path Computation Element (PCE) to initiate, provision, and manage bypass label-switched paths (LSPs) for a protected interface. Multiple bypass LSPs with bandwidth reservation can be initiated by the PCE to protect a resource.

With this feature, you can benefit from the LSP state synchronization of manual, dynamic, and PCE-initiated bypass LSPs from a PCE, and leverage on the PCE's global view of the network, resulting in better control over traffic at the time of a failure, and deterministic path computation of protection paths.

[See [Support of the Path Computation Element Protocol for RSVP-TE Overview](#).]

Software Installation and Upgrade

- **The curl binary is packaged and made available on all Junos OS variants (PTX Series)**—The curl binary is a command-line utility, used from the shell, that you can use to perform operations over several transport protocols, including the following: dict, file, ftp, gopher, http, imap, pop3, rtsp, smtp, telnet, tftp. The features enabled on Junos OS are curl version 7.59, libcurl version 7.59.

System Management

- **Support for transferring accounting statistics files and router configuration archives using HTTP URL (PTX Series)**—Starting in Junos OS Release 19.2R1, you can transfer accounting statistics files and router configuration archives to remote servers by using an HTTP URL. In addition to SCP and FTP, the following HTTP URL will be supported under the **archive-sites** statement:

`http://username@host:url-path password password`

- To transfer accounting statistics files, configure **archive-sites** under **[edit accounting-options file <filename>]** hierarchy.
- To transfer router configuration archival, configure **archive-sites** under **edit system archival configuration** hierarchy.
- To view the statistics of transfer attempted, succeeded, and failed, use the **show accounting server statistics archival-transfer** command.
- To clear the statistics of transfer attempted, succeeded, and failed, use the **clear accounting server statistics archival-transfer** command.

[See [archive-sites](#), [Backing Up Configurations to an Archive Site](#), [show accounting server statistics archival-transfer](#), and [clear accounting server statistics archival-transfer](#)].

- **Precision Time Protocol (PTP) transparent clock with IPv6 transport (PTX10001-20C and ACX6360-OR devices)**—Starting with Junos OS Release 19.2R1, PTP using IPv6 transport synchronizes clocks throughout a packet-switched network. With a transparent clock, the PTP packets are updated with residence time as the packets pass through the routers. There is no master clock-slave clock designation. With an end-to-end transparent clock, only the residence time is included. The residence time can be sent in a one-step process, which means that the timestamps are sent in one packet.

You can configure the transparent clock at the **[edit protocols ptp]** hierarchy level.
[See [Understanding Transparent Clocks in Precision Time Protocol.](#)]

SEE ALSO

What's Changed 221
Known Limitations 227
Open Issues 228
Resolved Issues 230
Documentation Updates 238
Migration, Upgrade, and Downgrade Instructions 238

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.2R3 | 222](#)
- [What's Changed in Release 19.2R2 | 222](#)
- [What's Changed in Release 19.2R1-S5 | 224](#)
- [What's Changed in Release 19.2R1 | 224](#)

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

What's Changed in Release 19.2R3

General Routing

- **Trigger alarms when a PTX10008 or PTX10016 router has a mix of AC and DC power supplies**—If you install a mix of AC and DC power supply units (PSUs), Junos OS raises an alarm to indicate that there is a mix of AC and DC power supplies in the router. To fix this alarm, you need to ensure that you install same type of power supplies.

[See [Understanding Chassis Alarms](#).]

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

See [traceoptions \(Services\)](#).

What's Changed in Release 19.2R2

General Routing

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

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

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

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

- **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.

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, 19.2R2, and later, PTX Series routers support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

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

- **Updates to the `show interfaces` and `show policer` commands (PTX Series)**—Starting in Junos OS Release 19.2R2, 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_`) even though it is not supported on PTX Series routers.

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 19.2R2, we have included an XML RPC equivalent for the `show bgp output-scheduler | display xml rpc` CLI command. In Junos OS releases before Release 19.2R2, 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)**—You can now configure a maximum of 2,147,483,647 seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

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

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

What's Changed in Release 19.2R1-S5

MPLS

- **IPv4 explicit-null label retained from the merged protocol MPLS label stack**—The IPv4 explicit-null label is retained from the merged protocol MPLS label stack, if the IPv4 explicit-null is at the bottom of the MPLS label stack.

What's Changed in Release 19.2R1

Interfaces and Chassis

- **Support to get optics loopback status for QSFP-100GE-DWDM2 transceivers (PTX Series)**—In Junos OS Release 19.2R1, and later, on PTX Series routers, you can obtain the optics loopback status of QSFP-100GE-DWDM2 transceivers along with the regular Ethernet loopback status by issuing the `show interfaces interface-name` or `show interfaces interface-name brief` command. A new output field, **Optics Loopback** is added under **Link-level type** when the `show interfaces interface-name` CLI command is executed.
- **Monitoring information available only in trace log (PTX Series)**—In Junos OS Release 19.2R1 and later, the Ethernet link fault management daemon (lfmd) in the peer router stops monitoring the locally occurred errors until unified ISSU is completed. You can view the monitoring-related details only through the trace log file.

- **Health check for power supplies (PTX10008 and PTX10016)**—Starting in Junos OS Release 19.2R1, on the PTX10008 and PTX10016 routers, the **show chassis environment pem** command displays the health check information about the DC or AC power supplies. For any power supply that does not support health check, the status is shown as **Unsupported**. The system starts health check of a power supply only if the power consumption exceeds 7 kW.

[See [show chassis environment pem](#).]

- **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 19.2R1 and later, 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.

MPLS

- **New debug statistics counter (PTX Series)**—The **show system statistics mpls** command has a new output field, called **Packets dropped, over p2mp composite nexthop**, to record the packet drops over composite point-to-multipoint next hops.

Network Management and Monitoring

- **Change in error severity (PTX10016)**—Starting in Junos OS Release 19.2R1, on PTX10016 routers, the severity of the FPC error, shown in the syslog as **PE Chip::FATAL ERROR!! from PE2[2]: RT: Clear Fatal if it is detected LLMEM Error MEM:llmem, MEMTYPE: 1**, is changed from fatal to nonfatal (or minor). In case of this error, only a message is displayed for information purposes. To view the error details, you can use the show commands **show chassis fpc errors** and **show chassis errors active**.

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

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (PTX Series)**—Starting in Junos OS Release 19.2R1, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.
- **Custom YANG RPC support for input parameters of type empty (PTX Series)**—Starting in Junos OS Release 19.2R1, custom YANG RPCs support input parameters of type **empty** when executing the RPC's command in the Junos OS CLI, and the value passed to the action script is the parameter name. In earlier releases, input parameters of type **empty** are supported only when executing the RPC in a NETCONF or Junos XML protocol session, and the value passed to the action script is the string '**none**'.

[See [Creating Action Scripts for YANG RPCs on Devices Running Junos OS](#).]

Routing Protocols

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

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

Services Applications

- **Support for enabling hardware timestamping of RPM probe messages (PTX Series)**—In Junos OS Releases 19.2R1, PTX Series routers support timestamping of RPM probe messages on the Packet Forwarding Engine. The following configuration statements at the `[edit services rpm probe owner test test-name]` hierarchy level are supported:

- **hardware-timestamp**—To enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor.
- **one-way-hardware-timestamp**—To enable timestamping of RPM probe messages for one-way delay and jitter measurements.

You can use these timestamping features only with `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp` probe types.

- **Hide HA information when the service set does not have HA configured**—When you run the `show services ha detail` command on a configuration with a service set, which does not have HA configured, the HA information is hidden for the service set in the output.

SEE ALSO

[What's New | 213](#)

[Known Limitations | 227](#)

[Open Issues | 228](#)

[Resolved Issues | 230](#)

[Documentation Updates | 238](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Known Limitations

IN THIS SECTION

- [General Routing | 227](#)
- [Interfaces and Chassis | 228](#)

Learn about known limitations in this release for the PTX Series.

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

General Routing

- On PTX Series with FPC3, if the filter action of policer and count are configured in the same filter, then it might cause the policer counter to be higher than the actual count. [PR1089330](#)
- In the specific case of semigraceful RCB reboot initiated by the internal shell command **vhclient init 0**, GRES takes longer to complete that is, 3 minutes as opposed to 21 seconds. The regular CLI command: **request vmhost reboot** (graceful) and a jack-out-jack-in of the Routing Engine (ungraceful) do not exhibit this delay. [PR1312065](#)
- The Routing Engine boots from the secondary disk when you:
 - Press the reset button on the RCB front panel, while the Routing Engine is booting up but before Junos OS reboots.
 - Upgrade software, by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network.
 - Upgrade BIOS and the upgrade fails.
 - Reboot and the system hangs before Junos OS reboots. [PR1344342](#)
- The command **request vmhost power-off** does not actually power off the system in the latest releases. It only does a reboot and the system comes back up. [PR1393061](#)
- Because of the small counter size present in the ASIC, the normal discard counter reported in the CLI is less than the actual packet drop rate. [PR1394979](#)

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after the upgrade. This is because of an old version of `/var/db/cfm.db`. [PR1281073](#)

SEE ALSO

[What's New | 213](#)

[What's Changed | 221](#)

[Open Issues | 228](#)

[Resolved Issues | 230](#)

[Documentation Updates | 238](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Open Issues

IN THIS SECTION

- [General Routing | 229](#)
- [Interfaces and Chassis | 230](#)
- [Layer 2 Ethernet Services | 230](#)
- [Routing Protocols | 230](#)

Learn about open issues in this release for the PTX Series.

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

General Routing

- On a PTX Series PIC with the CFP2-DCO-T-WDM transceiver installed, after repeated configuration rollback, the link sometimes takes a long time to come up. [PR1301462](#)
- When CFP2-DCO-T-WDM-1 is plugged in a PTX Series PIC, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)
- Output of J-Flow sensor is changed for a FPC that is not configured. [PR1379770](#)
- The rx_power value streamed to the telemetry server is the raw value (mW) returned directly from the transceiver driver. The Junos OS CLI value has been transformed in the transportd daemon into different units: (Rx input total power(0.01dBm). [PR1411023](#)
- Traffic loss for more than 15 seconds is seen when 50 percent of aggregated Ethernet links are brought down by restarting multiple FPCs. [PR1412578](#)
- Firewall counter for lo0 does not increment for gladiator. [PR1420560](#)
- The em2 interface configuration causes the FPC to crash during initialization and the FPC does not come online. Only after deleting the em2 configuration and restarting the router does the FPC come online. [PR1429212](#)
- More packet loss seen after unified ISSU with InterAS L3VPN OptionB configuration. [PR1435578](#)
- Optics-options syslog and link-down do not work on PTX5000 for FPC3. For low warning breach event, when configured action is syslog only, link goes down. [PR1461404](#)
- When plugging the optic QSFP-100G-LR4-T2(740-061409) to PTX3000 or PTX5000, the device displays an alarm **unsupported optic** and the interface will not be up. [PR1511492](#)

Interfaces and Chassis

- The cfmd process might continuously crash after upgrade. [PR1281073](#)

Layer 2 Ethernet Services

- The aggregated Ethernet members take more than 5 seconds from the physical interface up to LACP CD state when the laser is switched off and then turned on 25 percent of the aggregated Ethernet 10G members. [PR1415615](#)

Routing Protocols

- With Bidirectional Forwarding Detection (BFD) configured on an aggregated Ethernet interface, if you disable/enable the aggregated Ethernet interface, then that interface and the BFD session might not come up. [PR1354409](#)

SEE ALSO

What's New 213
What's Changed 221
Known Limitations 227
Resolved Issues 230
Documentation Updates 238
Migration, Upgrade, and Downgrade Instructions 238

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3 | 231](#)
- [Resolved Issues: 19.2R2 | 231](#)
- [Resolved Issues: 19.2R1 | 235](#)

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

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

Resolved Issues: 19.2R3

General Routing

- PTX1000 and PTX10002 routers might drop or get discarded after transient SIB/FPC voltage alarms. [PR1460406](#)
- **Optics-options** syslog and **link-down** do not work as expected on PTX5000 for FPC3 line card. [PR1461404](#)
- Sampling process might crash when MPLS/MPLS over UDP traffic is sampled. [PR1477445](#)
- Multicast routes add/delete events might cause adjacency and LSPs to go down. [PR1479789](#)
- BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)
- Outbound SSH connection flap or memory leak issue might be observed during pushing configuration to ephemeral DB with high rate. [PR1497575](#)
- Traffic might get dropped or discarded in the fast-reroute scenario. [PR1501817](#)
- Observed the error message **PFE_ERROR_FAIL_OPERATION: IFD et-1/0/8: RS credits failed to return: init=192 curr=193 chip=5**. [PR1502716](#)
- On a dual Routing Engine GRES/NSR enabled PTX10008/PTX10016 router, a few TCP-based application sessions like BGP/LDP might flap upon Routing Engine mastership switch. [PR1503169](#)

MPLS

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

Routing Protocols

- On PTX3000 and PTX5000, the ppmd core file is generated after configuring sbfd responder on RE-DUO-2600. [PR1477525](#)
- The BGP route-target family might prevent RR from reflecting L2VPN and L3VPN routes. [PR1492743](#)

Resolved Issues: 19.2R2

General Routing

- The agentd sensor transmits multiple interface telemetry statistics per FPC slot. [PR1392880](#)
- Confirmation message is missing when issuing **request vmhost reboot re***. [PR1397912](#)
- Kernel memory leak in virtual-memory because of interface flap. [PR1407000](#)

- Slow SNMP response time on entityMIB might be seen in the fully loaded setup with many SFPs. [PR1411062](#)
- PTX Series interface stays down after maintenance. [PR1412126](#)
- A core file **core-olympus-fpc3-sevfpc.elf.0.tgz** is generated after switchover with GRES in PTX3000. [PR1415145](#)
- Support for 140e NVM firmware upgrade through CLI for PTX10002-60c. [PR1418909](#)
- Packet Forwarding Engine wedge might occur after issuing the **show forwarding-options load-balance** command. [PR1422464](#)
- 4x10G interfaces on PTX3000 and PTX5000 third-generation FPCs might not come up after frequent flapping for a long time. [PR1422535](#)
- LACP packet does not pass through Layer 2 circuit. [PR1424553](#)
- The host-bound traffic might be dropped after performing change configuration related to prefix-list. [PR1426539](#)
- The jumbo frame size packets are dropped when maximum MTU is configured. [PR1428094](#)
- Interface does not come up after interface flapping and FPC reboot. [PR1428307](#)
- Inline J-Flow might cause PECHIP Major error [PR1429419](#)
- Reclassification policy applied on the route prefixes might not work on PTX Series platforms. [PR1430028](#)
- The l2cpd process might crash and generate a core file when interfaces are flapping. [PR1431355](#)
- IPFIX Flow timestamp is not matching with NTP synchronized system time. [PR1431498](#)
- **SIB Link Error** detected on a specific Packet Forwarding Engine might cause complete service impact. [PR1431592](#)
- The FPC might crash when a firewall filter is modified. [PR1432116](#)
- Traffic loss might be seen on PTX10000 platforms using line card. [PR1433300](#)
- Traffic drop might occur on PTX Series during filter change operation. [PR1433648](#)
- IPv6 neighbor solicitation packets getting dropped on PTX Series. [PR1434567](#)
- Router Engine switchover does not work as expected while SSD failure occurs. [PR1437745](#)
- No chassis alarm is raised when PEM is removed or power lost to PEM. [PR1439198](#)
- LACP state might get stuck in 'Attached' state after disabling peer active members. [PR1439268](#)
- FPC reboot might be observed in the events of jlock hog more than 5 seconds. [PR1439929](#)
- CPU might hang or interface might be stuck down on particular 100G port on PTX Series devices. [PR1440526](#)
- Interfaces on PTX Series devices might not come up after FPC restart or port flap. [PR1442159](#)
- The PTX Series devices BCM firmware needs to be upgraded to DE2E. [PR1445473](#)

- Receipt of a malformed packet for J-Flow sampling might create a FPC core file. [PR1445585](#)
- The use of wildcard <*> is not an available option at the group level of the Junos OS CLI. [PR1445651](#)
- The process jdncpd may crash after issuing the command "show access-security router-advertisement-guard" [PR1446034](#)
- The process vhostd might crash without generating a core file and vhostd process might fail restart automatically. [PR1448413](#)
- Egress sampling for sflow might stop working for more than 8 interfaces on PTX Series platforms. [PR1448778](#)
- Changing the hostname triggers the LSP on-change notification, not an adjacency on-change notification. Also, currently IS-IS is sending host-name instead of system-id in OC paths. [PR1449837](#)
- Interfaces might flap forever after deleting the interface disable configuration. [PR1450263](#)
- FPC generates **Voltage Tolerance Exceeded** major alarm. [PR1451011](#)
- Firewall filter applied at interface level might not work when MPLS label is present in certain scenarios. [PR1452716](#)
- The 100G interface might not come up after flapping on PTX5000. [PR1453217](#)
- The FPC might crash when the severity of error is modified. [PR1453871](#)
- Traffic might silently get dropped upon interface flap after DRD auto-recovery. [PR1459698](#)
- The "forwarding" option is missed in routing-instance type. [PR1460181](#)
- PTX1000/PTX10002 might silently get dropped or discarded after transient SIB/FPC voltage alarms. [PR1460406](#)
- Hardware failure in CB2-PTX causes traffic interruption. [PR1460992](#)
- The sample/syslog/log action in output firewall filter with packet of size less than 128 might cause ASIC wedge (all packet loss) on PTX Series platforms. [PR1462634](#)
- PIC might restart if the temperature of QSFP optics is overheated on PTX3000 and PTX5000. [PR1462987](#)
- FPC might restart during run time on PTX10000platforms. [PR1464119](#)
- EBUF parity interrupt is not seen on PTX Series platforms. [PR1466532](#)
- IPv6 traffic might get dropped in L3VPN network. [PR1466659](#)
- Packet Forwarding Engine error logs **prds_packet_classify_notification: Failed to find fwd nh for flabel 48** might be reported when IGMP packets got sampled on PTX5000 Platform. [PR1466995](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- Incorrect counter value for "Arrival rate" and "Peak rate" for DDoS commands. [PR1470385](#)
- Traffic loops for pure Layer 2 packets coming over EVPN tunnel with destination MAC matching IRB MAC. [PR1470990](#)

- A PTX5000 SIB3 might fail to come up in slot 0 and/or slot 8 when RE1 is master. [PR1471178](#)
- MACsec traffic over Layer 2 circuit might not work on PTX10000 and PTX1000 platforms after upgrading from Junos OS Release 15.1 to later. [PR1475089](#)
- Sampling process might crash when MPLS/MPLS over UDP traffic is sampled. [PR1477445](#)
- Multicast routes add or delete events might cause adjacency and LSPs to go down. [PR1479789](#)
- FPC might crash when dealing with invalid next-hops. [PR1484255](#)
- BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)

Forwarding and Sampling

- The pfd might crash and not be able to come up on the PTX Series or TVP platforms. [PR1452363](#)

Infrastructure

- FPC might reboot if jlock hog occurs on all Junos VM based platforms. [PR1439906](#)
- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. [PR1462986](#)

Interfaces and Chassis

- After member interface flapping aggregated Ethernet remains down on 5X100GE DWDM CFP2-ACO PIC. [PR1429279](#)

Layer 2 Ethernet Services

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

MPLS

- Traffic loss might be seen after LDP session flap rapidly. [PR1436119](#)
- The transit packets might be dropped if an LSP is added or changed on PTX Series device. [PR1447170](#)
- Kernel crash and device might restart. [PR1478806](#)

Platform and Infrastructure

- Packet drops, replication failure, or ksyncd crashes might be seen on the logical system of a Junos OS device after Routing Engine switchover. [PR1427842](#)
- REST API process become nonresponsive when a number of requests come at a high rate. [PR1449987](#)

Routing Protocols

- Routing Engine based micro BFD packets do not go out with configured source IP when the interface is in logical system. [PR1370463](#)
- PTX Series devices cannot intercept PIM BSR message. [PR1419124](#)
- The rpd might crash with SRTE configuration change. [PR1442952](#)

- L3VPN PE-CE link protection exhibits unexpected behavior on PTX10000 platforms. [PR1447601](#)
- SSH login might fail if a user account exists in both local database and RADIUS/TACACS+. [PR1454177](#)
- The "other querier present interval" timer cannot be changed in IGMP/MLD snooping scenario. [PR1461590](#)
- The rpd process might crash with BGP multipath and route withdraw occasionally. [PR1481589](#)
- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

VPNs

- In a specific CE device environment in which asynchronous-notification is used, after the link between the PE and CE devices goes up, the L2 circuit flap repeatedly. [PR1282875](#)
- Memory leak might be seen if PIM messages are received over an MDT (mt- interface) in the MVPN scenario. [PR1442054](#)

Resolved Issues: 19.2R1

General Routing

- Repeated log messages `%PFE-3 fpcX expr_nh_index_tree_ifl_get` and `expr_nh_index_tree_ipaddr_get` are observed when a sampling packet is discarded with the log (or syslog) statement under the firewall filter. [PR1304022](#)
- Disable reporting of the correctable single-bit error on Hybrid Memory Cube (HMC) and prevent a major alarm. [PR1384435](#)
- CPU overuse might be observed on PTX/QFX10000 Series platform. [PR1399369](#)
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. [PR1399683](#)
- Only one Packet Forwarding Engine could be disabled on an FPC with multiple Packet Forwarding Engines in an error/wedge condition. [PR1400716](#)
- The TCP connection between ppmmd and ppman might be dropped because of a kernel issue. [PR1401507](#)
- Log message **JAM HW data base open failed for ptx5kpic_3x400ge-cfp8** during commit. [PR1403071](#)
- The **Incorrect mem stats** message is seen in the logs of PTX Series Type 1 FPC. [PR1404088](#)
- On a PTX3000, FPCs are not able to come online for tens of minutes after a reboot of the chassis. [PR1404611](#)
- ZTP upgrade failed. Image file transfer got stalled during the ZTP process. [PR1404832](#)
- On PTX3000 and PTX5000, the backup CB's chassis environment status is displayed as Testing after the backup CB is removed and reinserted. [PR1405181](#)

- The 100-gigabit SR4 optics with part number 740-061405 should be displayed as **QSFP-100G-SR4-T2**. [PR1405399](#)
- Layer 2 VPN might flap repeatedly after the link up between the PE device and CE device comes up. [PR1407345](#)
- Ports 4, 5, 14, 15, 24, and 25 on PTX10K-LC1101/PTX10K-LC1105 might fail to come up after the device is upgraded to Junos OS Release 17.4R2-S3. [PR1407655](#)
- **openconfig-network-instance:network-instances** support for IS-IS must be hidden unless supported. [PR1408151](#)
- [PTX3000-CI] : Observing rpd crash @
`if_addr_link,krt_chnh_template_create_restart,krt_chnh_create_restart,krt_comp_add_comp_nh,
krt_build_comp_nh,krt_build_nexthop,krt_rt_add_sock,krt_decode_rt,krt_sysctl_read_consume,
krt_rt_read,krt_sys_rtread,krt_var_init,ctx_handle_node,ctx_walk_features,task_read_config,main`
[PR1409051](#)
- PTX Series Inline J-Flow : FPC goes offline when the sampling rate is changed at runtime to 80000;also dcpfe core file was generated. [PR1409502](#)
- The port on the FPC (for example, JNP10K-LC1101) might fail to come up. [PR1409585](#)
- Hostname does not update at FPC shell after system configuration change on CLI. [PR1412318](#)
- Junos OS PCC might reject PCUpdate/PCCreate message if the metric type is other than type 2. [PR1412659](#)
- The L2 circuit egress PE might drop the traffic in a FAT+CW enabled L2 circuit scenario when another FAT+CW enabled L2 circuit PW flaps. [PR1415614](#)
- Traffic loss could be seen for the duration of the hold-time down timer when flapping an interface with hold-time down timer configured. [PR1418425](#)
- Error messages might be seen on PTX10000/QFX10000 platforms during DFE tuning. [PR1421075](#)
- Virtual Chassis might become unstable and FXPC core files might be generated when there are a lot of configured filter entries. [PR1422132](#)
- While committing a huge configuration, the customer sees the error **error: mustd trace init failed**.[PR1423229](#)
- Traffic is dropped after an FPC reboot with aggregated Ethernet member links deactivated by the remote device. [PR1423707](#)
- A specific interface on the P3-15-U-QSFP28 PIC remains down until another interface comes up. [PR1427733](#)
- The output of **show chassis environment** shows Input0 and Input1. [PR1428690](#)

Infrastructure

- The **request system recover oam-volume** command might fail on PTX Series routers. [PR1425003](#)

Interfaces and Chassis

- Syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** upon LFM related configuration commit on ae interfaces. [PR1423586](#)
- Some ports on PTX Series routers might remain down after the FPC or device at the remote side is rebooted. [PR1429315](#)

MPLS

- An RSVP-signaled LSP might stay in the down state after a link in the path flaps. [PR1384929](#)
- LDP route flap during unrelated commit. [PR1416032](#)
- Bypass dynamic RSVP LSP tears down too soon when being used for protecting the LDP LSP with the **dynamic-rsvp-lsp** statement. [PR1425824](#)

Platform and Infrastructure

- RPM hardware-timestamp and **one-way-hardware-timestamp** statements are not enabled. [PR1399842](#)

Routing Protocols

- A syslog message is seen whenever the prefix SID coincides with the node sid. [PR1403729](#)
- An rpd memory leak might be seen in an IS-IS segment routing scenario. [PR1404134](#)
- Dynamic routing protocol flapping with vmhost Routing Engine switchover on NG-RE. [PR1415077](#)
- Route churn occurs when a nonrelevant configuration is changed. [PR1423647](#)
- RPD might crash with **ospf overload as-external** configuration. [PR1429765](#)

SEE ALSO

[What's New | 213](#)

[What's Changed | 221](#)

[Known Limitations | 227](#)

[Open Issues | 228](#)

[Documentation Updates | 238](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Documentation Updates

IN THIS SECTION

- [Installation and Upgrade Guide | 238](#)

This section lists the errata and changes in Junos OS Release 19.2R3 documentation for the PTX Series.

Installation and Upgrade Guide

- **Veriexec explained (PTX Series)**—Verified Exec (also known as veriexec) is a file-signing and verification scheme that protects the Junos operating system (OS) against unauthorized software and activity that might compromise the integrity of your device. Originally developed for the NetBSD OS, veriexec was adapted for Junos OS and enabled by default from Junos OS Release 7.5 onward.

[See [Veriexec Overview](#).]

SEE ALSO

[What's New | 213](#)

[What's Changed | 221](#)

[Known Limitations | 227](#)

[Open Issues | 228](#)

[Resolved Issues | 230](#)

[Migration, Upgrade, and Downgrade Instructions | 238](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 19.2 | 239](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 241](#)
- [Upgrading a Router with Redundant Routing Engines | 242](#)

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

Basic Procedure for Upgrading to Release 19.2

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 19.2R3:

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-x86-64-19.2R3.9.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-x86-64-19.2R3.9-limited.tgz
```

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

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

- `scp://hostname/pathname`

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).

NOTE: After you install a Junos OS Release 19.2 **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 | 213](#)

[What's Changed | 221](#)

[Known Limitations | 227](#)

[Open Issues | 228](#)

[Resolved Issues | 230](#)

[Documentation Updates | 238](#)

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- What's New | 243
- What's Changed | 253
- Known Limitations | 259
- Open Issues | 261
- Resolved Issues | 265
- Documentation Updates | 287
- Migration, Upgrade, and Downgrade Instructions | 288

These release notes accompany Junos OS Release 19.2R3 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- What's New in Release 19.2R3 | 244
- What's New in Release 19.2R2 | 244
- What's New in Release 19.2R1-S1 | 244
- What's New in Release 19.2R1 | 245

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

NOTE: The following QFX Series platforms are supported in Release 19.2R3: 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.

What's New in Release 19.2R3

There are no new features or enhancements to existing features for QFX Series in Junos OS Release 19.2R3.

What's New in Release 19.2R2

Routing Protocols

- **Option to pause BGP multipath computation during BGP peering churn (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R2, you can choose to defer multipath computation for all families during a BGP peering churn. In very large-scale network deployments during BGP peering churn there is a temporary spike in multipath computation, which takes a toll on the Packet Forwarding Engine resources. This feature allows you to pause the multipath computation and to resume after the peering churn settles down. Note that if there is no BGP peering churn, then multipath computation is not paused.

To enable the pause option for BGP multipath computation during BGP peering churn, include the **pause computation** statement at the **[edit protocols BGP multipath]** hierarchy level.

What's New in Release 19.2R1-S1

EVPN

- **Overlay load balancing in an EVPN-VXLAN network (QFX5200 and QFX5210 switches)**—Starting in Junos OS Release 19.2R1-S1, QFX5200 and QFX5210 switches that function as leaf or spine devices in an EVPN-VXLAN network (centrally-routed and edge-routed bridging overlays) support load balancing among different virtual tunnel endpoints (VTEPs). We support overlay load balancing 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). Overlay 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.

To enable overlay load balancing, enter the **vlan-overlay-load-balance** configuration statement at the **[edit forwarding-options]** hierarchy level.

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

Routing Protocols

- **Decouple RSVP for IGP-TE (MX Series, PTX Series, ACX Series, QFX Series, SRX Series, and EX Series)**—Starting in Junos OS Release 19.2R1-S1, 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.

What's New in Release 19.2R1

Hardware

- **5.5 KW Power Supplies (QFX10000 switches)**—Starting in Junos OS Release 19.2R1, QFX10000 modular chassis adds two 5.5 KW power supplies. The JNP10K-PWR-AC2 power supply supports AC, high-voltage alternating current (HVAC), DC, or high-voltage direct current (HVDC). The JNP10K-PWR-DC2 provides a 5.5 KW upgrade for DC users. Two new ordering SKUs are available for the QFX10008 switch: QFX10008-BASE-H and QFX10008-REDUND-H.

The JNP10K-PWR-AC2 takes AC input and provides DC output of 12.3 VDC, 5000 W with a single feed and 5500 W with a dual feed. For AC systems, the operating input voltage is 180 to 305 VAC and for DC systems, the operating input voltage is 190 to 410 VDC.

The JNP10K-PWR-DC2 power supply provides two power supplies in a single housing that accepts either 60 A or 80 A using four redundant input power feeds. PS_0 and PS_1 each have redundant input feeds: A0 and/or B0 for PS_0 and A1 and/or B1 for PS_1. The input is configured using a set of dip switches on the power supply faceplate. The output is dependent on the settings of these dip switches.

[See [QFX10008 System Overview](#).]

EVPN

- **EVPN-VXLAN support (QFX10002-60C switches)**—Starting in Junos OS Release 19.2R1, the QFX10002-60C switch can function as a Layer 2 or Layer 3 VXLAN gateway in both EVPN-VXLAN centrally-routed and edge-routed bridging overlays (EVPN-VXLAN topologies with two-layer and collapsed IP fabrics). In these roles, the switch supports the following features:
 - Enterprise style of Layer 2 interface configuration
 - Active/active multihoming
 - Default routing instance

- Multiple routing instances of type virtual switch, and VLAN-aware service on the virtual switch routing instance
- Pure type-5 routes
- Proxy ARP use and ARP suppression, and proxy NDP use and NDP suppression on an IRB interface
- ESIs on physical and aggregated Ethernet interfaces
- OSPF, IS-IS, BGP, and static routing on IRB interfaces
- DHCP relay
- IPv6 support for user data traffic
- EVPN-VXLAN with MPLS as transport layer
- MAC mobility

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

- **Unicast VXLAN with MC-LAG (QFX5120 switches)**—Instead of EVPN providing remote VXLAN tunnel endpoint (remote VTEP) reachability information, starting in Junos OS Release 19.2R1, Junos OS supports the static configuration of remote VTEPs on QFX5120 switches in a network that also includes the following elements:
 - Endpoints multihomed to a pair of QFX5120 switches, each of which functions as Layer 2 VXLAN gateways or leaf devices, and as MC-LAG peers
 - Spine devices functioning as Layer 3 devices that handle the QFX5120 switches' IPv4 traffic

In this environment, the QFX5120 switches also support the configuration of ingress node replication, which enables the replication of Layer 2 BUM traffic. In fact, when you configure ingress node replication, other multicast features are disabled.

Interfaces and Chassis

- **Domain Name System (DNS) is VRF aware (QFX Series)**—Starting in Junos OS Release 19.2R1, when the **management-instance** statement is configured at the **[edit system]** hierarchy level, you can use the non-default management routing instance `mgmt_junos` as the routing instance through which the DNS name server is reachable. To specify the routing instance `mgmt_junos`, configure our new configuration statement **routing-instance mgmt_junos**, at the **[edit system name-server server-ip]** hierarchy level.

[See [Management Interface in a Nondefault Instance](#), [Configuring a DNS Name Server for Resolving a Hostname into Addresses](#), [name-server](#), and [show host](#).]

- **Uplink failure detection debounce interval (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches)**—Starting with Junos OS Release 19.2R1, you can configure the debounce interval, which is an amount of time, in seconds, that elapses before the downlink interfaces are brought up after corresponding state change of the uplink interfaces. In the absence of a debounce interval configuration, the downlink interfaces are brought up immediately after a state change of the uplink interfaces, which

might introduce unnecessary state changes of the downlink interfaces, as well as unnecessary failovers on the servers connected to these ports.

You can configure the **debounce-interval** statement at the **[edit protocols uplink-failure-detection group group-name]** hierarchy level.

[See [Uplink Failure Detection](#).]

Junos OS XML, API, and Scripting

- **Automation script library additions and upgrades (QFX Series)**—Starting in Junos OS Release 19.2R1, devices running Junos OS that support the Python extensions package include new and upgraded Python modules. Python automation scripts can leverage new on-box Python modules, including the **requests**, **chardet**, and **urllib3** modules, as well as upgraded versions of the **idna**, **ipaddress**, and **six** modules. The Requests library provides additional methods for supporting initial deployments as well as for performing routine monitoring and configuration changes on devices running Junos OS.

[See [Overview of Python Modules Available on Devices Running Junos OS](#) and [Using the Requests Library for Python on Devices Running Junos OS](#).]

Junos Telemetry Interface

- **Sensor level statistics support on Junos Telemetry Interface (JTI) (MX960, MX2008, MX2010, MX2020, PTX5000, PTX1000, and PTX10000 routers and QFX5100 and QFX5200 switches)**—Starting with Junos OS Release 19.2R1, you can issue the Junos operational mode command **show network-agent statistics** to provide more information on a per-sensor level for statistics being streamed to an outside collector by means of remote procedure calls (gRPC) and JTI. Only sensors exported with gRPC are supported. The command does not support UDP-based sensors.

[See [show network-agent statistics](#) and [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **ONCE mode supported using gNMI services and JTI (QFX Series)**—Starting in Junos OS Release 19.2R1, you can include the "ONCE" mode with the **Subscribe** RPC when subscribing to gRPC Network Management Interface (gNMI) services to export statistics for telemetry monitoring and management using Junos telemetry interface (JTI). ONCE mode ensures that the collector is only streamed telemetry information one time at initial connection establishment? .

The subscribe RPC and subscription parameters are defined in the [gnmi.proto](#) file.

Streaming telemetry data through gNMI 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 JTI.

[See [Understanding OpenConfig and gRPC on Junos Telemetry Interface](#).]

- **Packet Forwarding Engine statistics export using gNMI and JTI (QFX5100 and QFX5200 switches)**—Starting in Junos OS Release 19.2R1, you can stream Packet Forwarding Engine statistics to an outside collector using gRPC Management Interface (gNMI) version 0.7.0 and Junos telemetry interface (JTI). Prior to this, these statistics were exported using OpenConfig gRPC and UDP protocol buffer (gpb)

format. OpenConfig gRPC and gNMI are both protocols used to modify and retrieve configurations as well as export telemetry streams from a device in order to manage and monitor it

To provision Packet Forwarding Engine sensors to export data through gNMI, use the subscribe RPC defined in the [gnmi.proto](#) to specify request parameters. This RPC already supports Routing Engine statistics to be exported by means of gNMI. Now, Packet Forwarding Engine sensors will also stream KV pairs in gNMI format for a majority of Packet Forwarding Engine sensors.

Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **gNMI support extended for JTI (QFX5110, QFX5120, QFX5200, and QFX5210 switches)**—Starting in Junos OS Release 19.2R1, Junos telemetry interface (JTI) sensor support extends the ability to configure the following resource paths that use gRPC for export to also use gRPC Management Interface (gNMI) for export. gNMI is a protocol for configuration and retrieval of state information.

JTI supports the following resource paths:

- `/components/component/properties/property/state/value`
- `/components/component/state/`
- `/interfaces/interface/state/`
- `/interfaces/interface/subinterfaces/subinterface/state/`

To provision the sensor to export data through gNMI, use the Subscribe RPC defined in the [gnmi.proto](#) to specify request parameters. Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **gNMI support for Routing Engine statistics for JTI (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 switches)**—Starting in Junos OS Release 19.2R1, Junos telemetry interface (JTI) supports the export of Routing Engine sensors using gRPC Management Interface (gNMI). gNMI is a protocol for configuration and retrieval of state information.

You can use gNMI to export the following statistics:

- LACP state export (resource path `/lacp/interfaces/interface[name='ae1']/members/member/`)
- LLDP statistics (resource path `/lldp/interfaces/interface[name='xe-0/0/9']/`)
- BGP peer information (for example, resource path `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/`)
- RSVP interface statistics (resource path `/junos/rsvp-interface-information/`)

- RPD task memory utilization (resource path `/junos/task-memory-information/`)
- LSP event export (resource path `/junos/task-memory-information/`)

To provision the sensor to export data through gNMI, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gNMI 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 [Guidelines for gRPC Sensors \(Junos Telemetry Interface\)](#).]

- **gNMI support for Packet Forwarding Engine sensors for JTI (QFX5200 switches)**—Starting in Junos OS Release 19.2R1, Junos telemetry interface (JTI) supports the export of Packet Forwarding Engine sensors using gRPC Management Interface (gNMI). gNMI is a protocol for configuration and retrieval of state information.

You can stream the following statistics using gNMI for export:

- Congestion and latency monitoring (resource path `/junos/system/linecard/qmon-sw/`)
- Logical interface usage (resource path `/junos/system/linecard/interface/logical/usage`)
- Filter statistics (resource path `/junos/system/linecard/firewall/`)
- Physical interface statistics (resource path `/junos/system/linecard/interface`)
- LSP statistics (resource path `/junos/services/label-switched-path/usage/`)
- NPU and line-card statistics (resource path `/junos/system/linecard/cpu/memory/`)

To provision the sensor to export data through gNMI, use the Subscribe RPC defined in the [gnmi.proto](#) to specify request parameters. Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **gNMI support for Routing Engine statistics for JTI (QFX5200 switches)**—Starting in Junos OS Release 19.2R1, Junos telemetry interface (JTI) supports export of Routing Engine sensors using gRPC Management Interface (gNMI). gNMI is a protocol for configuration and retrieval of state information. Both streaming and ON_CHANGE export is supported using gNMI.

Export the following statistics using gNMI:

- Network discovery, ARP table state (resource path `/arp-information/`)
- Network discovery, NDP table state (resource paths `/nd6-information/` and `/ipv6-ra/`)

To provision the sensor to export data through gNMI, use the Subscribe RPC defined in the [gnmi.proto](#) to specify request parameters. Streaming telemetry data through gNMI also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the JTI.

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

- **Specify Routing Instance for JTI (ACX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 19.2R1, you can specify the routing instance to use for remote procedure call (gRPC) services. Include the **routing-instance** *instance-name* at the **[edit system services extension-service request-response grpc]** hierarchy level. The routing instance name specified should match the name of the existing routing instance, such as a name configured under the **[routing-instances]** hierarchy level or **mgmt_junos** if **system management-instance** is configured (the dedicated management routing instance).

Configuring the routing instance lets you choose the VRF for gRPC services. When the routing instance is not configured, the default behavior is that all gRPC-related services are available through the management **fxp0/em0** interface.

MPLS

- **Support for MPLS firewall filter on loopback interface (QFX5100, QFX5110, QFX5200, QFX5210)**—Starting with Junos OS Release 19.2R1, you can apply an MPLS firewall filter to a loopback interface on a label-switching router (LSR). For example, you can configure an MPLS packet with **ttl=1** along with MPLS qualifiers, such as **label**, **exp**, and Layer 4 **tcp/udp** port numbers. Supported actions include **accept**, **discard**, and **count**. You configure this feature at the **[edit firewall family mpls]** hierarchy level. You can only apply a loopback filters on **family mpls** in the ingress direction.

[See [Overview of MPLS Firewall Filters on Loopback Interface](#).]

- **Support for IS-IS segment routing (QFX10002-60C)**—Starting in Junos OS Release 19.2R1, you can use IS-IS segment routing through MPLS. Currently, label advertisements are supported for IS-IS only. IS-IS creates an adjacency segment per adjacency, per level, and per address family (one each for IPv4 and IPv6). Junos OS IS-IS implementation allocates node segment label blocks in accordance with the IS-IS protocol extensions for supporting segment routing node segments. It provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:

- **source-packet-routing**—Enable the source packet routing feature.
- **node-segment**—Enable source packet routing at all levels.
- **use-source-packet-routing**—Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and for primary IS-IS source packet routing node segments.
- **no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.

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

Network Management and Monitoring

- **Support for displaying valid user input in the CLI for command options and configuration statements in custom YANG data models (QFX Series)**—Starting in Junos OS Release 19.2R1, the CLI displays the

set of possible values for a given command option or configuration statement in a custom YANG data model when you include the **action-expand** extension statement in the option or statement definition and reference a script that handles the logic. The **action-expand** statement must include the **script** child statement, which defines the Python action script that is invoked when a user requests context-sensitive help in the CLI for the value of that option or statement.

[See [Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules.](#)]

- **Remote port mirroring and remote port mirroring to an IP address (QFX10002-60C switch)**—Starting with Junos OS Release 19.2R1, use port mirroring to copy packets entering or exiting a port or entering a VLAN and send the copies to a VLAN for remote monitoring. You can also send mirrored packets to an IP address over a Layer 3 network (for example, if there is no Layer 2 connectivity to the analyzer device). Use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

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

Routing Policy and Firewall Filters

- **Support for interface, forwarding-class, and loss priority match conditions on egress interfaces (QFX10002-36Q, QFX10002-72Q, QFX10002-60C, QFX10008, QFX10016)**—Starting with Junos OS Release 19.2R1, you can apply the **interface**, **forwarding-class**, and **loss-priority** firewall filter match conditions in the egress direction on IPv4 and IPv6 interfaces. You configure the match conditions at the **[edit firewall]** hierarchy level. This feature was previously supported in an "X" release of Junos OS.

[See [Firewall Filter Match Conditions and Actions \(QFX10000\).](#)]

- **Loopback firewall filter scale optimization (QFX5120)**—Starting with Junos OS Release 19.2R1, you can increase the number of ingress firewall filters on the loopback interface from 384 to 768. To do this, you configure an ingress firewall filter, apply it to the loopback interface, and then use the **loopback-firewall-optimization** command at the **[edit chassis] hierarchy level** to enable optimization. When you configure the loopback filter, you must explicitly specify the terms for **reserved multicast destination** and **ttl** exception packets for this feature to work properly. Enabling or disabling optimization causes the PFE process to restart. This flaps the interfaces, meaning they go up and down, so traffic drops are expected.

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

Routing Protocols

- **Support for 512 ECMP next hops for BGP (QFX10000 switches)**—Starting with Junos OS Release 19.2R1, you can configure a maximum of 512 equal-cost multipath (ECMP) next hops for external BGP peers. (Previously, the maximum number supported was 128.) Having the ability to configure up to 512 ECMP next hops allows you to increase the number of direct BGP peer connections with the QFX10000 switches, thus improving latency and optimizing data flow. Optionally, you can configure those ECMP paths to use consistent load balancing (consistent hashing).

NOTE: This feature applies only to routes for external BGP peers. It does not apply to MPLS routes.

[See [Understanding Configuration of Up to 512 Equal-Cost Paths With Optional Consistent Load Balancing.](#)]

- **Support for export of BGP Local RIB through BGP Monitoring Protocol (BMP) (QFX Series)**—Starting in Junos OS Release 19.2R1, BMP is enhanced to support monitoring of local RIB (**loc-rib**) policy. The **loc-rib** policy is added to RIB types under the **bmp route-monitoring** statement.

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

Software Installation and Upgrade

- **The curl binary is packaged and made available on all Junos OS variants (QFX Series)**—The curl binary is a command-line utility, used from the shell, that you can use to perform operations over several transport protocols, including the following: dict, file, ftp, gopher, http, imap, pop3, rtsp, smtp, telnet, tftp. The features enabled on Junos OS are curl version 7.59, libcurl version 7.59.
- **In-service software upgrade (ISSU) and in-service software reboot (ISSR) (QFX5200 switches)**—Starting with Junos OS Release 19.2R1, you can perform an in-service software upgrade (ISSU) to upgrade between two different Junos OS releases with minimal data and control-plane traffic impact. You can also perform an in-service software reboot (ISSR), which enables you to reset the software state of the system with minimal disruption in data and control traffic.

You can perform an ISSU by issuing the **request system software in-service-upgrade package-name** command.

You can perform an ISSR by issuing the **request system reboot in-service** command.

[See [Understanding In-Service Software Upgrade \(ISSU\).](#)]

System Management

- **Support for transferring accounting statistics files and router configuration archives using HTTP URL (QFX Series)**—Starting in Junos OS Release 19.2R1, you can transfer accounting statistics files and router configuration archives to remote servers by using an HTTP URL. In addition to SCP and FTP, the following HTTP URL will be supported under the **archive-sites** statement:

http://username@host:url-path password password

- To transfer accounting statistics files, configure **archive-sites** under **[edit accounting-options file <filename>]** hierarchy.
- To transfer router configuration archival, configure **archive-sites** under **edit system archival configuration** hierarchy.

- To view the statistics of transfer attempted, succeeded, and failed, use the **show accounting server statistics archival-transfer** command.
- To clear the statistics of transfer attempted, succeeded, and failed, use the **clear accounting server statistics archival-transfer** command.

[See [archive-sites](#), [Backing Up Configurations to an Archive Site](#), [show accounting server statistics archival-transfer](#), and [clear accounting server statistics archival-transfer](#).]

SEE ALSO

What's Changed 253
Known Limitations 259
Open Issues 261
Resolved Issues 265
Documentation Updates 287
Migration, Upgrade, and Downgrade Instructions 288

What's Changed

IN THIS SECTION

- [What's Changed in Release 19.2R3 | 254](#)
- [What's Changed in Release 19.2R2 | 254](#)
- [What's Changed in Release 19.2R1 | 257](#)

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

What's Changed in Release 19.2R3

Class of Service

- **Priority-based flow control (PFC) support (QFX5120-32C)**—Starting with JunosOS 19.2R3, QFX5120-32C switches support priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic.

EVPN

- **IGMP snooping in EVPN-VXLAN multihoming environments (QFX5110)**—In an EVPN-VXLAN multihoming environment on QFX5110 switches, you can now selectively enable IGMP snooping only on those VLANs that might have interested listeners. In earlier releases, you must enable IGMP snooping on all VLANs associated with any configured VXLANs because all the VXLANs share VXLAN tunnel endpoints (VTEPs) between the same multihoming peers and require the same settings. This is no longer a configuration limitation.

Interfaces and Chassis

- **Autonegotiation status displayed correctly (QFX5120-48Y)**—In Junos OS Release 19.2R3, the **show interfaces interface-name <media> <extensive>** command displays the autonegotiation status only for the interface that supports autonegotiation. This is applicable when the switch operates at 1-Gbps speed. In the earlier Junos OS releases, incorrect autonegotiation status was displayed even when autonegotiation was disabled.

Juniper Extension Toolkit (JET)

- **Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)**— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

See [traceoptions \(Services\)](#).

What's Changed in Release 19.2R2

General Routing

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

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

Interfaces and Chassis

- **Support for creating Layer 2 logical interfaces independently (ACX Series, EX Series, MX Series, PTX Series, and QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, 19.1R1, 19.1R2, 19.2R2, and later, QFX

Series switches support creating Layer 2 logical interfaces independent of the Layer 2 routing-instance type. That is, you can configure and commit the Layer 2 logical interfaces separately and add the interfaces to the bridge domain or Ethernet VPN (EVPN) routing instance separately. Note that the Layer 2 logical interfaces work fine only when they are added to the bridge domain or EVPN routing instance.

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

- **Logical Interface is created along with physical Interface by default (QFX Series switches)**—In Junos OS Release 19.2R2 and later, logical interface is created on **ge**, **et**, **xe** interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces are created.

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

Layer 2 Features

- **New statement `input-native-vlan-push` (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—From Junos OS Release 19.2R2, the configuration statement `input-native-vlan-push` at the `[edit interfaces interface-name]` hierarchy level is introduced. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement `input-vlan-map` with a `push` operation is configured.

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

Management

- **`entPhysicalTable` fetched on QFX10002**—In Junos OS Release 19.2R2, the MIB data for `entPhysicalTable` will be fetched on a QFX10002-72Q or QFX10002-36Q switch.

[See [SNMP Explorer](#).]

Multicast

- **Multicast Layer 2 transit traffic statistics by multicast source and group (EX4600, EX4650, and the QFX5000 line of switches)**—Starting in Junos OS Release 19.2R2, EX4600, EX4650, and the QFX5000 line of switches provide statistics on the packet count for each multicast group and source when passing multicast transit traffic at Layer 2 with IGMP snooping. Run the **`show multicast snooping route extensive`** CLI command to see this count in the **Statistics: ... n packets** output field. The other statistics in that output field, **kBps** and **pps**, are not available (values displayed there are not valid statistics for multicast traffic at Layer 2). In earlier releases, all three values in the **Statistics** output field for **kBps**, **pps**, and **packets** did not provide valid statistics for multicast traffic at Layer 2.

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

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 19.2R2, we have included an XML RPC equivalent for the `show bgp output-scheduler | display xml rpc` CLI command. In Junos OS releases before Release 19.2R2, 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 (QFX Series)**—You can now configure a maximum of 2,147,483,647 seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

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

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

What's Changed in Release 19.2R1

EVPN

- **Support for disabling automatic ESI generation (MX Series and QFX Series)**—Starting with Junos OS Release 19.2R1, Junos OS supports disabling the automatic ESI generation for virtual gateway addresses. We recommend that you disable the automatic ESI generation for EVPN networks with edge-routed bridging to improve performance. To disable automatic ESI generation, include the **no-auto-virtual-gateway-esi** statement at the **[edit interfaces name irb unit logical-unit-number]** hierarchy level.

Interfaces and Chassis

- **The resilient-hash statement is no longer available under aggregated-ether-options (QFX5200 and QFX5210 switches)**—Starting in Junos OS Release 19.2R1, the **resilient-hash** statement is no longer available at the **[edit interfaces aex aggregated-ether-options]** hierarchy level. Resilient hashing is not supported on LAGs on QFX5200 and QFX5210.

[See [aggregated-ether-options](#).]

- **Logical interfaces created along with physical interfaces by default (QFX10000 and QFX5000 switches)**—On the QFX10000 line of switches, logical interfaces are created along with the physical et-, sxe-, xe-, and channelized xe- interfaces. In earlier releases, only physical interfaces are created.

On the QFX5000 line of switches, by default, logical interfaces are created on channelized xe- interfaces. In earlier releases, logical interfaces are not created by default on channelized xe- interfaces (xe-0/0/0:1, xe-0/0/0:2, and so on), but they are created on et-, sxe-, and nonchannelized xe- interfaces.

- **Health check for power supplies (QFX10008)**—Starting in Junos OS Release 19.2R1, on the QFX10008 switches, the **show chassis environment pem** command displays the health check information about the DC or AC Power supplies. For any power supply that does not support health check, the status is shown as **Unsupported**. The system starts health check of a power supply only if the power consumption exceeds 7 KW.

[See [show chassis environment pem](#)]

- **Deprecation of the [edit fabric protocols bgp] hierarchy level (QFX Series)**—Starting in Junos OS Release 19.2R1 and later, the **[edit fabric protocols bgp]** hierarchy level is deprecated.

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (QFX Series)**—Starting in Junos OS Release 19.2R1, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.
- **Custom YANG RPC support for input parameters of type empty (QFX Series)**—Starting in Junos OS Release 19.2R1, custom YANG RPCs support input parameters of type **empty** when executing the RPC's command in the Junos OS CLI, and the value passed to the action script is the parameter name. In earlier releases, input parameters of type **empty** are supported only when you execute the RPC in a NETCONF or Junos OS XML protocol session, and the value passed to the action script is the string **'none'**.

[See [Creating Action Scripts for YANG RPCs on Devices Running Junos OS](#).]

Services Applications

- **Commit check for incomplete tunnel encapsulation configuration on flexible tunnel interface (FTI) (QFX Series)**—Tunnel encapsulation configuration is mandatory for FTI interfaces. In Junos OS Release 19.2R1, when you try to commit any incomplete tunnel encapsulation configuration on an FTI, the CLI displays a commit error message.

Security

- **Firewall warning message (QFX5000 switches)**—Starting in 19.2R1, a warning message is displayed whenever a firewall term includes **log** or **syslog** with the **accept** filter action.

SEE ALSO

[What's New | 243](#)

[Known Limitations | 259](#)

[Open Issues | 261](#)

[Resolved Issues | 265](#)

[Documentation Updates | 287](#)

[Migration, Upgrade, and Downgrade Instructions | 288](#)

Known Limitations

IN THIS SECTION

- [EVPN | 259](#)
- [Layer 2 Features | 259](#)
- [Platform and Infrastructure | 259](#)
- [Routing Protocols | 260](#)

This section lists known limitations and system maximums in hardware and software in Junos OS Release 19.2R3 for the QFX Series.

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

EVPN

- Commit is blocked for VLAN-ID none with EVPN routing-instance and without routing-instance. [PR1287557](#)

Layer 2 Features

- Targeted-broadcast forward-only does not broadcast the traffic. [PR1359031](#)
- On the QFX5000 switches, the following message is reported in logs: **fpc0 Pools exhausted for Table:EGR_DVP_ATTRIBUTE_1**. [PR1479826](#)

Platform and Infrastructure

- On the QFX10016 switches, high interrupt for **swi1: netis** and **irq272: vir** with IPv4 or IPv6 unicast traffic, and telnet or ssh to the box do not work. [PR1332337](#)
- Media type displays as **fiber** for DAC. [PR1336416](#)
- After configuring and deleting the Ethernet loopback configuration, the interface goes down and does not come up. [PR1353734](#)
- Scaled VLAN configuration removal might cause high fxpc usage. [PR1363896](#)

- With 288000 MAC scale, the Routing Engine command output displays the learned scale entries after a delay of around 60 seconds. [PR1367538](#)
- The following error message is displayed in the error log in the vty mode:
nh_unilist_update_weight:2541NH: Failed to inc re-route counters for nh. [PR1387559](#)
- The mgmt port speed displays as 1 Gbps with peer interface speed set to 100 Mbps or 10 Mbps. [PR1401382](#)
- On the QFX5120-32C switches after the USB and PXE upgrade, the system boots from USB and PXE respectively and not from SSD. [PR1404717](#)
- Traffic issues are observed with FEC Scaling. [PR1406107](#)
- On the QFX5000 switches, the packets per second and bytes per second statistics for GRE interface are not supported. [PR1419321](#)
- GRE logical interface statistics might get reset when the ECMP child members get changed. [PR1421069](#)
- The following error message is observed while performing NSSU: **syntax error: request-package-validate message.** [PR1421378](#)
- VLAN does not get deleted in the hardware on disabling IRB, leading to ARP getting refreshed even though IRB is disabled. [PR1421382](#)
- The QFX5200 switch goes into amnesiac mode after downgrading from Junos OS Release 19.2 to 17.2X75-D42. [PR1427984](#)
- Memory spike or leakage is observed after the image upgrades to 19.2R1.8 in a mixed mode Virtual Chassis. [PR1464062](#)
- On the QFX10002-60C switches, 100 percent of Layer 2 MAC scaling traffic loss is observed after loading the EVPN-VXLAN collapsed profile configurations. [PR1489753](#)

Routing Protocols

- Maximum number of MPLS tunnels supported is max_mpls_tunnels - 1, as one label is reserved for explicit-null case. [PR1418733](#)
- On the QFX5120-48Y switches, the **show firewall** output gets stuck for around 1 minute when IPv6 lo filter is removed and reapplied. [PR1428087](#)

SEE ALSO

[What's New | 243](#)

[What's Changed | 253](#)

[Open Issues | 261](#)

[Resolved Issues | 265](#)[Documentation Updates | 287](#)[Migration, Upgrade, and Downgrade Instructions | 288](#)

Open Issues

IN THIS SECTION

- [EVPN | 261](#)
- [Infrastructure | 262](#)
- [Interfaces and Chassis | 262](#)
- [Layer 2 Features | 262](#)
- [Platform and Infrastructure | 262](#)
- [Routing Protocols | 264](#)
- [Virtual Chassis | 264](#)

This section lists the open issues in hardware and software in Junos OS Release 19.2R3 for the QFX Series switches.

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

EVPN

- On the QFX5110 switches, after restarting l2ald, traffic duplication for Layer 2 intrapod traffic is observed and it is not converged. [PR1384022](#)
- EVPN-VXLAN core isolation does not work upon system reboot or restart routing. [PR1461795](#)

Infrastructure

- The ftp error message is observed. [PR1315605](#)

Interfaces and Chassis

- On the QFX5110 MLAG switches, flooding of multicast packets for around 16-20 seconds is observed after disabling and enabling a member link of ICL after reboot. [PR1422473](#)

Layer 2 Features

- Adding one more subinterface Logical interface to an existing interface causes 20-50 milliseconds traffic drop on the existing logical interface. [PR1367488](#)
- Q-in-Q might malfunction if **vlan-id-lists** are configured. [PR1395312](#)

Platform and Infrastructure

- Configurations with IRB interfaces might produce Packet Forwarding Engine error messages during the interface deletions. [PR1054798](#)
- On the QFX10000 switches, traffic does not cover upto 100 percent after performing interface unshut even though all the joins gets reinstalled during VLAN replication with more than 2000 VLAN or 2000 IRB in the system.
- Port LEDs do not work on the QFX5100 switch in QFX5110-QFX5100 mixed mode Virtual Chassis. [PR1317750](#)
- In some cases, filter operation with log action gets logged with **Protocol 0 not recognized** in the firewall logs. [PR1325437](#)
- BFD session over an aggregated Ethernet flaps when the member link carrying the BFD transformation flaps. [PR1333307](#)
- The media type is displayed as **fiber** for DAC. [PR1336416](#)
- Backup Routing Engine might crash after multiple continuous GRES. [PR1348806](#)
- The QFX10000 switch drops the Aruba wireless AP heartbeat packets. [PR1352805](#)
- The mib2d process generates core file in **mib2d_write_snmpidx** at **snmpidx_sync.c** on both ADs while bringing up the base traffic profile. [PR1354452](#)
- Filter action to change VLAN does not work. [PR1362609](#)
- The following error messages are observed when channelization is deleted in the AS7816-64X product:
[Err] 0:_pm4x25_line_side_phymod_interface_get: [Wed Jun 13 08:22:45.845 LOG: Err] ERROR: u=0 p=81 interface type 16 not supported by internal SERDES for this speed 50000. [PR1366137](#)

- On the QFX1000 switches, **expr_nh_fwd_get_egress_install_mask:nh** type indirect of **nh_id: #** is invalid. [PR1367121](#)
- On the QFX5200 Virtual Chassis, **request virtual-chassis vc-port diagnostics optics** followed by **show virtual-chassis vc-port diagnostics optics** does not display information from members apart from the master. [PR1372114](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- On the QFX5110 Virtual Chassis, auto configured VCP links might not come up if the existing member, which is connected through auto configured VCP is removed and added back. [PR1375913](#)
- On the QFX5110 switches, the FEC error count does not get updated. [PR1382803](#)
- MLD snooping as MLD membership information gets overwritten to only the latest receiver if the stream group is the same with different source MLD version 2. [PR1386440](#)
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface with the Mini-PDT base configurations. [PR1388082](#)
- On the QFX5000, the DCPFE does not come up in some instances of abrupt power-off or power-on. [PR1393554](#)
- Layer 2 multicast/broadcast convergence is high while deleting and adding the scale configurations of VLANs and VXLAN back. [PR1399002](#)
- DCPFE crashes at **expr_l2_common_bd_ifbd_attach()**. [PR1401215](#)
- On the QFX10002 switches, traffic drop is observed with the MSTP configuration (65 instances and 64 interfaces with 3840 VLANs). [PR1408943](#)
- The following error message is observed after multiple triggers: **JTASK_IO_CONNECT_FAILED**. [PR1408995](#)
- Intermittently chassis alarms might not be raised after the power-cycle of the device. [PR1413981](#)
- On the QFX5110 and QFX5120, uRPF in strict mode does not work. [PR1417546](#)
- On the QFX10000, analyzer does not mirror after adding the child member to an aggregated Ethernet interface. [PR1417694](#)
- IP fragmentation over GRE tunnel on the QFX5000 switches is not supported. [PR1420803](#)
- MC-LAG MAC address synchronization does not occur for the local MACs when a new primary Lo0 IP is added and removed. [PR1424013](#)
- During switchover with LACP link protection configuration, 10 seconds of packet loss is observed. [PR1431034](#)
- The l2ald memory leakage is observed when the VTEP or overlay BGP session flaps. [PR1435561](#)
- On the QFX5200 switches, ISSU might fail. [PR1438690](#)
- The dcpfe process crashes during ISSU at **..bcm_esw_field_entry_reinstall (unit=0, entry=2609)**. [PR1440288](#)

- On the QFX5000 switches, port qualifier must be added. [PR1440980](#)
- After running the ISSR system back to back, the system is left with 2 VMs. [PR1442490](#)
- On the QFX10000 switches, removal of the EVPN-VXLAN Layer 3 Gateway on the IRB interface from the spine switches might silently discard the traffic. [PR1446291](#)
- The interface still sends mirrored traffic out even after it is removed from the output VLAN. [PR1452459](#)
- On the QFX5110 switches, VXLAN VNI (mcast) scaling traffic issue is observed. [PR1462548](#)
- System might stop new MAC learning and impact the Layer 2 traffic forwarding. [PR1475005](#)
- On the QFX10002-60C switches with MSTP scale that has 2 interfaces, 64 instances, and 3840 VLANs in which each MSTI is associated with 60 VLANs, traffic loss is observed. [PR1491161](#)
- On the QFX5000 switches with VXLAN enabled, the ARP request might get dropped if the storm control is configured. [PR1515254](#)

Routing Protocols

- On the QFX5100 switches, the FXPC process generates core file after reconfiguring the IS-IS overload bit. [PR1123116](#)
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, the following error message is observed: **Error BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(),128:l3 nh 6594 unintsall failed in h/w with Mini-PDT base configurations.** [PR1407175](#)
- On the QFX5110 switches, the following error message is observed after rebooting the node: **MCLAG: L2_L3_INTF_OPS_ERROR.** [PR1435314](#)
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, traffic loss on multiple traffic streams after reboot and interface flapping of the Virtual Chassis node are observed. [PR1500508](#)
- The RIP policy with multiple next hop does not work. [PR1515725](#)

Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on backup are not reported to the Routing Engine. [PR1490552](#)

SEE ALSO

[What's New | 243](#)

[What's Changed | 253](#)

[Known Limitations | 259](#)

[Resolved Issues | 265](#)

[Documentation Updates | 287](#)

[Migration, Upgrade, and Downgrade Instructions | 288](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 19.2R3 | 265](#)
- [Resolved Issues: 19.2R2 | 267](#)
- [Resolved Issues: 19.2R1 | 281](#)

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

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

Resolved Issues: 19.2R3

Class of Service (CoS)

- Traffic might be forwarded to the wrong queue when a fixed classifier is used. [PR1510365](#)

EVPN

- The ESI of IRB interfaces does not get updated after the autonomous-system number change if the interface is down. [PR1482790](#)
- The l2ald memory leak might be observed in any EVPN scenario. [PR1498023](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)

Interfaces and Chassis

- MC-LAG consistency check fails if multiple IRB units are configured with same VRRP group. [PR1488681](#)

Layer 2 Features

- The FPC goes down when 100-Gigabit Ethernet link comes up on the network port of the QFX5110-48S switches. [PR1499422](#)

Layer 2 Ethernet Services

- Issues with DHCPv6 relay processing confirm and reply packets are observed. [PR1496220](#)
- The MC-LAG might become down after disabling and then enabling the force-up. [PR1500758](#)

MPLS

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

Platform and Infrastructure

- Traceroute monitor with MTR version v.69 shows a false 10 percent loss. [PR1493824](#)
- ARP replies might be flooded through the EVPN-VXLAN network as unknown unicast ARP reply. [PR1510329](#)
- BFD sessions flap after deactivating or activating the aggregated Ethernet interface or executing GRES. [PR1500798](#)
- Virtual Chassis is not stable with 100-Gigabits Ethernet as well as 40-Gigabits Ethernet interface. [PR1497563](#)
- The following error message is generated at booting up: **CMQFX: Error requesting SET BOOLEAN, illegal setting 66.** [PR1385954](#)
- A VM might generate a core file. [PR1421250](#)
- SFP-LX10 stays down until auto negotiate is disabled. [PR1423201](#)
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB. [PR1442587](#)
- On the QFX10000 switches, the **show forwarding-options enhanced-hash-key** statement does not work. [PR1462519](#)
- On the QFX5100 switches, the interface output counter is double counted for self-generated traffic. [PR1462748](#)
- On the QFX5100 switches, traffic loss might be seen with framing errors or runts if MACsec is configured. [PR1469663](#)
- On the QFX5000 switches, the DSCP marking might not work as expected if the fixed classifiers are applied to interfaces. [PR1472771](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)

- The sFlow could not work correctly if the received traffic goes out of more than one interface. [PR1475082](#)
- The FPC major error is observed after the system boots up or restarts. [PR1475851](#)
- Egress port mirroring might not work when the analyzer port and mirrored port belong to a different FPC. [PR1477956](#)
- BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)
- On the QFX5120 switches, traffic loss could be seen under MC-LAG scenario. [PR1494507](#)
- ARP might not get refreshed after timeout. [PR1497209](#)
- Outbound SSH connection flap or memory leakage issue might be observed while pushing configuration to the ephemeral database with a high rate. [PR1497575](#)
- On the QFX5210064C switches, the lcmd process generates a core file. [PR1497947](#)
- Traffic might get dropped if an aggregated Ethernet member interface is deleted or added, or an SFP of the aggregated Ethernet member interface is unplugged or plugged. [PR1497993](#)
- On the QFX5210 switches, unexpected behavior for port LEDs lights is observed after the upgrade. [PR1498175](#)
- On the QFX10002 and QFX10008 switches, the inter-VNI and intra-VNI, or VRF traffic is dropped between CE devices when the interfaces are connected between top of rack and multihomed with disabled PE devices. [PR1498863](#)
- On the QFX5000 switches, ERPS might not work correctly. [PR1500825](#)
- LLDP is not acquired when **native-vlan-id** and tagged VLAN-ID are the same on a port. [PR1504354](#)
- The l2cpd crash might be seen if the ERP configuration is added or removed, and l2cpd is restarted. [PR1505710](#)

Routing Protocols

- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- On the QFX10000 switches, the rpd process might crash due to rpd resolver problem of INH. [PR1494005](#)
- Traffic drop might be observed after modifying FBF firewall filter. [PR1499918](#)
- Scale of filters with **egress-to-ingress** command is enabled. [PR1514570](#)

User Interface and Configuration

- The version information under configuration changes from Junos OS Release 19.1. [PR1457602](#)

Resolved Issues: 19.2R2

Class of Service (CoS)

- Shaping does not work after the reboot if **shaping-rate** is configured. [PR1432078](#)

- The traffic is placed in the network-control queue on the extended port even if it comes in with different DSCP marking. [PR1433252](#)
- When you move a unicast traffic to a multicast queue through the MF classifier, the **show interface queue <>** statement does not display any statistics on the QFX5120-32C. [PR1459281](#)

EVPN

- Unexpected next-hop operation error from kernel to l2ald in a Layer 2 gateway is observed during the MAC movement operation. [PR1430764](#)
- The VXLAN function might not work due to a timing issue. [PR1502357](#)
- ARP table and Ethernet switching table become asynchronous if the EVPN ESI link flaps multiple times. [PR1435306](#)
- The **multihomed mac-ip** table entry might not be cleaned when the host MAC is deleted from the MAC table. [PR1436712](#)
- ESI configured on a single-homed 25-Gigabit Ethernet port might not work. [PR1438227](#)
- When you use **no-arp-suppression**, an ARP request might not be sent out when an ARP entry ages out. [PR1441464](#)
- The MAC and IP entries might not get cleared after issuing the **clear** command. [PR1446957](#)
- On the QFX5100 line of switches, when there is a VXLAN with VLAN-ID of 2, ARP does not get resolved. [PR1453865](#)
- An ARP request or NS might be sent back to the local segment by the DF router. [PR1459830](#)
- The rpd might crash after changing the EVPN-related configuration. [PR1467309](#)

Forwarding and Sampling

- The dfwd might generate core files and commit error when you apply a firewall filter with the **then traffic-class** or **then dscp** action. [PR1452435](#)
- You might observe memory leakage with the l2ald process. [PR1455034](#)
- Type 1 ESI/AD route might not be generated locally on the EVPN PE devices in the all-active mode. [PR1464778](#)

High Availability and Resiliency

- The QFX5000 line of switches for some versions does not support ISSU. [PR1472183](#)

Interfaces and Chassis

- The VRRP IPv6 state might flap between init and idle states after configuring vlan-tagging. [PR1445370](#)
- On the QFX10000 line of switches, the ARP entries might not be synchronized between the MC-LAG devices. [PR1449806](#)
- The traffic might be forwarded to the wrong interfaces in an MC-LAG scenario. [PR1465077](#)

- The vrrpv3mibs does not poll the VRRPv6-related objects. [PR1467649](#)
- If the dcd process becomes nonresponsive, executing commit might not work. [PR1470622](#)
- Commit error does not thrown an error when a member link is added to the multiple aggregation group with different interface-specific options. [PR1475634](#)

Junos Fusion for Enterprise

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

Junos Fusion for Provider Edge

- On the QFX10000 line of switches, the support for Junos fusion is deprecated. [PR1448245](#)

Junos Fusion Satellite Software

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

Layer 2 Features

- Storm control configuration might be disabled for the interface. [PR1354889](#)
- On the QFX5000 line of switches, you can configure the hash function. [PR1397229](#)
- Packet loss might be observed when one of the spine switches fails or reboots. [PR1421672](#)
- Ethernet Ring Protection Switching nodes might not converge to an **Idle** state after the recovery or reboot fails. [PR1431262](#)
- On the QFX5100 line of switches, the interface driver might not be initialized correctly, which causes errors when the system halts. [PR1434687](#)
- On the QFX5100 and QFX5200 line of switches, the transit DHCPv6 packets might get dropped. [PR1436415](#)
- On the QFX5100 and QFX5110 line of switches, the physical layer and MAC/ARP learning might not work for copper base SFP-T. [PR1437577](#)
- On the QFX5120 line of switches, the MAC learning might not work correctly. [PR1441186](#)
- The operational status of the interface in the hardware and software might be out of synchronization in an EVPN setup with the **arp-proxy** feature enable. [PR1442310](#)
- The traffic leaving the QFX5000 line of switches might not be load-balanced properly over the aggregated Ethernet interfaces. [PR1448488](#)
- Unequal LAG hashing might occur. [PR1455161](#)
- The LLDP function might fail when a device running Junos OS connects to another device that does not run Junos OS. [PR1462171](#)
- On the QFX5000 line of switches, a few MAC addresses might be missed from the MAC table. [PR1467466](#)
- The fxpc might generate core files when you commit the configuration all together. [PR1467763](#)

- In the EVPN-VXLAN scenario, the ingress traffic might discard the traffic silently if the underlying interfaces flap. [PR1469596](#)
- Traffic might get affected if the composite next hop is enabled. [PR1474142](#)

Layer 2 Ethernet Services

- LACP PDU might be looped toward the peer MC-LAG nodes. [PR1379022](#)
- The DHCP DECLINE packets are not forwarded to the DHCP server when **forward-only** is set within **dhcp-reply**. [PR1429456](#)
- The **relay-source** command is applicable for the forward-only subscribers as well. [PR1455076](#)

MPLS

- On the QFX5110 line of switches, the Layer 2 circuit traffic might discard the traffic silently at the EVPN SPINE/MPLS LSP TRANSIT device if the VXLAN access interface flaps on the remote PE node. [PR1435504](#)
- On the QFX5000 line of switches, packet loss might occur when the ECMP resilient-hash is enabled. [PR1442033](#)
- The QFX5120 line of switches might drop the tunnel-encapsulated packets if the switch acts as a transit device. [PR1447128](#)
- On the QFX10002 line of switches, the **show mpls static-lsp | display xml** command produces invalid XML. [PR1469378](#)
- Traffic might be silently discarded on the PE devices when the CE devices sends traffic to the PE devices and the destination is resolved with two LSPs through one upstream interface. [PR1475395](#)
- The traffic might get lost over the QFX5100 line of switches acting as a transit PHP node in the MPLS network. [PR1477301](#)

Platform and Infrastructure

- The default VC MAC persistence timer is incorrectly set to 20 seconds instead of 20 minutes. [PR1478905](#)
- On the QFX5210 line of switches, unexpected behavior are observed on the port LED lights post the upgrade. [PR1498175](#)
- On the QFX5100 line of switches, traffic loss might be observed with framing errors or runts if MACsec is configured. [PR1469663](#)
- The host generated packets might get dropped in the EVPN-VXLAN scenario due to the **reject route** policy in Packet Forwarding Engine. [PR1451559](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- FPC major error is displayed after system reboot or FPC restart. [PR1475851](#)
- Kernel memory leak in the virtual-memory occurs because of the flapping interface. [PR1407000](#)
- Certain QFX Series devices are vulnerable to the Etherleak memory disclosure in the Ethernet padding data (CVE-2017-2304). [PR1063645](#)

- Packet drops, replication failure, or ksyncd crash might be observed on the logical system of a Junos OS device after the Routing Engine switchover. [PR1427842](#)
- REST API process becomes nonresponsive when a number of requests come at a high rate. [PR1449987](#)
- The SLAX script might get lost after upgrading the software. [PR1479803](#)
- On the QFX5100-VC line of switches, the following error message is observed: **MacDrainTimeOut and bcm_port_update failed: Internal error**. [PR1284590](#)
- On the QFX10002-60C line of switches, commit must be denied when mixed Layer 2 and Layer 3 or Layer 4 match conditions are configured on a Layer 2 filter. [PR1326715](#)
- On the QFX5100 line of switches, LR4 QSFP might take up to 15 minutes to come up after the VC reboots. [PR1337340](#)
- On the QFX10000 line of switches, the **show forwarding-options enhanced-hash-key** statement does not work. [PR1462519](#)
- Telemetry traffic might not be sent out when the telemetry server is reachable through different routing-instance. [PR1456282](#)
- When powering off an individual FPC, the other FPC Packet Forwarding Engine might go offline too. [PR1344395](#)
- On the QFX5100 and QFX5200 Virtual Chassis platform, the backup member switch might fail to become the master switch after switchover. [PR1372521](#)
- The new CLI command enables the copying of the Open vSwitch Database (OVSDb) to RAM on the Virtual Chassis backup Routing Engine instead of SSD. [PR1382522](#)
- Static default route with the next-table inet.0 does not work. [PR1383419](#)
- The following error message is generated while booting up: **CMQFX: Error requesting SET BOOLEAN, illegal setting 66**. [PR1385954](#)
- QSFP-100GBASE-SR4/LR4 might take longer time to come up after disabling or rebooting the interface. [PR1402127](#)
- On the QFX10000 line of switches, the ping over loopback might not work over TYPE 5 tunnel. [PR1405786](#)
- The QFX5200 or QFX5100 line of switches might not be able to send out control plane traffic to the peering device. [PR1406242](#)
- On the QFX10000 line of switches, no inner VLAN tag is added even with **input-vlan-map push** configured. [PR1407347](#)
- On the QFX5100-96S line of switches, fan failure alarms might be seen after upgrading to Junos OS Release 17.3R1 and later. [PR1408380](#)
- The 10-Gigabit Ethernet fiber interfaces might flap frequently when they are connected to other vendor switches. [PR1409448](#)

- The optic comes with Tx enabled by default. As the port is administratively disabled, the port is stopped but as the port has not been started, it does not disable Tx. [PR1411015](#)
- Storm control does not shut down the mc-ae interface. [PR1411338](#)
- Part of the routes are not be provided into the Packet Forwarding Engine when both IPv4 and IPv6 are used. [PR1412873](#)
- Traffic loss might be observed after the NSSU operation. [PR1418889](#)
- The **show interface** statement indicates **Media type: Fiber** on the QFX5100-48T line of switches that runs the QFX 5e Series image. [PR1419732](#)
- DHCP bindings for the clients might fail when the QFX5000 line of switches acting as the DHCP local server with Layer 2 channelizes the interface. [PR1421110](#)
- IPv6 multicast traffic received on one VC member might be dropped when exiting on another VC 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 the QFX5000 and QFX10000 line of switches, packet drops might be seen for the traffic that has to go over the TYPE-5 overlay tunnel. [PR1423928](#)
- The host-bound traffic might be dropped after committing change configuration related to the prefix-list. [PR1426539](#)
- The dcpfe or Packet Forwarding Engine might not start on the AS7816-64X and QFX5000 Series devices. [PR1426737](#)
- On the QFX5210 line of switches, the received LLDP frames on em0 are not displayed in the LLDP neighbor output. [PR1426753](#)
- Rebooting or halting of the Virtual Chassis member might cause traffic on the RTG link to be down for about 30 seconds. [PR1427500](#)
- On the QFX5100-VCF line of switches, rollback for uncommitted configuration takes 1 hour. [PR1427632](#)
- The dcpfe process might crash and restart in a MC-LAG scenario when the ARP/NDP next hop is changed. [PR1427994](#)
- The jumbo frame size packets are dropped when the maximum MTU is configured. [PR1428094](#)
- The licenses that are used to flag ovsdb on the **show system license** statement are not flagged even though ovsdb is configured and working. [PR1428207](#)
- The l2ald crash is observed after dot1x gets deleted when dot1x and PVLAN (private VLAN) are enabled simultaneously. [PR1428469](#)
- The **global-mac-limit** and **global-mac-ip-limit** might allow more entries than the configured values. [PR1428572](#)

- On the QFX10008 line of switches, after the Routing Engine switchover, the LED status is not set for the missing fan tray. [PR1429309](#)
- On the QFX10000 line of switches with an EVPN-VXLAN scenario, the DHCP-relay might not work. [PR1429506](#)
- On the QFX5110 line of switches with an EVPN-VXLAN scenario, the DHCP-relay might not work. [PR1429536](#)
- CoS rewrite rules applied under an aggregated Ethernet interface might not get affected after NSSU. [PR1430173](#)
- On the QFX10000 line of stitches, traffic impact might be seen with the interface hold-down timer configured. [PR1430722](#)
- The l2cpd process might crash and generate a core file when the interface flaps. [PR1431355](#)
- The following error on a specific Packet Forwarding Engine might cause complete service impact: **SIB Link Error**. [PR1431592](#)
- On the QFX1000 line of switches, the dcpfe might crash on all line cards in a scaled setup. [PR1431735](#)
- All ingress traffic might be dropped on 100m fixed speed port with **no-auto-negotiation** enabled. [PR1431885](#)
- On the QFX5110 and QFX5120 line of switches, the optical power of interface might gradually reduce the optical power for almost 3 minutes after issuing the **request system reboot at now** statement. [PR1431900](#)
- On the QFX10000 line of switches, the Layer 2 traffic drops with an interface MTU lower than 270 bytes. [PR1431902](#)
- 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 a Q-in-Q scenario. [PR1432703](#)
- Line card might crash due to the plug-in of the unsupported SFP-T module. [PR1432809](#)
- On the QFX10000 line of switches using the LC1105 line card, traffic loss might be observed. [PR1433300](#)
- The VC Mezz temp and QIC sensor fails. [PR1433525](#)
- Traffic drop might occur during the filter change operation. [PR1433648](#)
- Layer 3 filters applied to the PVLAN IRB interface might not work after ISSU. [PR1434941](#)
- Traffic drops might be seen when the MACsec session key rolls over between the primary and fallback for more than 10 times. [PR1435277](#)
- On the QFX5100-VC line of switches, there might be approximately 1 to 5 minutes traffic loss during NSSU with the LACP link protection configuration. [PR1435519](#)
- On the QFX10000 line of switches, the SIB or FPC Link error alarms might be observed due to a single CRC. [PR1435705](#)

- The mc-ae interface might get stuck in the **Waiting** state in a dual mc-ae scenario. [PR1435874](#)
- Traffic drop might occur when SXE interface is used. [PR1435963](#)
- DHCP discover packets sent to IP addresses in the same subnet as the IRB interface causes the QFX5110 line of switches to send incorrect traffic, out of the dhcp-snooping enabled interfaces. [PR1436436](#)
- The unknown SNMP trap (1.3.6.1.4.1.2636.3.69.1.0.0.1) is sent when you restart the QFX5110 line of switches. [PR1436968](#)
- The FPC might crash if both the aggregated Ethernet bundle flaps on the local device and the configuration change on the peer device occur at the same time. [PR1437295](#)
- In the QFX5110, QFX5200, and QFX5210 line of switches, there is no jnxFruOK SNMP trap message when you disconnect the power cable and then connect it back. [PR1437709](#)
- The Routing Engine switchover does not work as expected while the SSD failure occurs. [PR1437745](#)
- The BGP neighborhood 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. [PR1438351](#)
- On the QFX5210 line of switches, the port LED turns red when the cable is connected. [PR1438359](#)
- Interfaces configured with **flexible-vlan-tagging** might lose connectivity. [PR1439073](#)
- The xSTP recognizes 1-Gigabit Ethernet SFP-T optic interface as LAN type resulting in a slow STP convergence. [PR1439095](#)
- LACP state might get stuck in the **Attached** state after disabling the peer active members. [PR1439268](#)
- The default logical interface on the channelized IFD might not be created after ISSU/ISSR. [PR1439358](#)
- The VC of the QFX5100 line of switches might not come up after replacing the VC port fiber connection with the DAC cable. [PR1440062](#)
- MAC addresses learned on RTG might not get aged out after the Virtual Chassis member reboots. [PR1440574](#)
- Traffic drop might be seen after disabling and enabling the mc-ae interface in a MC-LAG scenario. [PR1440732](#)
- Interface match conditions of firewall filter might not work on egress direction with the IRB over an aggregated Ethernet. [PR1441230](#)
- The Layer 3 communication might break on an interface that is configured with **flexible-ethernet-services**. [PR1441690](#)
- On the QFX5110 line of switches, the flow control does not work as expected on the 100-Gigabits Ethernet interface. [PR1442522](#)
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB. [PR1442587](#)
- The DHCPv6 client might fail to get an IP address. [PR1442867](#)

- When a line card reboots, the MC-LAG might not get programmed after the line card comes back online. [PR1444100](#)
- On the QFX5200 line of switches, the following error message is observed when you change the UFT profile in the FPC logs: **DCBCM[bcore_init]: ioctl call failed ret:0**. [PR1445855](#)
- On the QFX10008 line of switches, you might observe traffic impact when you use the JSRV interface. [PR1445939](#)
- CoS classifier might not work as expected. [PR1445960](#)
- Long IPv6 addresses are not displayed fully on the IPv6 neighbor table. [PR1447115](#)
- Unicast ARP requests are not replied with the **no-arp-trap** option. [PR1448071](#)
- Rebooting of the QFX5120-48Y line of switches using the **request system reboot** statement does not take the physical links offline immediately. [PR1448102](#)
- On the QFX10000 line of switches with QSFP28 100G AOC, 740-065632 and QSFP+ 40G, or 740-043308 transceiver, the port LED remains green after disconnecting one end. [PR1448121](#)
- The process vehostd might crash without generating a core file and the automatic restart of the vehostd might fail. [PR1448413](#)
- On the VCPs of the QFX5100 VC line of switches, the CRC error might be observed. [PR1449406](#)
- Except for one aggregated Ethernet member link, the other links do not send out sFlow sample packets for the ingress traffic. [PR1449568](#)
- On the QFX10008 line of switches, the FPC0 generates core files after running the Packet Forwarding Engine **show cos sched-usage** command. [PR1449645](#)
- The em0 route might be rejected after the em0 interface is disabled and then enabled again. [PR1449897](#)
- FPC does not restart immediately after rebooting the system, which causes packet loss. [PR1449977](#)
- On the QFX10000 line of switches, the CoS classification does not work. [PR1450265](#)
- On the QFX5000 line of switches, when the dual VLAN tag feature is configured on the physical interface, no error message is displayed. [PR1450455](#)
- Tunneling encapsulated packets are dropped on the Layer 3 VPN MPLS PE-CE interface. [PR1451032](#)
- The DHCP snooping static binding does not get affected after deleting and readding the entries. [PR1451688](#)
- FPC might generate core files after changing the PTP/SyncE configuration. [PR1451950](#)
- On the QFX10008 line of switches, the Packet Forwarding Engine **show cos scheds-per-pfe** and **show cos pfe-scheduler-ifds** commands restart to forward planes. [PR1452013](#)
- Vgd might generate core files when tunnel gets deleted twice. [PR1452149](#)
- There might be an interface reachability issue on AS7816. [PR1452433](#)

- The l2ald and eventd are consumed 100 percent after committing the **clear ethernet-switching table** statement. [PR1452738](#)
- DHCP offer packet with the unicast flag set gets dropped by the QFX10000 line of switches in a VXLAN multihomed setup using anycast IP. [PR1452870](#)
- Configuration change in the VLAN all option might affect the per-VLAN configuration. [PR1453505](#)
- The classifier configuration does not get applied to the interface in an EVPN-VXLAN environment. [PR1453512](#)
- The **show chassis led** statement shows wrong status. [PR1453821](#)
- VGD process consumes the CPU when **switch-options vtep-source-interface lo0.0** is not configured. [PR1454014](#)
- In the EVPN-VXLAN scenario, changing the VLAN name associated with the access ports might prevent the MAC addresses from being learned. [PR1454095](#)
- Master FPC might come up in the **Master** state again after the reboot instead of backup. [PR1454343](#)
- On the QFX5000 line of switches, dcpfe might crash when data that is not correctly NULL terminated is processed. [PR1454527](#)
- On the QFX10002-60C line of switches with EVPN-VXLAN, the MAC+IP count is displayed as 0. [PR1454603](#)
- The untagged ARP/NS requests might not be resolved when the host is connected on the encapsulation ethernet-bridge interface. [PR1454804](#)
- A firewall filter might not be applied in a particular VC/VCF member as the TCAM space runs out. [PR1455177](#)
- In a 16 and more member of the QFX5100 VCF, the **FROM** column under the **show system users** output reports feb0/1/2/3 for fpc16/17/18/19, respectively. [PR1455201](#)
- On the QFX10000 line of switches, the PFC feature does not work. [PR1455309](#)
- The cosd crash might be observed if the **forwarding-class-set** is directly applied on the child interface of an aggregated Ethernet interface. [PR1455357](#)
- Link up delay and traffic drop might be observed on the mixed SP Layer 2 or Layer 3, and EP Layer 2 type configurations. [PR1456336](#)
- The laser from the 10-Gigabits Ethernet SFP+ interface still remains on when the interface is disabled or the device is rebooted. [PR1456742](#)
- The QFX5110 line of switches, the interface on the QSFP-100GBASE-SR4 made by Avago vendor cannot be linked up. [PR1457266](#)
- On the QFX10000 line of switches, the Packet Forwarding Engine process might crash after the Routing Engine switchover. [PR1457414](#)

- Over-temperature SNMP trap messages are displayed after update even though the temperatures are within the system thresholds. [PR1457456](#)
- On the QFX5110 line of switches, the PTP uses port 52 and port 53 but it does not have any FPGA register address. [PR1457516](#)
- Dual tag Q-in-Q does not work with EVPN-VXLAN. [PR1458206](#)
- On the QFX5210 line of switches, the LED does not light on the port 64 and port 65 after upgrading to Junos OS Release 19.2R1. [PR1458514](#)
- The BPDU packet might be looped between leaf DF switch and non-DF switch that causes traffic blocking. [PR1458929](#)
- On the QFX5000 line of switches, the lightweight DHCPv6 relay agent functionality might be broken. [PR1459499](#)
- The fxpc process might crash due to several BGP IPv6 session flaps. [PR1459759](#)
- The **forwarding** option is missed in the routing-instance type. [PR1460181](#)
- On the QFX5000 line of switches, the **accept-source-mac** feature with VXLAN does not work. [PR1460885](#)
- On the QFX10002-72Q or QFX10002-36Q line of switches, the **entPhysicalTable** MIB does not fetch the expected data. [PR1462582](#)
- On the on QFX5000 line of switches, the fxpc process might generate core file when you change MTU in a VXLAN scenario with firewall filters applied. [PR1462594](#)
- While cleaning up EVPN-VXLAN configurations with Mini-PDT base configurations, the following error message is observed: **Error BRCM-VIRTUAL,brcm_vxlan_walk_svp(),6916:Failed to find L2-iff for ifl.** [PR1463939](#)
- On the QFX10000 line of switches, FPC might restart during run-time. [PR1464119](#)
- On the QFX5000 line of switches, the dcpfe might crash when you change the firewall filter. [PR1464352](#)
- On the QFX10000 line of switches, the interface might not come up when FPC restarts. [PR1464650](#)
- On the QFX5100-48S line of switches, when you try to apply a firewall filter that contains a **then dscp** action to a Layer 3 inet subinterface, an error is displayed when you try to commit. Applying the same filter to an IRB interface succeeds as does applying the same filter to a Layer 3 subinterface. [PR1464883](#)
- On the QFX5210 line of switches, PEM is not present spontaneously. [PR1465183](#)
- A 10-Gigabits Ethernet interface might not come up on the QFX5100-48T line of switches or negotiate at speed 1-Gigabits Ethernet when connected with the Broadcom 57800-T daughter card. [PR1465196](#)
- On the QFX5110 or QFX5200 line of switches, the QSFP-100G-PSM4 modules are not correctly identified. [PR1465214](#)
- The physical interface of the aggregated Ethernet interface might take time to come up after disabling or enabling the physical interface. [PR1465302](#)

- The numbering on the AS7816-64X platform changes from 0 to 1 and 1 to 0. The fan numbering also changes from 0,1,2,3 to 3,2,1,0. [PR1465327](#)
- The broadcast and multicast traffic might be dropped over the IRB or LAG interface. [PR1466423](#)
- BGP BMP messages are sent to the BMP collector with BGP optional capabilities truncated. As a result, the BMP collector does not register the correct information exchanged during BGP session establishment. [PR1466477](#)
- On the QFX10000 line of switches, the EBUF parity interrupt is not observed. [PR1466532](#)
- IPv6 traffic might get dropped in the Layer 3 VPN network. [PR1466659](#)
- On the QFX5000 line of switches, slow packet drops might be observed. [PR1466770](#)
- On the QFX10000 line of switches, the following error message might be observed that causes the protocols to go down: **EPR iCRC**. [PR1466810](#)
- A few of the DHCPvX inform messages specific to a particular VLAN do not receive any acknowledgment from the server. [PR1467182](#)
- Ingress drops must be included at the CLI command from the interface statistics and added to InDiscards. [PR1468033](#)
- Optics measurements might not be streamed for interfaces of a PIC over JTI. [PR1468435](#)
- MAC address might not be learned on a new extended port after VMotion in the Junos fusion for data center environment. [PR1468732](#)
- If continuous interface flaps occur at ingress or egress of the PE devices, the IP routed packets might be looped on the MPLS PHP node. [PR1469998](#)
- Incorrect counter values are displayed for the arrival rate and peak rate for the DDoS commands. [PR1470385](#)
- The speed 10m might not be configured on the Gigabit Ethernet interface. [PR1471216](#)
- Traffic loss might occur when the VTEP source interface is configured in the multiple routing instances. [PR1471465](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- The detached interface in LAG might process the xSTP BPDUs. [PR1473313](#)
- The l2ald crash might be seen when around 16,000 VLAN-IDs share the same VXLAN tunnel and Packet Forwarding Engine is rebooted. [PR1473521](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)
- On the QFX5000 line of switches in a EVPN-VXLAN scenario, continuous error log messages might be raised. [PR1474545](#)
- On the QFX5000 line of switches, the Layer 2 circuit might fail to communicate through VLAN 2. [PR1474935](#)

- On the QFX10000 line of switches, the MACsec traffic over Layer 2 circuit might not work after upgrading from Junos OS Release 15.1 and later. [PR1475089](#)
- The DAC cables are not being properly detected in the Packet Forwarding Engine in the QFX5200 line of switches on Junos OS Release 18.4R2-S2.4. [PR1475249](#)
- On the QFX5110 or QFX5120 line of switches, there might be traffic drop acting as a leaf switch in a multicast environment with VXLAN. [PR1475430](#)
- The QFX Series devices exhibit invalid Packet Forwarding Engine PG counter pairs to copy src 0xfffff80, dst 0. [PR1476829](#)
- On the QFX10002-36Q or QFX10002-72Q line of switches, the following continuous error logs on the device are observed: **prds_ptc_wait_adoption_status: PECHIP[1] PTC[1]: timeout on getting adoption valid bit[8] asserted.** [PR1477192](#)
- The remaining interface might be still in the **Down** state even though the number of channelized interfaces is no more than 5. [PR1480480](#)
- The ARP request packets for the unknown host might get dropped in the remote PE devices in the EVPN-VXLAN scenario. [PR1480776](#)
- VLAN creation failure might be observed with scaled VLAN and Layer 3 configuration. [PR1484964](#)
- The BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)

Routing Protocols

- Some storm control error logs might be observed. [PR1355607](#)
- The OSPF VRF sessions take a long time to come up when the host table is full and host routes are in the LPM table. [PR1358289](#)
- Invalid VRRP mastership election on the QFX5110-VC peers is observed. [PR1367439](#)
- Value added in the Hexa after an unknown Ext-Community gets reset to 0. [PR1371448](#)
- Host-destined packets with filter log action might not reach the Routing Engine if the log or syslog is enabled. [PR1379718](#)
- The IRB transit traffic might not be counted for the EVPN or VXLAN traffic. [PR1383680](#)
- On the QFX5100 line of switches, the BGP v4 or BGP v6 convergences, and the RIB installs or deletes time degraded in Junos OS Releases 19.1R1, 19.2R1, and later. [PR1414121](#)
- On the QFX5000 line of switches, the same traffic flow might be forwarded to different ECMP next hops. [PR1422324](#)
- The BGP **multipath multiple-as** statement does not work. [PR1430899](#)
- On the QFX5100 line of switches, the fxpc might generate core files during the reboot of device. [PR1432023](#)
- Ping fails over Type-5 tunnel on the IRB interfaces under the EVPN-VXLAN scenario. [PR1433918](#)

- The IPv4 fragmented packets might be broken if the PTP transparent clock is configured. [PR1437943](#)
- The bandwidth value of the DDoS-protection might cause packet loss after the device reboots. [PR1440847](#)
- Traffic might be dropped after the Q-in-Q enabled interface is flapped or a change is made to the vlan-id-list. [PR1441402](#)
- On the QFX5210 line of switches, the firewall filter DSCP Action Modifier does not work when the firewall filter is mapped to IRB. [PR1441444](#)
- The rpd process might crash in an inter-AS option B Layer 3 VPN scenario if CNHs is used. [PR1442291](#)
- The IPv6 connectivity between the MC-LAG peers might fail when multiple IRB interfaces are present. [PR1443507](#)
- PIM (S,G) joins cause MSDP to incorrectly announce source active messages. [PR1443713](#)
- On the QFX5100 Virtual Chassis, the CRC errors might be observed. [PR1444845](#)
- Loopback address exported into other VRF instance might not work. [PR1449410](#)
- MPLS LDP might still use stale MAC of the neighbor even though the neighbor of the LDP MAC changes. [PR1451217](#)
- Core files might be generated during addition or removal of the EVPN Type-5 routing instance. [PR1455547](#)
- A few seconds of traffic drop might be observed on the existing receivers when another receiver joins or leaves. [PR1457228](#)
- The egress interface in the Packet Forwarding Engine for some end-hosts might not be correct on the Layer 3 gateway switch after it is rebooted. [PR1460688](#)
- The **other querier present interval** timer cannot be changed in the IGMP/MLD snooping scenario. [PR1461590](#)
- When deleting IRB on the Layer 3 gateway, IRB does not get removed from the Packet Forwarding Engine and discards the traffic silently to the IRB MAC address. [PR1463092](#)
- The mcsnoopd crash might be observed if one BD/VLAN is configured as a part of EVPN and if it has any multicast router interfaces (static/dynamic). [PR1468737](#)
- Traffic might not be forwarded over the ECMP link in the EVPN-VXLAN scenario. [PR1475819](#)
- ARP packets are always sent to the CPU regardless of whether the storm-control is activated. [PR1476708](#)
- GRE transit traffic does not get forwarded in the VRRP scenario. [PR1477073](#)
- MUX state in the LACP interface does not go to the **collecting and distributing** state and remains attached after enabling the aggregated Ethernet interface. [PR1484523](#)
- The following multicast statistics related errors are observed during ISSU:
`brcm_ipmc_route_counter_delete:3900Multicast stat destroy failed (-10:Operation still running) y".`
[PR1460791](#)

- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

User Interface and Configuration

- Switch might be unable to commit baseline configuration after returning to zero. [PR1426341](#)

Resolved Issues: 19.2R1

Authentication and Access Control

- Without configuring anything related to dot1x, the syslog **dot1xd[2192]: task_connect: task PNACAUTH./var/run/authd_control addr /var/run/authd_control: Connection refused** is generated repeatedly. [PR1406965](#)

Class of Service (CoS)

- Error message **STUCK_BUFF : port_sp not empty for port 35 sp 1 pkts:1**. [PR1346452](#)

EVPN

- The rpd process might crash with EVPN type-3 route churn. [PR1394803](#)
- EVPN routes might show **Route Label: 0** in addition to the real label. [PR1405695](#)
- The rpd might crash after an NSR switchover in an EVPN scenario. [PR1408749](#)
- ARP entry is still pointing to the failed VTEP after the PE-CE link fails for a multihomed remote ESI. [PR1420294](#)
- Multicast MAC addresses being learned in the Ethernet switching table with VXLAN through an ARP packet in a pure L2 configuration. [PR1420764](#)
- The device might proxy the ARP probe packets in an EVPN environment. [PR1427109](#)
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs. [PR1429821](#)
- ESI is configured on a single-homed 25-gigabit port might not work. [PR1438227](#)

General Routing

- The 1-gigabit copper module interface shows **Link-mode: Half-duplex** on the QFX10000 line of devices. [PR1286709](#)
- Interface flap on 100GBASE-LR4 is seen during an unified ISSU. [PR1353415](#)
- On QFX5120 switches, the convergence delay between PE1 and P router link is more than the expected delay value. [PR1364244](#)
- RIPv2 update packets might not be sent when **IGMP snooping** is enabled. [PR1375332](#)

- The **overlay-ecmp** configuration might not work as expected on QFX5110 in an EVPN-VXLAN environment. [PR1380084](#)
- There is an inconsistency in applying a scheduler map with **excess-rate** on the physical interface and the aggregated Ethernet interface. [PR1380294](#)
- Traffic get silently dropped and discarded When the FPC is taken offline in an MC-LAG scenario. [PR1381446](#)
- Last reboot reason is incorrect if the device is rebooted because of power cycling. [PR1383693](#)
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent the major alarm. [PR1384435](#)
- The configuration statement **show chassis errors active detail** is not supported on QFX5000 platforms. [PR1386255](#)
- The rpd and KRT queue get stuck in a VRF scenario. [PR1386475](#)
- ARP received on an SP-style interface is not sent to all RVTEPs in case of QFX5100 Virtual Chassis only; the normal BUM traffic works fine. [PR1388811](#)
- The input rate (in pps) do not increase on EX2300-MP uplink ports when the packet is a pure L2 packet such as non-etherII or non-EtherSnap. [PR1389908](#)
- 10-gigabit copper link flapping might happen during a TISSU operation of QFX5100-48T switches. [PR1393628](#)
- **BRCM_NH-,brcm_bcm_mpls_tunnel_initiator_clear(),226:bcm_mpls_tunnel_initiator_get failed intf = 4 failure** error logs might seen in syslog. [PR1396014](#)
- On QFX5110 Fan LED turns amber randomly. [PR1398349](#)
- The interrupt process consumes high CPU because of the intr{swi4: clock (0)} on QFX5100-48t-6Q running a QFX5100 Series image and Junos OS Release 18.x code. [PR1398632](#)
- The DHCPv6 relay packets are dropped when both the UDP source and destination ports are 547. [PR1399067](#)
- CPU hog might be observed on the QFX10000 line of switches. [PR1399369](#)
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. [PR1399683](#)
- PEM I2C failure alarm might be shown incorrectly as failed. [PR1400380](#)
- MAC limit with persistent MAC is not working after reboot. [PR1400507](#)
- On QFX5120-32C error logs for flex counter are seen with GRE configuration. [PR1400515](#)
- Only one Packet Forwarding Engine might be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. [PR1400716](#)
- File permissions are changed for **/var/db/scripts** files after reboot. [PR1402852](#)

- The DHCP discover packets are forwarded out of an interface incorrectly if **DHCP snooping** is configured on that interface. [PR1403528](#)
- The VRRP VIP might not work when it is configured on the LAG interface. [PR1404822](#)
- Commit warning message is seen on QFX5100 switches. [PR1405138](#)
- Executing the **request system configuration rescue save** command might fail with error messages. [PR1405189](#)
- DHCP does not work for some clients in dual Junos fusion aggregated device setup on extended ports (EP). [PR1405495](#)
- VXLAN transit traffic over tagged underlay L3 interface and underlay IRB interface gets dropped because of the hardware limitation. [PR1406282](#)
- The ARP request might not be resolved successfully if **arp-suppression** is enabled and **vlan-id-list** is configured on the QFX10000 node. [PR1407059](#)
- The DHCP discover packets might be dropped over VXLAN tunnel if DHCP relay is enabled for other VXLAN or VLANs. [PR1408161](#)
- MAC address movement might not happen in flexible Ethernet services mode when family inet/inet6 and vlan-bridge are configured on the same physical interface. [PR1408230](#)
- Fan failure alarms might be seen on QFX5100-96S after upgrading to Junos OS Release 17.3R1. [PR1408380](#)
- Restarting the line card on QFX10008 and QFX10016 with MC-LAG enhanced-convergence might cause intra-VLAN traffic to get silently dropped and discarded. [PR1409631](#)
- The FPC might crash and might not come up if the interface number or next hop is set to maximum value under **vxlan-routing** on QFX Series platforms. [PR1409949](#)
- LLDP memory leak when IEEE dcbx packet is received in auto-negotiation mode followed by another dcbx packet with none of ieee_dcbx TLVs present. [PR1410239](#)
- On EX2300-24P, error message **dc-pfe: BCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family**. [PR1410717](#)
- Fix jfirmware support to upgrade primary BIOS from a system booted from secondary BIOS. [PR1411603](#)
- Traffic loss might be observed after VXLAN configuration change. [PR1411858](#)
- The spfe on satellite device in Junos fusion setup might crash and it might cause the satellite device to go offline. [PR1412279](#)
- PEM alarm for a backup FPC will remain on the master FPC though the backup FPC is detached from Virtual Chassis. [PR1412429](#)
- Junos PCC might reject PCUpdate or PCCreate message if there is a metric type other than type 2. [PR1412659](#)
- On QFX5000 line of switches, the EVPN/VXLAN mutlicast next-hop limit is 4000. [PR1414213](#)

- VC ports using DAC might not establish a link on QFX5200. [PR1414492](#)
- DC output information is missing in the **show chassis environment pem** output for whitebox. [PR1414703](#)
- VXLAN encapsulation next hop (VENH) does not get installed during BGP flap or restart routing. [PR1415450](#)
- FEC change from FEC91 to NONE does not take effect on 100-gigabit Ethernet interfaces with QSFP-100GBASE-SR4 optics. [PR1416376](#)
- Two instances of Junos OS are running after an upgrade to Junos OS Release 18.1R3-S3.7. [PR1416585](#)
- On restarting routing, the dcpfe might generate a core file at **nh_composite_change**. [PR1416925](#)
- ERSPAN traffic does not tag when output interface is trunk port. [PR1418162](#)
- Traffic loss might be seen on the aggregated Ethernet interface on QFX10000 platforms. [PR1418396](#)
- Rebooting QFX5200-48Y using **request system reboot** does not take physical links offline immediately. [PR1419465](#)
- On QFX5120-48Y or QFX5120-32C, 100-gigabit PSM4 optics connected ports went down randomly. [PR1419826](#)
- Ping fails over type-5 tunnel on IRB interfaces under EVPN-VXLAN scenario. [PR1420785](#)
- Error messages might be seen on QFX10000 platforms during DFE tuning. [PR1421075](#)
- On QFX5120-32C, DHCP binding on client might fail when QFX5120-32C acting as DHCP server is seen only for channelized port. [PR1421110](#)
- ETS configuration does not apply on non cascade ports when the AD is rebooted. [PR1421429](#)
- BFD might get stuck in slow mode on QFX10002, QFX10008, and QFX10016 platforms. [PR1422789](#)
- QFX5100-48T 10-Gbps interface might be auto negotiated at 1-Gbps speed instead of 10-Gbps. [PR1422958](#)
- The interface cannot get up when the **remote-connected interface** only supports 100M in QFX5100 VC setup. [PR1423171](#)
- BUM traffic coming over IRB underlay interface gets dropped on destination VTEP in a PIM-based VXLAN. [PR1423705](#)
- Traffic drops when an FPC reboots with aggregated Ethernet member links deactivated by a remote device. [PR1423707](#)
- Ping over EVPN type-5 route to QFX10000 does not work. [PR1423928](#)
- All interfaces will be down and the dcpfe might crash if SFP-T is inserted in a QFX5210. [PR1424090](#)
- IPv6 neighbor solicitation packets for link-local address are dropped when passing through QFX10002-60C. [PR1424244](#)
- QFX5120 QSFP-100G-PSM4 interfaces are undetected and come back up as channelized interfaces. [PR1424647](#)

- All interfaces creation fails after NSSU. [PR1425716](#)
- Heap memory leak might be seen on QFX10000 platforms. [PR1427090](#)
- On QFX5120-48Y, the interfaces with the QSFP-100GBASE-ER4L optics do not come up in Junos OS Release 18.3R1-S2.1. [PR1428113](#)
- The configuration statement **show chassis environment** shows **Input0** and **Input1**. [PR1428690](#)
- The l2ald process crashes and generates a core file when the number of VXLAN HW IFBDS exceeds the maximum limit of 16,382. [PR1428936](#)
- An interface on a QFX Series switch does not come up after the transceiver is replaced with one having different speed. [PR1430115](#)
- When the IRB interface is trying to broadcast an ARP request, the ARP request might not go out of the chip because of the SDK bug, which might lead to ARP failure in QFX5120. [PR1430327](#)
- On QFX Series switches, the **Validation of meta data files failed** message is seen on the hypervisor. [PR1431111](#)
- Transit DHCPv6 packets might be dropped on QFX5000 platforms. [PR1436415](#)

Interfaces and Chassis

- Changing the value of **mac-table-size** to default might lead all FPCs to reboot. [PR1386768](#)
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces a misleading error message. [PR1402606](#)
- EVPN aggregated Ethernet interface flap followed by a commit. [PR1425339](#)

Junos Fusion Satellite Software

- Extended port (EP) LAG might go down on the satellite devices (SDs) if the related cascade port (CP) links to an aggregation device (AD) goes down. [PR1397992](#)

Layer 2 Ethernet Services

- The malfunction of the core isolation feature in EVPN-VXLAN scenarios causes traffic to get silently dropped and discarded. [PR1417729](#)

Layer 2 Features

- On the QFX Series switches, the error message **Failed with error (-7) while deleting the trunk 1 on the device 0** is seen. [PR1393276](#)
- On the QFX5110-48S switch, the FPC goes down when 100-Gigabits Ethernet link comes up. [PR1499422](#)
- On the QFX5000 line of switches, symmetric hash is observed. [PR1397229](#)
- On the QFX5000 line of switches, dcpfe process crash might be observed during restart of the Packet Forwarding Engine on a system with scaled EVPN/VXLAN configuration. [PR1403305](#)

- On the QFX Series EVPN-VXLAN, the unicast IPv6 NS message gets flooded on L3GW. Both IPv4 and IPv6 traffic gets dropped on L2SW. [PR1405814](#)
- The IPv6 NS/NA packets received over VTEP from an ESI host are incorrectly flooded back to the host. [PR1405820](#)
- **IGMP-snooping** on EVPN-VXLAN might impact OSPF hello packets flooding after a VTEP leaf reboot. [PR1406502](#)
- With cut through configuration enabled after the device is rebooted, cut through mode is disabled on the channelized interfaces in releases before Junos OS Release 19.1R1. [PR1407706](#)
- The QFX5110 Virtual Chassis generates DDoS messages of different protocols on inserting a 1-gigabit or 10-gigabit SFP transceiver or after forming a VCP connection. [PR1410649](#)
- With **arp-suppression** enabled, QFX5000 might not forward IPv6 router solicitations or advertisements packets. [PR1414496](#)

Network Management and Monitoring

- The chassisd might crash and restart after the AGENTX session timeout between master (snmpd) and subagent times out. [PR1396967](#)
- Log files might not get compressed during the upgrade. [PR1414303](#)

Routing Protocols

- Host-destined packets with **filter log** action might reach the Routing Engine. [PR1379718](#)
- BUM packets might get looped if EVPN multihoming interfaces flap. [PR1387063](#)
- EVPN-VXLAN NON-COLLAPSED AUTONEG errors and flush operation failed errors are seen after the device is power cycled. [PR1394866](#)
- On the QFX5110 and QFX5200 line of switches EVPN-VXLAN **non-collapsed** state, the dcfpe process generates a core file at `brcm_pkt_tx_flush`, `l2alm_mac_ip_timer_handle_expiry_event_loc` after a random event. [PR1397205](#)
- The FPC or dcfpe process might crash because of the interface flapping. [PR1408428](#)
- The rpd crashes on static route configuration for multicast source. [PR1408443](#)
- The ERACL firewall group operates in double wide mode for QFX5110 in Junos OS Release 19.1R1. [PR1408670](#)
- Host-generated ICMPv6 RA packets might be dropped on the backup member of Virtual Chassis if **igmp-snooping** is configured. [PR1413543](#)
- The QFX Series switches might not install all IRB MAC addresses in the initialization. [PR1416025](#)
- After an IRB logical interface is deleted, MAC entry for the IRB interface is deleted for the IRB hardware address; as a result, packets destined to other IRB logical interfaces where MAC is not configured are impacted. [PR1424284](#)

Spanning Tree Protocols

- The l2cpd might crash if the VSTP traceoptions and VSTP VLAN all commands are configured. [PR1407469](#)

SEE ALSO

- [What's New | 243](#)
- [What's Changed | 253](#)
- [Known Limitations | 259](#)
- [Open Issues | 261](#)
- [Documentation Updates | 287](#)
- [Migration, Upgrade, and Downgrade Instructions | 288](#)

Documentation Updates

IN THIS SECTION

- [Installation and Upgrade guide | 287](#)

This section lists the errata and changes in Junos OS Release 19.2R3 for the QFX Series switches documentation.

Installation and Upgrade guide

- **Veriexec explained (QFX Series)**—Verified Exec (also known as veriexec) is a file-signing and verification scheme that protects the Junos operating system (OS) against unauthorized software and activity that might compromise the integrity of your device. Originally developed for the NetBSD OS, veriexec was adapted for Junos OS and enabled by default from Junos OS Release 7.5 onward.

[See [Veriexec Overview](#).]

SEE ALSO

- [What's New | 243](#)

What's Changed	253
Known Limitations	259
Open Issues	261
Resolved Issues	265
Migration, Upgrade, and Downgrade Instructions	288

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 288
- Installing the Software on QFX10002-60C Switches | 291
- Installing the Software on QFX10002 Switches | 291
- 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 | 292
- Installing the Software on QFX10008 and QFX10016 Switches | 294
- Performing a Unified ISSU | 298
- Preparing the Switch for Software Installation | 299
- Upgrading the Software Using Unified ISSU | 299
- Upgrade and Downgrade Support Policy for Junos OS Releases | 301

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 **19.2** 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 19.2 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-x86-64-19.2-R3.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 19.2jinstall 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-x86-64-19.2R3.9.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-x86-64-19.2R3.9.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-x86-64-19.2R3.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-x86-64-19.2R3.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-x86-64-19.2R3.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-x86-64-19.2R3.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 299](#)
- [Upgrading the Software Using Unified ISSU on page 299](#)

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

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

To verify that nonstop active routing is enabled:

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

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

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

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

Upgrading the Software Using Unified ISSU

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

To upgrade the switch using unified ISSU:

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

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

where *package-name.tgz* is, for example, *jinstall-host-qfx-10-f-x86-64-18.3R1.n-secure-signed.tgz*.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

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

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R3.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 | 243](#)

[What's Changed | 253](#)

[Known Limitations | 259](#)

[Open Issues | 261](#)

[Resolved Issues | 265](#)

[Documentation Updates | 287](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 303](#)
- [What's Changed | 313](#)
- [Known Limitations | 317](#)
- [Open Issues | 319](#)
- [Resolved Issues | 321](#)
- [Documentation Updates | 340](#)
- [Migration, Upgrade, and Downgrade Instructions | 341](#)

These release notes accompany Junos OS Release 19.2R3 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

- [New and Changed Features: 19.2R3 | 303](#)
- [New and Changed Features: 19.2R2 | 303](#)
- [New and Changed Features: 19.2R1-S1 | 304](#)
- [New and Changed Features: 19.2R1 | 304](#)

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

New and Changed Features: 19.2R3

There are no new features in Junos OS Release 19.2R3 for the SRX Series devices.

New and Changed Features: 19.2R2

Intrusion Detection and Prevention (IDP)

- **HTTP X-Forwarded-For header support in IDP (SRX Series)**—Starting in Junos OS Release 19.2R2, we've introduced the **log-xff-header** option to record the x-forward-for header (xff-header) information. When this option is enabled. During the traffic flow, IDP saves the source IP addresses (IPv4 or IPv6) from the contexts for HTTP and SMTP traffics and displays in attack logs.

The xff-header is not processed unless it's enabled through sensor-configuration.

- To enable the xff-header, use the **set security idp sensor-configuration global log-xff-header** command.
- To disable the xff-header, use the **delete security idp sensor-configuration global log-xff-header** command.

In previous releases, when you accessed the internet, to lessen the external bandwidth the servers used transparent proxies. It was difficult to identify the originating source IP address because the proxy server converted it into an anonymous source IP address.

[See [Understanding Multiple IDP Detector Support](#).]

New and Changed Features: 19.2R1-S1

Routing Protocols

- **Decouple RSVP for IGP-TE (MX Series, PTX Series, ACX Series, QFX Series, SRX Series, and EX Series)**—Starting in Junos OS Release 19.2R1-S1, 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.

New and Changed Features: 19.2R1

Application Security

- **Application-based multipath support (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX4100, SRX4200, and vSRX)**—Starting in Junos OS Release 19.2R1, application-based multipath routing is supported on SRX Series devices.

Multipath routing allows the sending device to create copies of packets and to send each copy through two or more WAN links. On the other end, multipath calculates the jitter and packet loss for the combined links and estimates the jitter and packet loss for the same traffic on individual links. You can compare the reduction in packet loss when combined links instead of individual links are used. Sending multiple copies of traffic ensures timely delivery of the sensitive application traffic.

Multipath support in SD-WAN use cases enhances application experience.

[See [Application Quality of Experience](#).]

- **Application-level logging for AppQoE (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX4100, SRX4200, and vSRX)**—Starting in Junos OS Release 19.2R1, SRX Series devices support application-level logging for AppQoE. This feature reduces the impact on the CSO or log collector device while processing a large number of system log messages generated at the session-level. The SRX Series device maintains session-level information and provides system log messages for the session level. Replacing session-level logging with application-level logging decreases the overhead on the SRX Series device and increases AppQoE throughput.

[See [AppQoE](#).]

- **Secure Web proxy (SRX Series and vSRX)**—Starting in Junos OS Release 19.2R1, SRX Series devices support secure Web proxy service.

The secure Web proxy feature enables you to specify dynamic Web applications for which the system performs proxy service. In this deployment, the SRX Series device receives a request from the client,

examines the HTTP header for the application, and redirects the request directly to the webserver based on the application.

As a result, the SRX Series device performs transparent proxy between the client and the webserver for the specified applications and provides better quality of service for the application traffic.

[See [SSL Proxy](#).]

- **Application identification of micro-applications (SRX Series, vSRX)**—Starting in Junos OS Release 19.2R1, SRX Series devices support micro-applications with the application identification (AppID) feature.

AppID detects the applications at the subfunction level on your network and the security policy leverages the application identity information determined from the AppID module. After a particular application is identified, an action such as permit, deny, reject, or redirect is applied to the traffic according to the policy configured on the device.

[See [Application Identification](#).]

- **JDPI-Decoder engine version upgrade (SRX Series)**—Starting in Junos OS Release 19.2R1, the Juniper Networks Deep Packet Inspection-Decoder (JDPI-Decoder) engine comes with a default application signature package version 999 that includes the protobundle version 1.380.0-64.005 and the JDPI-Decoder engine version 5.3.0-56. You can also upgrade the application signature package when a new signature package version is available.

[See [show services application-identification status](#).]

Flow-Based and Packet-Based Processing

- **PowerMode IPsec fragment support (SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 19.2R1, PowerMode IPsec (PMI) is enhanced to handle the incoming and outgoing fragment packets in first path or fast path processing.

PMI supports first path and fast path processing both for fragment handling and for unified encryption. You can enable PowerMode IPsec processing by using the **set security flow power-mode-ipsec** command.

See [[Improving IPsec Performance with PowerMode IPsec](#).]

- **Multiple J-Flow Server (SRX Series)** —On SRX Series devices, the J-Flow version 9 can export flow records to only one collector. Starting from Junos OS Release 19.2R1, the J-Flow version 9 can configure up to 4 collectors under a family.

Packet Forwarding Engine exports flow record, flow record template, option data, and option data template packet to up to four collectors under a family. The template that is mapped, and the export version across the collectors under a family should be same.

- **Per-flow CoS support for GTP-U in PMI mode (SRX5000 line of devices with SPC3)**— Starting in Junos OS Release 19.2R1, Junos OS supports per-flow CoS functions for GTP-U traffic in PowerMode IPsec (PMI) mode. This feature introduces tunnel endpoint identifier (TEID)-based hash distribution for creating GTP-U sessions to multiple cores on the anchor PIC when both PMI and IPsec session affinity are enabled. TEID-based hash distribution helps split a fat GTP session into multiple slim GTP sessions and process

them on multiple cores in parallel. With this enhancement, per-flow CoS for GTP-U traffic is enabled even when the traffic carries multiple streams with different DSCP code within one GTP tunnel.

[See [PMI Flow Based CoS functions for GTP-U](#).]

Intrusion Detection and Prevention (IDP)

- **Support for IDP intelligent inspection (SRX Series and vSRX)**—Starting in Junos OS Release 19.2R1, you can enable IDP intelligent inspection and tune it dynamically to reduce IDP inspection load. IDP intelligent inspection helps the device to recover from overload state when the configured CPU and memory threshold values exceed the resource limits. Prior to Junos OS Release 19.2R1, when the device exceeds the configured CPU and memory threshold limit, IDP either rejects or ignores new sessions.

[See [IDP Intelligent Inspection](#).]

Juniper Sky ATP

- **Juniper Sky ATP Support for Encrypted Traffic Inspection and Server Name Identification**—Starting in Junos OS 19.2, SRX Series devices support inspection of encrypted traffic (HTTPS) in security-intelligence policies. Server name identification (SNI) checks are also supported. Note that these changes do not introduce any new CLI commands. All existing commands and configurations can make use of this expanded functionality.

Junos Telemetry Interface

- **Support for JTI (SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 19.2R1, you can stream statistics through junos telemetry interface (JTI) to an outside collector using remote procedure call (gRPC) services. gRPC is a protocol for configuration and retrieval of state information.

JTI supports the following sensors:

- Log messages (resource path `/junos/events`)
- Border Gateway Protocol (BGP) peer information (resource path `/network-instances/network-instance/protocols/protocol/bgp/`)
- Memory utilization for a routing protocol task (resource path `/junos/task-memory-information/`)
- Operational state of hardware components (resource path `/components/`)
- Operational state of the AE interface (resource path `/lACP/`)
- Operational state of Ethernet interfaces enabled with Link Layer Discovery Protocol (LLDP) (resource path `/lldp/`)
- Address Resolution Protocol (ARP) statistics (resource path `/arp-information/`)
- Routing Engine internal interfaces, such as `fxp0`, `em0`, and `em1` (resource path `/interfaces/interface[name=' interface-name ']/`)
- Network Discovery Protocol (NDP) table state (resource path `/nd6-information/`)

- NDP router advertisement statistics (resource path `/ipv6-ra/`)
- Intermediate System to Intermediate System (IS-IS) protocol statistics (resource path `/network-instances/network-instance/protocols/protocol/isis/levels/level/`)
- IS-IS protocol (resource path `/network-instances/network-instance/protocols/protocol/interfaces/interface isis/levels/level/`)

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

J-Web

- **Threats Map (Live) (SRX Series except SRX5000 line of devices and vSRX)**—Starting in Junos OS Release 19.2R1, you can monitor Threat Maps (Live). You can view blocked and allowed threat events based on feeds from intrusion prevention systems (IPSs), antivirus, antispam engines, Juniper Sky ATP, and screen options. You can also choose a country and view the total threat events for that country since midnight, followed by the number of inbound and outbound threat events, and see the top five IP addresses, either inbound or outbound. With View Details, you can see the additional details of the selected country.

[See [Monitor Threats Map \(Live\)](#).]

- **Quick Setup wizard enhancement (SRX Series except SRX5000 line of devices)**—Starting in Junos OS Release 19.2R1 after the configuration is completed, you see a notification message when you access J-Web again through a new browser tab or window with the configured IPv4 or IPv6 address.

[See [Understanding the J-Web CLI Terminal](#).]

- **Getting Started panel (SRX Series)**—Starting in Junos OS Release 19.2R1, you have quick access to the important configurations using a Getting Started panel on the J-Web UI. For logical systems users and tenant users, this option is available only in SRX1500, SRX4100, SRX4200, SRX4600, and SRX5000 line of devices.

By default, this panel appears when you log in. If you choose **Don't show this again**, then you can access this panel using the help (?) icon.

[See [Security J-Web Getting Started](#).]

- **HA Mode wizard (SRX Series)**—Starting in Junos Release 19.2R1, you can configure chassis cluster using a new HA Mode wizard when the devices are in factory default. You can create HA using the same wizard from Configure > Device Settings > Cluster (HA) Setup when the devices are already in the network.

[See [Configuring Cluster \(HA\) Setup](#).]

- **IPS Sensor enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, you can configure IP Sensor using the following settings:
 - Basic—Supports protection mode, IDP intelligent inspection, and basic IDP flow configuration.
 - Advanced—Supports IDP flow, global, IPS, log, reassembler, and packet log configuration.
 - Detectors settings—Supports the configuration for a specific service. You can also add or edit the configuration inline.

[See [Sensor Configuration Page Options](#).]

- **Active Directory enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, for SRX4200, SRX1500, SRX550M, and vSRX, and for the SRX5000 and SRX300 lines of devices, you can configure the integrated user firewall in a maximum of two domains. For the other SRX Series devices, you can create only one domain.

[See [Configuring Active Directory](#).]

- **Certificate management enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, you can now configure device certificates, trusted certificate authorities (CAs), and CA groups. You can view information about the local certificate, trusted CA profiles, and CA groups that are configured on the device. You can manually generate self-signed certificate. You can enroll online, export, import, manually load, and delete the local certificate or certificate signing request (CSR).

[See [Managing Device Certificates](#).]

- **Forwarding mode enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, flow mode is the default mode for processing traffic. You can now configure an SRX Series devices as a border router by changing the flow-based processing to packet-based processing.

[See [Forwarding Configuration Page Options](#).]

- **Dashboard enhancement (SRX Series except SRX5000 line of devices)**—Starting in Junos OS 19.2R1, you can view the Web filtering, Antispam, Content filtering, Application & Users, and Threat monitoring widgets in the J-Web dashboard for root, logical systems, and tenant users.

[See [Monitoring the Dashboard](#).]

- **Security policy rules enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, when you create rules for the destination traffic, you can:
 - Add an application or application group for a dynamic application using the Add New Application/Group button.
 - Add a service for Service(s) using the Add New Service button.

[See [Configuring Firewall Security Policy Rules](#).]

- **Monitoring firewall events enhancement (SRX Series except SRX5000 line of devices)**—Starting in Junos OS Release 19.2R1, you can now see that an application displays the same value as a nested application (if the application supports nested applications).

[See [Monitoring Firewall Events](#).]

- **Monitoring events enhancement (SRX Series except SRX5000 line of devices)**—Starting in Junos OS Release 19.2R1, you can monitor the following new events:
 - ATP—Top Malware Source Countries, Infected File Categories, and Malwares Identified widgets are shown in the chart view and detailed advanced anti-malware (AAMW) logs are shown in the grid view.
 - Security Intelligence—Top Infected Hosts and C&C Servers widgets are shown in the chart view and detailed secintel logs are shown in the grid view.

- Screens—Top Screen Attacks, Screen Victims, and Screen Hits widgets are shown in the chart view and detailed screen logs are shown in the grid view.

[See [Monitoring ATP Events](#), [Monitoring Security Intelligence Events](#), and [Monitoring Screen Events](#).]

- **Juniper Sky ATP enrollment enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, you can view the detailed enrollment steps on the SKY ATP Enrollment page.

[See [Sky ATP Enrollment](#).]

- **Link aggregation enhancement (Standalone SRX Series)**—Starting in Junos OS Release 19.2R1, VLAN tagging is enabled by default when you add an AE interface.

[See [Link Aggregation Configuration Page Options](#).]

- **OSPF enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, you can configure OSPF area in two ways:

- Basic—You can add new routing instances.
- Advanced—You can group a policy and trace options.

[See [OSPF Configuration Page Options](#).]

- **VLAN enhancement (SRX Series)**—Starting in Junos OS Release 19.2R1, Bridge domain is the new name for VLANs in Layer 2 transparent mode. You can assign an interface for the created VLANs. You can view all the available VLANs with their IDs, interfaces assigned, and status.

[See [VLAN Configuration Page Options](#).]

Logical Systems and Tenant Systems

- Starting in Junos OS Release 19.2R1, the following features that are supported on the logical systems are now extended to tenant systems:

- **Default routing-instance support for tenant systems (SRX Series)**—Starting in Junos OS Release 19.2R1, you can use the **ping**, **telnet**, **ssh**, **traceroute**, **show arp**, **clear arp**, **show ipv6 neighbors**, and **clear ipv6 neighbors** commands to pass the virtual router configured in a tenant system as a default routing instance.

[See [Tenant Systems Overview](#).]

- **UTM support for tenant systems (SRX Series)**—Starting in Junos OS Release 19.2R1, SRX Series devices support unified threat management (UTM) on tenant systems. Use the **set utm default-configuration** command under the **[edit security]** hierarchy level to create a default UTM profile for tenant systems. Configure policies, profiles, and custom objects for each tenant system in the UTM profile.

[See [UTM for Tenant Systems](#).]

- **On-box logging support for tenant systems (SRX Series)**—Starting in Junos OS Release 19.2R1, SRX Series devices support on-box logging configurations for each tenant system, and handle logs based on these configurations. Configure the **set log mode event** and **set log mode stream** commands under

the **[edit security]** hierarchy level to enable on-box logging. Tenant systems also support binary format log in event mode.

[See [Security Log for Tenant Systems](#).]

- **IDP for tenant systems (SRX Series and vSRX)**—Starting in Junos OS Release 19.2R1, tenant systems support intrusion detection and prevention (IDP). The IDP policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a tenant system.

[See [IDP for Tenant Systems](#).]

Network Management and Monitoring

- **Support for displaying valid user input in the CLI for command options and configuration statements in custom YANG data models (SRX Series)**—Starting in Junos OS Release 19.2R1, the CLI displays the set of possible values for a given command option or configuration statement in a custom YANG data model when you include the **action-expand** extension statement in the option or statement definition and reference a script that handles the logic. The **action-expand** statement must include the **script** child statement, which defines the Python action script that is invoked when a user requests context-sensitive help in the CLI for the value of that option or statement.

[See [Displaying Valid Command Option and Configuration Statement Values in the CLI for Custom YANG Modules](#).]

Security

- **Support to configure micro-applications in a unified policy (SRX Series and vSRX)**—Starting in Junos OS Release 19.2R1, you can configure micro-applications in a unified policy. Micro-applications are subfunctions of a particular application.

You can configure micro-applications at the same hierarchy as predefined dynamic applications in a security policy and take the action based on the policy rules.

[See [Configuring Micro-Applications in Unified Policies](#).]

Unified Threat Management (UTM)

- **SRX5K-SPC3 support Avira scan engine on antivirus module (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 19.2R1, SRX Series devices support an on-device antivirus Avira scan engine. The on-device antivirus Avira scan engine scans the data by accessing the virus pattern database. The antivirus scan engine is provided as a unified threat management (UTM) module that you can download and install on an SRX Series device either manually or by using the Internet to connect to a Juniper Networks-hosted URL or a user-hosted URL.

NOTE: The SRX5000 line of devices with SRX5K-SPC-4-15-320 or SRX5K-SPC-2-10-40 cards do not support the on-device antivirus Avira scan engine.

[See [On-Device Antivirus Scan Engine](#).]

VPN

- **PIM using point-to-multipoint mode support for AutoVPN and Auto Discovery VPN (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, and vSRX)**—Starting in Junos OS Release 19.2R1, Protocol Independent Multicast (PIM) using point-to-multipoint (P2MP) mode supports AutoVPN and Auto Discovery VPN in which a new **p2mp** interface type is introduced for PIM. The **p2mp** interface tracks all PIM joins per neighbor to ensure that multicast forwarding or replication happens only to those neighbors that are in joined state. In addition, the PIM using point-to-multipoint mode supports chassis cluster mode.

[See [Multicast Overview](#), [Understanding AutoVPN](#), [Understanding Auto Discovery VPN](#), and [Understanding Multicast Routing on a Chassis Cluster](#).]

- **PowerMode IPsec for NAT-T (SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 19.2R1, SRX Series devices equipped with SRX5K-SPC3 Services Processing Cards (SPCs) support PowerMode IPsec (PMI) for Network Address Translation-Traversal (NAT-T).

[See [Understanding PowerMode IPsec](#).]

- **IPsec Distribution Profile (SRX5400, SRX5600, and SRX5800)**—Starting with Junos OS Release 19.2R1, you can manage the tunnel distribution through the configuration. You can create a profile for a VPN object to handle the distribution of tunnels. In a profile, mention the slot and thread-id where the tunnels from the VPN object should be distributed. The same profiles can be used for different VPN objects.

To add profiles for distributing IPsec SAs, use the new **distribution-profile profile-name** statement.

[See [IPsec Distribution Profile](#) and [distribution-profile](#).]

- **Anti-replay window (SRX Series 5000 line of devices with SPC3 cards)**—Starting from Junos OS Release 19.2R1, you can configure the anti-replay window size within the range of 64 to 8192 (power of 2). If you do not configure the anti-replay window size, the default window size remains as 64.

To configure the window size, use the new **anti-replay-window-size** option.

[See [Replay Protection](#).]

SEE ALSO

[What's Changed | 313](#)

[Known Limitations | 317](#)

[Open Issues | 319](#)

[Resolved Issues | 321](#)

[Documentation Updates | 340](#)

[Migration, Upgrade, and Downgrade Instructions | 341](#)

What's Changed

IN THIS SECTION

- [Release 19.2R3 Changes in Behavior and Syntax | 314](#)
- [Release 19.2R2 Changes in Behavior and Syntax | 314](#)
- [Release 19.2R1 Changes in Behavior and Syntax | 316](#)

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

Release 19.2R3 Changes in Behavior and Syntax

General Routing

- Repetition of WALinuxAgent logs causing file size increase (vSRX 3.0)—The Azure WALinuxAgent performs the provisioning job for the vSRX instances. When a new vSRX instance is deployed, the continued increasing size of the waagent log file might cause the vSRX to stop. If the vSRX is still operating, then delete the /var/log/waagent.log directly or run the clear log waagent.log all command to clear the log file. Or you can run the set groups azure-provision system syslog file waagent.log archive size 1m and set groups azure-provision system syslog file waagent.log archive files 10 commands to prevent the growing of the waagent logs. These configurations will cause the rotation of log of waagent with the size bigger than 1MB and set a maximum of 10 backups. See [vSRX with Microsoft Azure](#).

Juniper Extension Toolkit (JET)

- Set the trace log to only show error messages (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series)— You can set the verbosity of the trace log to only show error messages using the error option at the **edit system services extension-service traceoptions level** hierarchy.

See [traceoptions \(Services\)](#).

Network Management and Monitoring

- Packet Forwarding Engine data display command (SRX Series)— Starting in Junos OS Release 19.2R3, you can now view the Packet Forwarding Engine data using the **usp flow config** and **usp flow stats** options for the **show pfe data** command.

Juniper Sky ATP

- Dynamic address entries on SRX Series devices in chassis cluster mode—Starting in Junos OS Release 19.2R3, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Sky Advanced Threat Prevention (ATP).

Release 19.2R2 Changes in Behavior and Syntax

Authentication and Access Control

- Enhanced user firewall support—In Junos OS Release 19.2R2, for SRX300 devices with eUSB (SRX300, SRX320, SRX340, and SRX345), the SRX Series user firewall (UserFW) module tries to synchronize user entries from the domain controller or Juniper Identity Management Service (JIMS) after booting up. If the historical login events expired on the domain controller, then the SRX Series UserFW module is unable to retrieve those user entries after the UserFW module boots up.

[See [User Authentication Entries in the ClearPass Authentication Table](#).]

- SSH protocol version v1 option deprecated from CLI (SRX Series)—Starting in Junos OS Release 19.2R2, we've removed the nonsecure SSH protocol version 1 (v1) option from the **[edit system services ssh protocol-version]** hierarchy level. You can use the SSH protocol version 2 (v2) as the default option to

remotely manage systems and applications. With the **v1** option deprecated, Junos OS is compatible with OpenSSH 7.4 and later versions.

Junos OS releases earlier than Release 19.2R2, continue to support the **v1** option to remotely manage systems and applications.

[See [protocol-version](#).]

- **Enabling and disabling SSH login password or challenge-response authentication (SRX Series)**—Starting in Junos OS Release 19.2R2, you can disable either the SSH login password or challenge-response authentication at the **[set system services ssh]** hierarchy level.

In Junos OS releases earlier than 19.2R2, you can enable or disable both SSH login password and challenge-response authentication simultaneously at the **[set system services ssh]** hierarchy level.

[See [Configuring SSH Service for Remote Access to the Router or Switch](#).]

Security

- **Disable the do not fragment flag from packet IP header (SRX Series and vSRX)**—Starting in Junos OS Release 19.2R2, we've introduced the **clear-dont-frag-bit** option at the **[edit security alg alg-manager]** hierarchy level to disable the do not fragment flag from the packet IP header, which allows the packet to be split after NAT is performed.

In Junos OS releases earlier than Release 19.2R2, when the ALG performs payload-NAT, sometimes the size of the packet becomes bigger than the outgoing interface maximum transmission unit (MTU). If the packet IP header has the do not fragment flag, this packet cannot be sent out.

[See [alg-manager](#).]

Release 19.2R1 Changes in Behavior and Syntax

Application Security

- Starting in Junos OS Release 19.2R1, the SSL decryption mirroring feature is supported on redundant Ethernet (reth) interface on SRX Series devices operating in a chassis cluster.

Ethernet Switching and Bridging

- **Support for double tagged VLANs being pushed out the egress interface (SRX300, SRX320, SRX340, SRX345, SRX550, and SRX1500)**—Starting in Junos OS Release 19.2R1, in a Q-in-Q scenario, double tagged VLANs are pushed out the egress interface. In previous releases, when two VLANs were added at the ingress interface, with the **native-vlan-id** *vlan-id* assigned to the user-to-network interface (UNI) interface and the **vlan-id** *vlan-id-list* assigned to the network-to-network interface (NNI) interface, the VLAN with the **native-vlan-id** tag did not exit from the egress interface. Now both VLAN tags exit from the egress interface.

Flow-Based and Packet-Based Processing

- **Power Mode IPsec (SRX Series)**—On SRX Series devices, when Power Mode IPsec is enabled, the **show security flow statistics** and **show security flow session tunnel summary** commands does not count, or display the number of packets that are processed within the Power Mode IPsec as these packets do not go through the regular flow path.

[See [show security flow statistics](#).]

Network Management and Monitoring

- **The show system schema command and <get-yang-schema> RPC require specifying an output directory (SRX Series)**—Starting in Junos OS Release 19.2R1, when you issue the **show system schema** operational mode command in the CLI or execute the **<get-yang-schema>** RPC in a remote session to retrieve schema files, you must specify the directory in which to generate the output files by including the **output-directory** command option in the CLI or the **<output-directory>** element in the RPC. In earlier releases, you can omit the **output-directory** argument when requesting a single module to display the module in standard output.

- **Custom YANG RPC support for input parameters of type empty (SRX Series)**—Starting in Junos OS Release 19.2R1, custom YANG RPCs support input parameters of type **empty** when executing the RPC's command in the Junos OS CLI, and the value passed to the action script is the parameter name. In earlier releases, input parameters of type **empty** are only supported when executing the RPC in a NETCONF or Junos XML protocol session, and the value passed to the action script is the string '**none**'.

[See [Creating Action Scripts for YANG RPCs on Devices Running Junos OS](#).]

- **NSD Restart Failure Alarm (SRX Series)**—Starting in Junos OS Release 19.2R1, a system alarm is triggered when the Network Security Process (NSD) is unable to restart due to the failure of one or more NSD subcomponents. The alarm logs about the NSD are saved in the messages log. The alarm is automatically cleared when NSD restarts successfully.

The **show chassis alarms** and **show system alarms** commands are updated to display the following output when NSD is unable to restart - **NSD fails to restart because subcomponents fail**.

[See [Alarm Overview](#).]

VPNs

- **Certificate revocation list (SRX Series)**—Local certificates are being validated against certificate revocation list (CRL) even when CRL check is disabled. Starting in Junos OS Release 19.2R1, this can be stopped by disabling the CRL check through the Public Key Infrastructure (PKI) configuration. When CRL check is disabled, PKI will not validate local certificate against CRL.

[See [revocation-check \(Security PKI\)](#) and [Understanding Online Certificate Status Protocol and Certificate Revocation Lists](#).]

SEE ALSO

[What's New | 303](#)

[Known Limitations | 317](#)

[Open Issues | 319](#)

[Resolved Issues | 321](#)

[Documentation Updates | 340](#)

[Migration, Upgrade, and Downgrade Instructions | 341](#)

Known Limitations

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

J-Web

- CLI terminal is not working in Java version 1.8 because of a security restriction in the running applet. [PR1341956](#)
- After you generate the Default Trusted CA profile group under Certificate Management>Trusted Certificate Authority in J-Web, J-Web does not display the CA profile group local under Certificate Management>Certificate Authority Group page. [PR1424131](#)
- The CA profile group imported using J-Web does not populate the group on 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](#)
- Country logo is not displaying in Threats Map page and Events page for some countries. Time slider is not displayed properly in Screen/ATP/Security Intelligence events pages. [PR1435124](#)

VPNs

- In the HA design for SRX Series devices, the antireplay window is synchronized to the backup only when the total incoming packet count is an odd multiple of 128 packets. When a failover occurs, the antireplay bitmap is not synchronized. Again, when the node comes back online, the SA is installed but the antireplay bitmap is reset to 0 along with the in and out sequence number. [PR1420521](#)
- In a chassis cluster, ESP or AH packet sequence number is not synchronized to the backup node after the backup node is rebooted. [PR1433424](#)
- On the SRX5000 line of devices with SPC3 installed, the IPsec VPN antireplay sequence number might be reset to zero after the crash of the SPC3 card or the flowd process. Traffic drop is seen due to the mismatch of the sequence number. [PR1433568](#)
- Per tunnel debugging configuration is not synchronized to backup node. It needs to be configured again after RGO failover. [PR1450393](#)

SEE ALSO

[What's New | 303](#)

[What's Changed | 313](#)

[Open Issues | 319](#)

[Resolved Issues | 321](#)

[Documentation Updates | 340](#)

[Migration, Upgrade, and Downgrade Instructions | 341](#)

Open Issues

IN THIS SECTION

- [Chassis Clustering](#) | 319
- [Flow-Based and Packet-Based Processing](#) | 319
- [Intrusion Detection and Prevention \(IDP\)](#) | 320
- [J-Web](#) | 320
- [Platform and Infrastructure](#) | 320
- [Routing Policy and Firewall Filters](#) | 320
- [VPNs](#) | 320

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

Chassis Clustering

- When GTP profile with the same name is deleted and then added, the profile ID will be changed. So, if this profile is being used by policy, you need to reconfigure the policy application bounding; otherwise, the GTP will not work as you expect. [PR1409213](#)

Flow-Based and Packet-Based Processing

- On an SRX4600 device, when the next hop is set to the st0 interface, the output of the show route forwarding-table command displays the next-hop IP address twice. [PR1290725](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, it is possible that when multiple core files are generated in quick succession, the cold-sync-monitored status is displayed and cannot be removed even though cold-sync has finished. You must reboot the affected node to recover. [PR1403000](#)
- Automatic completion is not working on proxy terminator profile name. [PR1424822](#)
- Syslog severity level of msg subtype is end of policy is set to error although this message can be ignored. [PR1435233](#)
- Some packets are dropped due to the FPGA issue. [PR1443600](#)
- SSL-based AppID simplification effort (removal of HTTPS, POP3S, IMAPS, SMTPS). [PR1444767](#)
- TCP session cannot time out properly upon receiving the TCP RESET packet, and the session timeout does not change to two seconds. [PR1467654](#)

Intrusion Detection and Prevention (IDP)

- On SRX Series devices, commit or show command for IDP might not work if SNMP queries are run when large-scale IDP is used. [PR1444043](#)

J-Web

- Due to set chassis auto-image-upgrade in factory configuration, from phone home page you are not able to skip to J-Web and get the error Bootstrap is in progress, Can't Skip!!.. [PR1420888](#)
- SECINTEL_ACTION_LOG events with subcategories such as Infected-Hosts and C are not shown on Juniper Sky ATP threat count on Monitor>Threats Map page in J-Web. [PR1425795](#)
- On SRX Series devices, until Junos OS Release 19.2, phone home UI portal is displayed by default. The J-Web UI should be the default page to be launched when the device is in factory default. [PR1428717](#)
- Country logo is not displaying in Threats Map page and Events page for some countries. Time slider is not displayed properly in Screen/ATP/Security Intelligence events pages. [PR1435124](#)

Platform and Infrastructure

- Under certain heavy traffic conditions srpxfe process might crash and result in a denial-of-service (DoS) condition for the SRX1500 device. Repeated crashes of the srpxfe can result in an extended DOS condition. The SRX Series device might fail to forward traffic when this condition occurs. [PR1277363](#)
- CDN-based dynamic application classification has been deprecated in this release. To restore previous behavior, you can configure **set services application-identification enable-cdn-application-detection**. [PR1375442](#)

Routing Policy and Firewall Filters

- If a huge number of policies are configured on SRX Series devices and some policies are changed, the traffic that matches the changed policies might be dropped. [PR1454907](#)

VPNs

- On SRX Series devices, if multiple traffic selectors are configured for a peer with IKEv2 reauthentication, only one traffic selector is rekeyed at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors are cleared without immediate rekeying. New negotiation of these traffic selectors is triggered through other mechanisms such as traffic or by peer. [PR1287168](#)
- VPN tunnels flap after a group is added or deleted in edit private mode in a clustered setup. [PR1390831](#)

- On SRX5400, SRX5600, and SRX5800 devices, during in-service software upgrade (ISSU), the IPsec tunnels flap, causing a disruption of traffic. The IPsec tunnels recover automatically after the ISSU process is completed. [PR1416334](#)
- On the SRX5000 line of devices with SPC3 cards, sometimes IKE SA is not seen on the device when st0 binding on VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)
- In IPsec VPN scenario on SRX5400, SRX5600, and SRX5800 platforms, the IKED treats retransmission of IKE_INIT request packets as new connections when SRX Series device acts as responder of IKE negotiation. This causes IKE tunnel negotiation to fail and IPsec VPN traffic might be impacted. [PR1460907](#)
- The SRX5000 line of devices with SPC3 was not supporting simultaneous IKE negotiation in Junos OS Release 19.2, Junos OS Release 19.3, Junos OS Release 19.4, and Junos OS Release 20.1. [PR1497297](#)

SEE ALSO

[What's New | 303](#)

[What's Changed | 313](#)

[Known Limitations | 317](#)

[Resolved Issues | 321](#)

[Documentation Updates | 340](#)

[Migration, Upgrade, and Downgrade Instructions | 341](#)

Resolved Issues

Learn which issues were resolved in Junos OS main and maintenance releases for SRX Series devices. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Resolved Issues: 19.2R3

Chassis Clustering

- If a cluster ID of 16 or multiples of 16 is used, the chassis cluster might not come up. [PR1487951](#)

Flow-Based and Packet-Based Processing

- The show security group-vpn server statistics |display xml is not in expected format. [PR1349959](#)
- ECMP load balancing does not happen when RG1 node 0 is secondary. [PR1475853](#)

- The flowd or srxpfe process might stop when deleting user firewall local authentication table entry. [PR1477627](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- On Web proxy, memory leak occurs in association hash table and DNS hash table. [PR1480760](#)
- GRE or IPsec tunnel might not come up when set security flow no-local-favor-ecmp command is configured. [PR1489276](#)
- Outbound SSH connection flap or memory leak issue might be observed during pushing configuration to ephemeral database with high rate. [PR1497575](#)
- Traffic interruption happens due to MAC address duplication between two Junos OS devices. [PR1497956](#)
- Don't use capital characters for source-identity when using show security match-policies command. [PR1499090](#)
- J-Flow v9 does not display correct outgoing interface for APBR traffic. [PR1502432](#)

Interfaces and Chassis

- All interfaces remain in the down status after the SRX300 line of devices power up or reboot. [PR1488348](#)
- Continuous drops seen in control traffic, when high data queues in one SPC2 PIC. [PR1490216](#)

Intrusion Detection and Prevention (IDP)

- IDP's custom-attack time-binding interval command was mistakenly hidden within the CLI. [PR1506765](#)

J-Web

- The J-Web users might not be able to configure PPPoE using PPPoE wizard. [PR1502657](#)

Multiprotocol Label Switching (MPLS)

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

Platform and Infrastructure

- Packets get dropped when next hop is IRB over LT interface. [PR1494594](#)

Routing Policy and Firewall Filters

- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. [PR1492436](#)
- Traffic fails to hit the policies with match source-end-user-profile. [PR1505002](#)

Routing Protocols

- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

VPNs

- With NCP remote access solution, in a PathFinder case (for example, where IPsec traffic has to be encapsulated as TCP packets), TCP encapsulation for transit traffic is failing. [PR1442145](#)
- On SRX Series devices with SPC3, when overlapping traffic-selectors are configured, multiple IPsec SAs get negotiated with peer device. [PR1482446](#)

Resolved Issues: 19.2R2

Application Layer Gateways (ALGs)

- Sometimes unexpected forwarding sessions appear for tenant ALG SIP traffic in cross tenant. [PR1409748](#)
- The H.323 connection might not be established when the H.323 packet passes SRX Series devices twice through different virtual routers. [PR1436449](#)
- On the SRX5000 line of devices, the H323 call with NAT64 could not be established. [PR1462984](#)
- The flowd or srpxfe process might stop when an ALG creates a gate with an incorrect protocol value. [PR1474942](#)
- SIP messages that need to be fragmented might be dropped by SIP ALG. [PR1475031](#)

Application Security

- The AAMW diagnostic script generates incorrect error: **Error: Platform does not support SkyATP: srx300.** [PR1423378](#)
- Automatic application-identification download stops after going over the year and reboot. [PR1436265](#)
- The flowd or srpxfe process might crash when advanced anti-malware service is used. [PR1437270](#)
- The applications that get declassified in the middle of a session are not identified properly. [PR1437816](#)
- Unable to get more than 60 Gbps of AppQoS throughput. [PR1439575](#)
- The flowd process core files might be seen when the traffic hits an AppQoS policy. [PR1446080](#)

Authentication and Access Control

- The CPU utilization of the uacd is high, about 100 percent, in the output of **show chassis routing-engine**. [PR1424971](#)
- The srpxfe or flowd process might stop if a UAC policy is removed. [PR1437892](#)
- Same-source IP sessions are cleared when the IP entry is removed from the UAC table. [PR1457570](#)

Chassis Clustering

- IP monitoring might fail on the secondary node. [PR1468441](#)
- The chassis cluster failover to a secondary node does not happen after Packet Forwarding Engine stops on the primary node [PR1451091](#)
- Hardware failure is seen on both nodes in **show chassis cluster status**. [PR1452137](#)
- On SRX Series devices with chassis cluster, the control link remains up even though the control link is actually down. [PR1452488](#)
- An unhealthy node might become primary in SRX4600 devices with chassis cluster scenario. [PR1474233](#)

Class of Service (CoS)

- Frequent issuance of the **show class-of-service spu statistics** command causes the rtlogd process to be busy. [PR1438747](#)
- The CoS rewrite rule does not work for st0 interface. [PR1439401](#)

Flow-Based and Packet-Based Processing

- Throughput or latency performance of all traffic drops when TCP traffic is passing through the device. [PR1403727](#)
- Juniper Sky ATP does not escape the \ inside the username before the metadata is sent to the cloud. [PR1416093](#)
- Blacklist compilation failures are reported. [PR1418980](#)
- Group VPN IKE security associations cannot be established before RGO failover. [PR1419341](#)
- SSL proxy did not correctly warn users about unsupported certificates. [PR1419485](#)
- The trusted-ca and root-ca names or IDs should not be the same within an SSL proxy configuration. [PR1420859](#)
- Failed to clear sessions on SPC2 with error message error: **usp_ipc_client_recv::ipc_pipe_read() failed read timed out after 5 second(s)**. [PR1426090](#)
- When configuring a GRE tunnel (GRE-over-IPsec-tunnel) or an IPsec tunnel on an SRX Series device, the MTU of the tunnel interface is calculated incorrectly. [PR1426607](#)
- Junos OS: SRX1500: Denial of service due to crash of srxpfe process under heavy traffic conditions. (CVE-2019-0050) [PR1428657](#)
- The X2 traffic cannot be encrypted after the traffic is decrypted when PMI is enabled. [PR1429473](#)
- The flowd process might stop on the SRX5000 line of devices. [PR1430804](#)
- VPN traffic fails after primary node reboot or power off. [PR1433336](#)
- Intermittent packet drop might be observed if IPsec is configured. [PR1434757](#)
- Traffic drop when session key rolls over between primary and fallback for more than 10 times. [PR1435277](#)

- The second IPsec ESP tunnel might not be able to establish between two IPv6 IKE peers. [PR1435687](#)
- Control logical interface is not created by default for LLDP. [PR1436327](#)
- On an SRX4600 device, core file generation might be observed and SPM might be in present state. [PR1436421](#)
- The ipfd process might crash when SecIntel is used. [PR1436455](#)
- Some webpages cannot be fully rendered. [PR1436813](#)
- When running SSL proxy on the firewall, the locally generated certificate is not validated by OpenSSL client. [PR1436831](#)
- Packet reorder does not work when sending traffic over IPsec tunnel with session-affinity. [PR1436720](#)
- Decryption traffic doesn't take PMI path after IPsec rekey (initiated by peer) when loopback interface is configured as external interface. [PR1438847](#)
- The flowd process stops and generates a core file when processing SSL proxy traffic. [PR1437783](#)
- Member of dynamically created VLANs information is not displaying on show VLANs. [PR1438153](#)
- The probe of Ethernet switching always shows down in a chassis cluster scenario. [PR1438277](#)
- The flowd process stops and generates core files. [PR1438445](#)
- Security logs cannot be sent to the external syslog server through TCP. [PR1438834](#)
- When lmd is rotating database, there is a possibility of reading a NULL db at the same time, which generates core files. [PR1439186](#)
- LACP MUX state stuck in "Attached" after disabling peer active members when link protection is enabled on local along with force-up. [PR1439268](#)
- The IKE pass-through packet might be dropped after the source address has been changed. [PR1440605](#)
- While checking the flow session XML for source NAT under tenant, there is no value identifier for tenant-name (< tenant>< /tenant>). [PR1440652](#)
- Performance improvements were made to Screens, which benefit multiset socket systems. [PR1440677](#)
- Support inspection for pass-through IP-IP tunnel traffic on TAP mode. [PR1441226](#)
- SPC2 wrongly forwarded packet to SPC3 core0 and core14. [PR1441234](#)
- New CLI option to show only useful group information for an Active Directory user. [PR1442567](#)
- The SRX300 line of devices does not have MIB that can retrieve the fan status. [PR1443649](#)
- In the BERT test for E1 interface, bits counts number is not within the range. [PR1445041](#)
- Junos OS: SRX5000 Series: flowd process crash due to receipt of specific TCP packet (CVE-2019-0064) [PR1445480](#)
- There is no active connection with Juniper Sky ATP server. [PR1446481](#)

- The flowd process might stop on SRX Series devices when chassis cluster and IRB interface are configured. [PR1446833](#)
- 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](#)
- 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](#)
- On SRX Series devices with SSL proxy service used, a memory leak issue might occur, which results in the flowd or srpxfe process stopping. [PR1450829](#)
- AAWM policy rules for IMAP traffic sometimes might not get applied when passed through an SRX Series device. [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 is not observed properly on R0. [PR1451474](#)
- The rpd process might stop and restart and an rpd core file is generated when committing the configuration. [PR1451860](#)
- Update SRX300 traffic default logging to stream mode. [PR1453074](#)
- The fxp0 interface might redirect packets not destined to itself. [PR1453154](#)
- Introduction of default inspection limits for application identification to optimize CPU usage and improve resistance to evasive applications. [PR1454180](#)
- The SRX Series devices stop and generate several core files. [PR1455169](#)
- When you try to reset the system configuration on an SRX1500 device using the reset config button, it does not work properly. [PR1458323](#)
- The security flow traceoptions fills in with RTSP ALG related information. [PR1458578](#)
- Optimizations were made to improve the connections-per-second performance of SPC3. [PR1458727](#)
- SRX Series device might not be reachable when initiating offline command for PIC. [PR1459037](#)
- 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](#)
- The srpxfe or flowd process might stop if the sampling configuration is changed. [PR1462610](#)
- 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](#)

- A core file might be generated when you perform an ISSU on SRX Series devices. [PR1463159](#)
- Fragmented traffic might get looped between the fab interface in a rare case. [PR1465100](#)
- The PKI daemon keeps leaking memory on SRX Series devices. [PR1465614](#)
- HTTP block message stops working after SNI check for HTTPS session. [PR1465626](#)
- Loading CA certificate causes PKI daemon core file to be generated. [PR1465966](#)
- 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](#)
- Packet Forwarding Engine might generate core files because SSL proxy is enabled on NFX Series and SRX Series devices. [PR1467856](#)
- Server unreachable is detected; ensure that port 443 is reachable. [PR1468114](#)
- Tail drop on all ports is observed when any switch-side egress port gets congested. [PR1468430](#)
- FTP data connection might be dropped if SRX Series devices send the FTP connection traffic through the dl interface. [PR1468570](#)
- RPM test probe fails to show that round-trip time has been exceeded. [PR1471606](#)
- On SRX Series devices, Packet Forwarding Engine memory might be used up if the security intelligence feature is configured. [PR1472926](#)
- Support LLDP protocol on reth interface. [PR1473456](#)
- Certificate error while configuration validation during Junos OS upgrade. [PR1474225](#)
- Packet drop might be observed on the SRX300 line of devices when adding or removing an interface from MACsec. [PR1474674](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- Recent changes to JDPI's classification mechanism caused a considerable performance regression (more than 30 percent). [PR1479684](#)
- The flowd or srpxfe process might crash when advanced anti-malware services are used. [PR1480005](#)
- IMAP curl sessions stuck in the active state if AAMW IMAP block mode is configured. [PR1484692](#)
- The **show chassis temperature-thresholds** command displays a lot of FPC 0 output. [PR1485224](#)
- After installation through boot loader at cluster setup, primary node cannot proceed commit. [PR1487831](#)
- If a cluster-id of 16 or multiples of 16 is used, the cluster might not come up. [PR1487951](#)
- On SRX1500, CPU board inlet increases after Junos OS upgrade from Junos OS Release 15.1X49 to Junos OS Release 18.x. [PR1488203](#)
- Risk of service interruption on SRX Series devices with a dual-stacked CA server. [PR1489249](#)

Installation and Upgrade

- SPMC version mismatch errors after Junos OS install using USB method. [PR1437065](#)
- Junos OS upgrade fails when partition option is used on SRX Series devices. [PR1449728](#)

Interfaces and Chassis

- Both nodes in the SRX Series chassis cluster go into DB mode after downgrading to Junos OS Release 18.1. [PR1407295](#)
- 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](#)
- LFM remote loopback is not working as expected. [PR1428780](#)
- The LACP interface might flap if performing a failover. [PR1429712](#)
- Certain interfaces might drop all unicast traffic when LTE PIM is used. [PR1430403](#)
- Static route through dl0.0 interface is not active. [PR1465199](#)
- MAC limiting on Layer 3 routing interfaces does not work. [PR1465366](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srpxfe process stops and generates core files when processing IDP packets. [PR1416275](#)
- NSD fails to push security zone to the Packet Forwarding Engine after reboot, if there is an active IDP rule configured with FQDN. [PR1420787](#)
- The flowd or srpxfe process stops and generates core files. [PR1437569](#)
- Updating the IDP security package offline might fail in SRX Series devices. [PR1466283](#)

J-Web

- 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 nonconfigured routing-instance field. [PR1461599](#)

Layer 2 Ethernet Services

- DHCP requests might get dropped in a DHCP relay scenario. [PR1435039](#)
- The metric is not changing when configured under the DHCP. [PR1461571](#)

Network Address Translation (NAT)

- The nsd process might stop when SNMP queries deterministic NAT pool information. [PR1436775](#)
- Core files are generated while using NAT PBA in AA mode. [PR1443148](#)
- RTSP resource session is not found during NAT64 static mapping. [PR1443222](#)

- On SRX5000 line of devices with SPC3 card, when using source NAT, under high traffic load, a small fraction of TCP-SYN packets might be dropped due to the source NAT port failing to be allocated. Also, the NAT pool resources might leak over time. [PR1443345](#)
- Packet loss happens during cold synchronization from the secondary node after rebooting. [PR1448252](#)
- 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](#)
- The flowd or srpxfe process might stop when traffic is processed by both ALGs and NAT. [PR1471932](#)
- Issuing the **show security nat source paired-address** command might return an error. [PR1479824](#)

Network Management and Monitoring

- 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](#)
- The flowd or srpxfe process might stop immediately after committing the jflowv9 configuration or after upgrading to affected releases. [PR1471524](#)
- SNMP trap coldStart agent-address becomes 0.0.0.0. [PR1473288](#)

Platform and Infrastructure

- Memory leak might occur on the data plane during composite next-hop installation failure. [PR1391074](#)
- The **show security flow session** command fails with error messages when SRX4600 has over a million routing entries. [PR1408172](#)
- On 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](#)
- SPC in slot1 of node0 remained in offline state for more than 1 hour after the cluster was upgraded from Junos OS Release 18.2R2-S1.3 to Junos OS Release 18.2X41.1. [PR1423169](#)
- 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](#)
- Packet loss is caused by FPGA back pressure on SPC3. [PR1429899](#)
- REST API does not work properly. [PR1430187](#)
- Packet Forwarding Engine pause might be seen on the SRX1500 device. [PR1431380](#)

- The false license alarm might be seen even if there is a valid license. [PR1431609](#)
- When changing the decrypt mirror interface in the SSL proxy service configuration, it does not reflect properly in the Packet Forwarding Engine. [PR1434595](#)
- On SRX4100 and SRX4200 devices, when LACP is configured on the reth interface, the interface flaps when Routing Engine is busy. [PR1435955](#)
- LACP traffic is distributed evenly on ingress child links but not on egress links. [PR1437098](#)
- The ksyncd process might crash and restart on SRX Series devices. [PR1440576](#)
- The chassis cluster might stuck at CS FL state after rebooting. [PR1440938](#)
- The configured RPM probe server hardware timestamp does not respond with correct timestamp to the RPM client. [PR1441743](#)
- The RPM udp-ping probe does not work in a multiple routing instance scenario. [PR1442157](#)
- ARP resolution might fail after ARP HOLD NHs are added and deleted continuously. [PR1442815](#)
- On the SRX300 line of devices, the interface LED does not work properly. [PR1446035](#)
- 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](#)
- REST API process will become unresponsive when a number of requests come at a high rate. [PR1449987](#)
- Traffic loss might occur when there are around 80,000 routes in FIB. [PR1450545](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- VM core files might be generated if the configured sampling rate is more than 65,535. [PR1461487](#)
- On the SRX300 line of devices, you might encounter Authentication-Table loading slowly while using user-identification. [PR1462922](#)
- The AE interface cannot be configured on an SRX4600 device. [PR1465159](#)
- On SRX1500 and the SRX4000 line of devices, physically disconnecting the cable from fxp0 interface causes hardware monitor failure and redundancy group failover, when the device is the primary node in a chassis cluster. [PR1467376](#)
- The RGx might fail over after RG0 failover in a rare case. [PR1479255](#)

Routing Policy and Firewall Filters

- The NSD process might stop due to a memory corruption issue. [PR1419983](#)
- The ipfd generates a core file while scaling cases 6-1. [PR1431861](#)
- An SRX1500 device allows only a maximum of 256 policies with counting enabled. [PR1435231](#)

- Two ipfd processes appear in ps command and the process pauses. [PR1444472](#)
- During commit, the nsd_vrf_group_config_lsys log messages are displayed. [PR1446303](#)
- Security policies cannot synchronize between Routing Engine and Packet Forwarding Engine on SRX Series devices. [PR1453852](#)
- 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](#)
- Some domains are not resolved by the SRX Series devices when using DNS address book. [PR1471408](#)
- The count option in security policy does not take effect even if the policy count is enabled. [PR1471621](#)
- Support for dynamic tunnels on SRX Series devices was mistakenly removed. [PR1476530](#)

Routing Protocols

- SSH login might fail if a user account exists in both local database and RADIUS or TACACS+. [PR1454177](#)
- The rpd might stop when both instance-import and instance-export policies contain as-path-prepend action. [PR1471968](#)
- The routing protocol process (rpd) crashes while processing a specific BGP update information. [PR1448425](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)

Services Applications

- The flowd process stops when SRX5800 devices work at SPC3 mix mode with 1 SPC3 card and 7 SPC2 cards. [PR1448395](#)

Unified Threat Management (UTM)

- The command **show security utm web-filtering status** now provides additional context when the status of EWF is down. [PR1426748](#)
- Memory issue due to SSL proxy whitelist or whitelist URL category. [PR1430277](#)
- Adjust core allocation ratio for on-box antivirus. [PR1431780](#)
- On SRX Series devices, memory might leak if Websense Redirect Web Filtering is configured. [PR1445222](#)
- Increase the scale number of UTM profile or policy for the SRX1500 device, and the SRX4000 and SRX5000 lines of devices. [PR1455321](#)
- The utmd process might pause after deactivating UTM configuration with predefined category upgrading used. [PR1478825](#)

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](#)
- IPsec VPN is missing half of the IKE SA and IPsec SA is showing incorrect port number when scaling to 1000 IKEv1 AutoVPN tunnels. [PR1399147](#)
- The IKE and IPsec configuration under groups is not supported. [PR1405840](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when the 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](#)
- The established tunnels might remain unchanged when an IKE gateway is changed from AutoVPN to Site-to-Site VPN. [PR1413619](#)
- The iked process might crash due to misconfiguration in IPsec VPN network [PR1416081](#)
- The IKED process might stop when IKE and IPsec SA rekey happen simultaneously. [PR1420762](#)
- The VPN tunnel might flap when IKE and IPsec rekey happen simultaneously. [PR1421905](#)
- Old tunnel entries might be observed in the output of show security, IPsec or IKE SA. [PR1423821](#)
- The **show security ipsec statistics** command output displays buffer overflow and wraps around 4,---,---,--- count. [PR1424558](#)
- 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](#)
- On the SRX5000 line of devices with SPC3, with P2MP and IKEv1 configured, if negotiation fails on the peer device, then multiple IPSec SA entries are created on the device if the peer keeps triggering new negotiation. [PR1432852](#)
- IPsec rekey trigger is not working 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](#)

- On SRX Series devices, fragments exit VPN traffic earlier than required by ingress packet sizes. [PR1435700](#)
- The IKED stops on the SRX5000 line of devices with SPC3 when IPsec VPN or IKE is configured. [PR1443560](#)
- IPsec VPN traffic drop might be seen on SRX Series platforms with NAT-T scenario. [PR1444730](#)
- After a long time (a few hours) of traffic during a mini PDT test, the number of IPsec tunnels is much higher than expected. [PR1449296](#)
- Some IPsec tunnels flap after RGs fail over on the 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 the traffic selector randomly while the tunnel is still up, causing complete outage. [PR1456301](#)
- Traffic is not sent out through an IPsec VPN after update to Junos OS Release 18.2 or later. [PR1461793](#)
- The IPsec VPN tunnels cannot be established if overlapped subnets are configured in traffic selectors. [PR1463880](#)
- IPsec tunnels might lose connectivity on SRX Series devices after chassis cluster failover when using AutoVPN point-to-multipoint mode. [PR1469172](#)
- The kmd process might crash continually after the chassis cluster failover in the IPsec ADVPN scenario. [PR1479738](#)

Resolved Issues: 19.2R1

Application Firewall

- Fail to match permit rule in Application Firewall (AppFW) rule set. [PR1404161](#)

Application Identification

- IDP install failing on secondary node due to AI installation failure. [PR1336145](#)

Application Layer Gateways (ALGs)

- DNS requests with the EDNS option might be dropped by the DNS ALG. [PR1379433](#)
- On all SRX Series platforms, SIP/FTP ALG does not work when SIP traffic with source NAT goes through the SRX Series devices. [PR1398377](#)
- The TCP rst packet is dropped when any TCP proxy-based feature and **rst-invalidate-session** are enabled simultaneously. [PR1430685](#)

Chassis Clustering

- The SNMP trap sends wrong info with **Manual failover**. [PR1378903](#)
- Traffic with domain name address might fail for 3-5 minutes after RG0 failover on SRX Series platforms. [PR1401925](#)
- The flowd process stops when updating or deleting a GTP tunnel. [PR1404317](#)
- Mixed mode (SPC3 coexisting with SPC2 cards) high availability (HA) IP-Monitoring fails on secondary node with **secondary arp entry not found** error [PR1407056](#)
- The SRX Series devices might be potentially overwritten with an incorrect buffer address when detailed logging is configured under the GTPv2 profile. [PR1413718](#)
- Starting with Junos OS Release 18.4, at most, 6 pdn connects can be contained in a pdp context response; otherwise, the response will be dropped. [PR1422877](#)
- Memory leaks might be seen on the jsqsyncd process on SRX chassis clusters [PR1424884](#)
- RG0 failover sometimes causes FPC offline/present status. [PR1428312](#)

Flow-Based and Packet-Based Processing

- Control traffic loss might be seen on SRX4600 platform. [PR1357591](#)
- On SRX1500 devices, the activity LED (right LED) for 1-Gigabit Ethernet/10-Gigabit Ethernet port is not on although traffic is passing through that interface. [PR1380928](#)
- Password recovery menu is not shown up on SRX device. [PR1381653](#)
- Large file downloads slow down for many seconds. [PR1386122](#)
- On the SRX300 line of devices default configuration changed. [PR1393683](#)
- Switching interface mode between family **ethernet-switching** and **family inet/inet6** might cause traffic loss. [PR1394850](#)
- SRX to not strip vlan added by native vlan id command. [PR1397443](#)
- Increase DAG feed scale number to 256 from 63. [PR1399314](#)
- CPU is hitting 100 percent with fragmented traffic. [PR1402471](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when PowerMode IPsec is enabled, the **show security flow statistics** and **show security flow session tunnel summary** commands will not count or display the number of packets processed within PowerMode IPsec, because these packets do not go through regular flow path. [PR1403037](#)
- Downloads might stall and/or completely fail when utilizing services that are reliant on TCP proxy. [PR1403412](#)
- The flowd process stops and all cards are brought offline. [PR1406210](#)
- The RG1 failover does not happen immediately when the SPC3 card crashes. [PR1407064](#)

- The flowd process might crash if **enable-session-cache** knob is configured under the SSL termination profile. [PR1407330](#)
- Support for LAG interface with PowerMode IPsec. [PR1407231](#)
- The kernel might stop on the secondary node when committing **set system management-instance** command. [PR1407938](#)
- On SRX1500 platform, traffic is blocked on all interfaces after configuring **interface-mac-limit** on one interface. [PR1409018](#)
- Memory leak if AAMW is enabled. [PR1409606](#)
- Packets might get dropped in chassis cluster Z mode with local interface configured. [PR1410233](#)
- Session capacity of SRX340 does not match SRX345. [PR1410801](#)
- While PMI is ON, IPsec encrypted statistics on the Routing Engine **show security ipsec statistics** are not working anymore for fragment packets. [PR1411486](#)
- PEM 0 or PEM 1 or FAN, I2C failure major alarm might be set and cleared multiple times. [PR1413758](#)
- HA packets might be dropped on SRX5000 line of devices with IOC3 or IOC2 cards. [PR1414460](#)
- On SRX1500, SRX4100, SRX4200, SRX4600, and SRX5000 line of devices with SPC3 card, if SSL proxy is configured, the firewall FPC CPU might spike above 80 percent and traffic might be lost. [PR1414467](#)
- Any traffic originated from the device itself might be dropped in the IPsec tunnel. [PR1414509](#)
- The input and output bytes or bps statistic values might not be identical for the same size of packets. [PR1415117](#)
- The reth interfaces are now supported when configuring SSL decryption mirroring (**mirror-decrypt-traffic interface**) [PR1415352](#)
- Force clearing **Client Session** from flow does not clean up **Proxy session**. [PR1415756](#)
- Traffic would be dropped if SOF is enabled in a chassis cluster in active/active mode. [PR1415761](#)
- Juniper Sky ATP does not escape the \ inside the username before the metadata is sent to the cloud. [PR1416093](#)
- The flowd process stops on the SRX5000 Series or SRX4000 lines of devices when large-size packets go through IPsec tunnel with the post-fragment check. [PR1417219](#)
- TCP segmented client side session fails to create transparent proxied relay session, and session stays idle. [PR1417389](#)
- Best path selected keeps changing at regular intervals even when no violation is reported. [PR1417926](#)
- Traffic might be lost on the SRX Series device if IPsec session affinity is configured with **ipsec-performance-acceleration** command. [PR1418135](#)
- Group VPN IKE security-associations can not establish before RG0 failover. [PR1419341](#)

- On all SRX Series devices firewalls, if the traffic-log feature is configured, logs might incorrectly display IPv4 addresses in an IPv6 format [PR1421255](#)
- The **show security flow session session-identifier < sessID>** is not working if the session ID is bigger than 10M on SRX4600 platform. [PR1423818](#)
- The tunnel-id information is displayed in the flow session. [PR1423889](#)
- Replace **bypass-on-dns-cache-miss** command with **drop_on_dns_error** command in Web proxy profile. If **drop_on_dns_error** command is not set and DNS failure occurs for a session, that session passes through bypass mode. If **drop_on_dns_error** command is set and DNS failure occurs for a session, that session is dropped by Web proxy plug-in. [PR1430425](#)
- Support IPv6 session through Web proxy. [PR1433088](#)
- The applications which get declassified in the middle of session will not be identified properly. [PR1437816](#)
- Partial traffic might get dropped on an existing LAG. [PR1423989](#)
- Alarms due to high temperature when operating with expected temperatures. [PR1425807](#)
- PIM neighbors might not come up on SRX Series chassis cluster [PR1425884](#)
- The IPsec traffic going through SRX5000 line of devices with SPC2 cards installed causes SPU CPU utilization to be high. [PR1427912](#)
- Uneven distribution of CPU with high PPS on device. [PR1430721](#)
- SRX550M running Junos OS Release 18.4R1 shows PEM 1 output failure message where as with Junos OS Release 15.1X49 or Junos OS Release 18.1R3.3 it does not show any alarms. [PR1433577](#)
- Some webpages cannot be fully rendered. [PR1436813](#)

Infrastructure

- Increase in Junos OS image size for Junos OS Release 19.1R1. [PR1423139](#)

Interfaces and Routing

- On SRX4600 platform, the 40-Gigabit Ethernet might flap continuously by MAC local fault. [PR1397012](#)
- SRX Series devices cannot obtain IPv6 address through DHCPv6 when using a PPPoE interface with a logical unit number greater than 0. [PR1402066](#)

Intrusion Detection and Prevention (IDP)

- IDP might crash with the custom IDP signature. [PR1390205](#)
- Unable to configure **dynamic-attack-group** command. [PR1418754](#)

Installation and Upgrade

- ISSU failed from Junos OS Release 18.3R1.9 to Junos OS Release 18.4R1.4. [PR1405556](#)
- SRX1500 devices running Junos OS Release 15.1X49-D160 are unable to be upgraded or downgraded successfully to all releases built before 17 February 2019 [PR1407556](#)

J-Web

- In the J-Web dashboard, the **Security Resources** widget did not display absolute values. [PR1372826](#)
- J-Web now supports defining SSL-Proxy and redirect (block page) profiles when a policy contains dynamic applications. [PR1376117](#)
- Configuring using the CLI editor in J-Web generates an mgd core file. [PR1404946](#)
- The httpd-gk process stops, leading to dynamic VPN failures and high Routing Engine CPU utilization (100 percent). [PR1414642](#)
- Risk report, when generated in IE browser, appears completely out of alignment and XML tags are displayed. [PR1415767](#)
- J-Web configuration change for an address set using the search function results in a commit error. [PR1426321](#)
- J-Web not working when logged in as read-only user. [PR1428520](#)
- IRB interface is not available in zone option of J-Web. [PR1431428](#)

Logical Systems and Tenant Systems

- Tenant system administrator can change vlan assignment beyond the allocated tenant system. [PR1422058](#)

Multiprotocol Label Switching (MPLS)

- RPD might restart unexpectedly when **no-cspf** is configured and lo0 is not included under protocol rsvp. [PR1366575](#)

Network Address Translation (NAT)

- SRX SPC3 mix mode, NAT SPC3 core files are generated at `../sysdeps/unix/sysv/linux/raise.c:55`. [PR1403583](#)

Network Management and Monitoring

- The **set system no-redirects** setting does not take effect for the reth interface. [PR894194](#)
- The chassisd might crash and restart after the AGENTX session timeout between master(snmpd) and sub-agent. [PR1396967](#)

Platform and Infrastructure

- In chassis cluster redundancy group failover scenario, on SRX5600 and 5800 platforms, if the failover is caused by interface monitoring failure, the failover on Packet Forwarding Engine side (that is, data plane) might be slow (for example, impact on BFD session up to several seconds). This issue might result in protocol and traffic outage. [PR1385521](#)
- The flowd process might crash if there are too many IPsec tunnels [PR1392580](#)
- Complete device outage might be seen when an SPU VM core file is generated. [PR1417252](#)

- Some applications might not be installed during upgrade from lower version which does not support FreeBSD 10 to FreeBSD 10(based system). [PR1417321](#)
- On SRX Series devices, flowd process stops might be seen. [PR1417658](#)
- Routing Engine CPU utilization is high and eventd process is consuming a lot of resources. [PR1418444](#)
- On SRX4600 device, commit failed while configuring 2047 VLAN IDs on the reth interface. [PR1420685](#)

Routing Policy and Firewall Filters

- Memory leak in nsd prevents change from taking effect. [PR1414319](#)
- The flowd process (responsible for traffic forwarding in SRX Series devices) stops on SRX Series devices while deleting a lot of policies from Junos Space. [PR1419704](#)
- A commit warning will now be presented to the user when a traditional policy is placed below a unified policy. [PR1420471](#)
- The dynamic-address summary's IP entry count does not include IP entries in root logical system. [PR1422525](#)
- If restarting NSD fails, there is no any indication or symptom, and users don't know it. So a new alarm is added to indicate this failure. [PR1422738](#)
- The ipfd generates a core file while scaling cases 6-1. [PR1431861](#)

Unified Threat Management (UTM)

- Whitelist/blacklist does not work for HTTPS traffic going through Web proxy. [PR1401996](#)
- On SRX Series devices, when configuring Enhanced Web Filtering on the CLI, the autocomplete function did not properly handle or suggest custom categories. [PR1406512](#)
- On SRX Series devices, when using Unified Policies and Web filtering (EWF) without SSL proxy, the Server Name Indication (SNI) might not be identified correctly and the RT_UTM logs were recording incomplete information. [PR1410981](#)
- Unable to achieve better Avira AV TP on SRX4600 due to reaching mbuf high watermark. [PR1419064](#)
- UTM Web filtering status shows down when using Hostname [routing-instance synchronization failure]. [PR1421398](#)
- When using Unified Policies, the base-filter for certain UTM profiles might not be applied correctly. [PR1424633](#)
- The custom-url-categories are now pushed correctly to the Packet Forwarding Engine under all circumstances. [PR1426189](#)

User Interface and Configuration

- Tenant system administrator cannot view the configuration with **Empty Database** message when configuring tenant system using **groups**. [PR1422036](#)

VPNs

- On SRX1500 device, when configuring IPsec VPN and BGP simultaneously, the kmd process might stop and generate a core file if BGP peers reach approximately 350. All of the VPN tunnels will be disconnected during the pause. [PR1336235](#)
- SPC3 IKE SA detail output is not showing proper traffic statistics. [PR1371638](#)
- The pkid process might stop after RG0 failover. [PR1379348](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, the **show security ike security-association detail** command does not display local IKE-ID field correctly. [PR1388979](#)
- A few VPN tunnels do not forward traffic after RG1 failover. [PR1394427](#)
- The kmd process might stop when SNMP polls for the IKE SA. [PR1397897](#)
- VPN tunnels flap after adding or deleting a configuration group in edit private mode on a clustered setup. [PR1400712](#)
- Syslog is not generated when IKE gateway rejects duplicate IKE ID connection. [PR1404985](#)
- Idle IPsec VPN tunnels without traffic and with ongoing DPD probes can be affected during RG0 failover. [PR1405515](#)
- Not all the tunnels are deleted when authentication algorithm in IPsec proposal is changed. [PR1406020](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, the incoming packet's flow context information is not reset correctly when the packet is dropped in IPsec acceleration module. This will cause subsequent packets to be incorrectly processed as IPsec packets and results in the crash. To address this issue, SRX Series device now resets the flow context before dropping the packet in all relevant modules including IPsec acceleration module. [PR1407910](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when SRX Series device is configured in IKEv1 and NAT traversal is active, after a successful IPsec rekey, IPsec tunnel index might change. In such a scenario, there might be some traffic loss for a few seconds. [PR1409855](#)
- Traffic drops on peer due to bad SPI after first reauthentication. [PR1412316](#)
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when SRX Series device is configured to initiate IKEv2 reauthentication when NAT traversal is active, occasionally reauthentication might fail. [PR1414193](#)
- The flowd/srxpfe process might stop when traffic selector is used for IPsec VPN [PR1418984](#)
- The **show security ike sa detail** command shows incorrect value in **IPSec security associations** column. [PR1423249](#)

- On the SRX5000 line of devices with SPC3, with P2MP and IKEv1 configured, if negotiation fails on the peer device, then multiple IPsec SA entries are created on the device if the peer keeps triggering new negotiation. [PR1432852](#)
- On SRX Series devices with SPC3, should send IKE delete notification to peer when traffic selector configuration is changed for a specific AutoVPN. [PR1426714](#)
- The kmd process stops and generates a core file after running the **show security ipsec traffic-selector** command. [PR1428029](#)
- IPsec rekey triggers for when sequence number in AH and ESP packet is about to exhaust is not working. [PR1433343](#)

SEE ALSO

[What's New | 303](#)

[What's Changed | 313](#)

[Known Limitations | 317](#)

[Open Issues | 319](#)

[Documentation Updates | 340](#)

[Migration, Upgrade, and Downgrade Instructions | 341](#)

Documentation Updates

There are no errata or changes in Junos OS Release 19.2R3 documentation for the SRX Series.

SEE ALSO

[What's New | 303](#)

[What's Changed | 313](#)

[Known Limitations | 317](#)

[Open Issues | 319](#)

[Resolved Issues | 321](#)

[Migration, Upgrade, and Downgrade Instructions | 341](#)

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 | 303](#)

[What's Changed | 313](#)

[Known Limitations | 317](#)

[Open Issues | 319](#)

[Resolved Issues | 321](#)

[Documentation Updates | 340](#)

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the [High Availability User Guide](#).

For additional information about using ISSU on security devices, see the [Chassis Cluster User Guide for SRX Series Devices](#).

For information about ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\)](#) Web application.

Licensing

Starting in 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that have been developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see [Juniper Flex Program](#).

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. <https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home

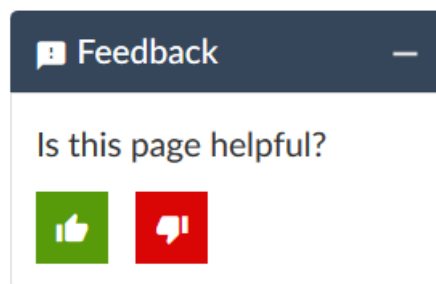
NOTE: To obtain information about the components that are supported on the devices, and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **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 J-Care or JNASC support contract, or are covered under warranty, and need post sales technical support, you can access our tools and resources online or open a case with JTAC.

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://support.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://support.juniper.net/support/>
- Search for known bugs: <https://kb.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://support.juniper.net/support/downloads/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>

- Join and participate in the Juniper Networks Community Forum: <https://forums.juniper.net>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

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

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

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

Revision History

22 April 2021—Revision 3, Junos OS Release 19.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 2, Junos OS Release 19.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 July 2020—Revision 1, Junos OS Release 19.2R3— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 May 2020—Revision 1, Junos OS Release 19.2R2— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 March 2020—Revision 10, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 January 2020—Revision 9, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 November 2019—Revision 8, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 November 2019—Revision 7, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 October 2019—Revision 6, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 September 2019—Revision 5, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 July 2019—Revision 4, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 July 2019—Revision 3, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 July 2019—Revision 2, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

27 June 2019—Revision 1, Junos OS Release 19.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.