

Release Notes: Junos[®] OS Release 20.3R3 for the ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX

25 November 2022

Contents	Introduction 13
	Junos OS Release Notes for ACX Series 13
	What's New 14
	What's New in Release 20.3R3 14
	What's New in Release 20.3R2 14
	What's New in Release 20.3R1 15
	What's Changed 19
	What's Changed in Release 20.3R3 19
	What's Changed in Release 20.3R2 19
	What's Changed in Release 20.3R1 20
	Known Limitations 21
	General Routing 22
	Open Issues 23
	General Routing 24
	Interfaces and Chassis 26
	Network Management and Monitoring 26
	Virtual Chassis 26

Resolved Issues | 26**Resolved Issues: 20.3R3 | 27****Resolved Issues: 20.3R2 | 28****Resolved Issues: 20.3R1 | 30****Documentation Updates | 33****Migration, Upgrade, and Downgrade Instructions | 34****Upgrade and Downgrade Support Policy for Junos OS Releases | 34****Junos OS Release Notes for cSRX | 35****What's New | 36****What's New in Release 20.3R3 | 36****What's New in Release 20.3R2 | 36****What's New in Release 20.3R1 | 37****What's Changed | 37****What's Changed in Release 20.3R3 | 37****What's Changed in Release 20.3R2 | 37****What's Changed in Release 20.3R1 | 38****Known Limitations | 39****Open Issues | 39****Resolved Issues | 40****Resolved Issues: 20.3R3 | 40****Resolved Issues: 20.3R2 | 40****Resolved Issues: 20.3R1 | 40****Junos OS Release Notes for EX Series | 40****What's New | 41****What's New in Release 20.3R3 | 41****What's New in Release 20.3R2 | 41****What's New in Release 20.3R1 | 41****What's Changed | 53****What's Changed in Release 20.3R3 | 54****What's Changed in Release 20.3R2 | 54****What's Changed in Release 20.3R1 | 55****Known Limitations | 56****General Routing | 57****EVPN | 57**

Infrastructure	57
Platform and Infrastructure	57
Open Issues	58
General Routing	58
Authentication and Access Control	60
EVPN	60
Forwarding and Sampling	61
Infrastructure	61
Layer 2 Features	61
Network Management and Monitoring	62
Platform and Infrastructure	62
User Interface and Configuration	62
Resolved Issues	63
Resolved Issues: 20.3R3	63
Resolved Issues: 20.3R2	67
Resolved Issues: 20.3R1	69
Documentation Updates	72
Migration, Upgrade, and Downgrade Instructions	72
Upgrade and Downgrade Support Policy for Junos OS Releases	72
Junos OS Release Notes for JRR Series	74
What's New	74
What's New in Release 20.3R3	75
What's New in Release 20.3R2	75
What's New in Release 20.3R1	75
What's Changed	76
Known Limitations	76
Open Issues	76
Resolved Issues	77
Resolved Issues: 20.3R3	77
Resolved Issues: 20.3R2	77
Resolved Issues: 20.3R1	78
Documentation Updates	78
Migration, Upgrade, and Downgrade Instructions	79
Upgrade and Downgrade Support Policy for Junos OS Releases	79

Junos OS Release Notes for Juniper Secure Connect | 80

What's New | 81

What's New in Release 20.3R3 | 81

What's New in Release 20.3R2 | 81

What's New in Release 20.3R1 | 81

What's Changed | 82

What's Changed in Release 20.3R3 | 82

What's Changed in Release 20.3R2 | 83

What's Changed in Release 20.3R1 | 83

Known Limitations | 83

Open Issues | 83

Resolved Issues | 83

Resolved Issues: 20.3R3 | 83

Resolved Issues: 20.3R2 | 83

Resolved Issues: 20.3R1 | 84

Junos OS Release Notes for Junos Fusion for Enterprise | 84

What's New | 84

What's New in Release 20.3R3 | 85

What's New in Release 20.3R2 | 85

What's New in Release 20.3R1 | 85

What's Changed | 86

What's Changed in Release Junos OS 20.3R3 | 86

What's Changed in Release Junos OS 20.3R2 | 86

What's Changed in Release Junos OS 20.3R1 | 86

Known Limitations | 87

Open Issues | 87

Resolved Issues | 88

Resolved Issues: Release 20.3R3 | 88

Resolved Issues: Release 20.3R2 | 88

Resolved Issues: Release 20.3R1 | 88

Documentation Updates | 89

Migration, Upgrade, and Downgrade Instructions | 89

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 89

Upgrading an Aggregation Device with Redundant Routing Engines | 91

- Preparing the Switch for Satellite Device Conversion | 92
- Converting a Satellite Device to a Standalone Switch | 93
- Upgrade and Downgrade Support Policy for Junos OS Releases | 93
- Downgrading Junos OS | 94

Junos OS Release Notes for Junos Fusion Provider Edge | 95

What's New | 95

What's Changed | 96

Known Limitations | 96

Open Issues | 97

Resolved Issues | 97

- Resolved Issues: 20.3R3 | 97
- Resolved Issues: 20.3R2 | 98
- Resolved Issues: 20.3R1 | 98

Documentation Updates | 98

Migration, Upgrade, and Downgrade Instructions | 99

- Basic Procedure for Upgrading an Aggregation Device | 99
- Upgrading an Aggregation Device with Redundant Routing Engines | 102
- Preparing the Switch for Satellite Device Conversion | 102
- Converting a Satellite Device to a Standalone Device | 104
- Upgrading an Aggregation Device | 106
- Upgrade and Downgrade Support Policy for Junos OS Releases | 106
- Downgrading from Junos OS Release 20.1 | 107

Junos OS Release Notes for MX Series | 108

What's New | 108

- New and Changed Features: 20.3R3 | 109
- New and Changed Features: 20.3R2 | 109
- New and Changed Features: 20.3R1 | 109

What's Changed | 131

- What's Changed in Release 20.3R3 | 131
- What's Changed in Release 20.3R2 | 132
- What's Changed in Release 20.3R1 | 134

Known Limitations | 137

- EVPN | 138
- General Routing | 138

Interfaces and Chassis	139
MPLS	139
Platform and Infrastructure	140
Routing Protocols	140
Network Management and Monitoring	140
Platform and Infrastructure	141
Routing Protocols	142
Subscriber Management and Services	142
Open Issues	143
Class of Service (CoS)	144
EVPN	144
Flow-based and Packet-based Processing	144
Forwarding and Sampling	144
General Routing	145
High Availability (HA) and Resiliency	150
Infrastructure	150
Interfaces and Chassis	150
Intrusion Detection and Prevention (IDP)	151
Layer 2 Ethernet Services	151
MPLS	151
Network Management and Monitoring	152
Platform and Infrastructure	152
Routing Protocols	153
Subscriber Access Management	155
User Interface and Configuration	155
Resolved Issues	156
Resolved Issues: 20.3R3	156
Resolved Issues: 20.3R2	165
Resolved Issues: 20.3R1	174
Documentation Updates	188
Migration, Upgrade, and Downgrade Instructions	188
Basic Procedure for Upgrading to Release 20.3R3	189
Procedure to Upgrade to FreeBSD 11.x-Based Junos OS	190
Procedure to Upgrade to FreeBSD 6.x-Based Junos OS	192

Upgrade and Downgrade Support Policy for Junos OS Releases | **194**

Upgrading a Router with Redundant Routing Engines | **195**

Downgrading from Release 20.3R3 | **195**

Junos OS Release Notes for NFX Series | **196**

What's New | **196**

What's New in Release 20.3R3 | **197**

What's New in Release 20.3R2 | **197**

What's New in Release 20.3R1 | **197**

What's Changed | **198**

What's Changed in Release 20.3R3 | **198**

What's Changed in Release 20.3R2 | **199**

What's Changed in Release 20.3R1 | **199**

Known Limitations | **199**

Open Issues | **200**

Platform and Infrastructure | **200**

Resolved Issues | **201**

Resolved Issues: 20.3R3 | **201**

Resolved Issues: 20.3R2 | **201**

Resolved Issues: 20.3R1 | **202**

Documentation Updates | **203**

Migration, Upgrade, and Downgrade Instructions | **203**

Upgrade and Downgrade Support Policy for Junos OS Releases | **203**

Basic Procedure for Upgrading to Release 20.3 | **204**

Junos OS Release Notes for PTX Series | **206**

What's New | **206**

What's New in Release 20.3R3 | **207**

What's New in Release 20.3R2 | **207**

What's New in Release 20.3R1 | **207**

What's Changed | **218**

What's Changed in Release 20.3R3 | **219**

What's Changed in Release 20.3R2 | **220**

What's Changed in Release 20.3R1 | **222**

Known Limitations | 224**General Routing | 225****MPLS | 226****Routing Protocols | 226****Open Issues | 227****General Routing | 227****Layer 2 Ethernet Services | 229****MPLS | 229****Routing Protocols | 229****Resolved Issues | 229****Resolved Issues: 20.3R3 | 230****Resolved Issues: 20.3R2 | 231****Resolved Issues: 20.3R1 | 232****Documentation Updates | 234****Migration, Upgrade, and Downgrade Instructions | 235****Basic Procedure for Upgrading to Release 20.3 | 235****Upgrade and Downgrade Support Policy for Junos OS Releases | 238****Upgrading a Router with Redundant Routing Engines | 239****Junos OS Release Notes for QFX Series | 239****What's New | 240****What's New in Release 20.3R3 | 241****What's New in Release 20.3R2 | 241****What's New in Release 20.3R1 | 241****What's Changed | 253****What's Changed in Release 20.3R3 | 254****What's Changed in Release 20.3R2 | 255****What's Changed in Release 20.3R1 | 257****Known Limitations | 259****Class of Service (CoS) | 259****Layer 2 Features | 259****Layer 2 Ethernet Services | 259****Platform and Infrastructure | 260****Routing Protocols | 260**

Open Issues | 261

EVPN | 261

High Availability (HA) and Resiliency | 262

Infrastructure | 262

Layer 2 Features | 262

Platform and Infrastructure | 262

Routing Protocols | 264

User Interface and Configuration | 265

Resolved Issues | 265

Resolved Issues: Release 20.3R3 | 266

Resolved Issues: Release 20.3R2 | 272

Resolved Issues: Release 20.3R1 | 276

Documentation Updates | 280

Migration, Upgrade, and Downgrade Instructions | 281

Upgrading Software on QFX Series Switches | 281

Installing the Software on QFX10002-60C Switches | 284

Installing the Software on QFX10002 Switches | 284

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 | 285

Installing the Software on QFX10008 and QFX10016 Switches | 287

Performing a Unified ISSU | 291

Preparing the Switch for Software Installation | 292

Upgrading the Software Using Unified ISSU | 292

Upgrade and Downgrade Support Policy for Junos OS Releases | 294

Junos OS Release Notes for SRX Series | 296

What's New | 296

What's New in Release 20.3R3 | 297

What's New in Release 20.3R2 | 297

What's New in Release 20.3R1 | 297

What's Changed | 306

What's Changed in Release 20.3R3 | 307

What's Changed in Release 20.3R2 | 307

What's Changed in Release 20.3R1 | 309

Known Limitations | 311**General Routing | 312****J-Web | 312****VPNs | 312****Open Issues | 313****Flow-Based and Packet-Based Processing | 313****General Routing | 313****Interfaces and Chassis | 314****Intrusion Detection and Prevention (IDP) | 314****J-Web | 314****User Interface and Configuration | 315****VPNs | 315****Resolved Issues | 315****Resolved Issues: 20.3R3 | 315****Resolved Issues: 20.3R2 | 318****Resolved Issues: 20.3R1 | 321****Documentation Updates | 324****Migration, Upgrade, and Downgrade Instructions | 324****Upgrade and Downgrade Support Policy for Junos OS Releases | 324****Junos OS Release Notes for vMX | 325****What's New | 326****What's New in Release 20.3R3 | 326****What's New in Release 20.3R2 | 326****What's New in Release 20.3R1 | 326****What's Changed | 327****What's Changed in Release 20.3R3 | 328****What's Changed in Release 20.3R2 | 328****What's Changed in Release 20.3R1 | 328****Known Limitations | 328****Open Issues | 329****Platform and Infrastructure | 329****Resolved Issues | 329****Resolved Issues: 20.3R3 | 329****Resolved Issues: 20.3R2 | 330**

Resolved Issues: 20.3R1	330
Licensing	330
Upgrade Instructions	330
Junos OS Release Notes for vRR	331
What's New	331
What's New in Release 20.3R3	331
What's New in Release 20.3R2	332
What's New in Release 20.3R1	332
What's Changed	333
What's Changed in Release 20.3R3	333
What's Changed in Release 20.3R2	333
What's Changed in Release 20.3R1	333
Known Limitations	334
Routing Protocols	334
Open Issues	334
Resolved Issues	335
Resolved Issues: 20.3R3	335
Resolved Issues: 20.3R2	335
Resolved Issues: 20.3R1	335
Junos OS Release Notes for vSRX	335
What's New	336
What's New in Release 20.3R3	336
What's New in Release 20.3R2	336
What's New in Release 20.3R1	337
What's Changed	340
What's Changed in Release 20.3R3	341
What's Changed in Release 20.3R2	341
What's Changed in Release 20.3R1	341
Known Limitations	341
Intrusion Detection and Prevention (IDP)	342
J-Web	342
Open Issues	342
J-Web	342
User Access and Authentication	343

Resolved Issues | 343

Resolved Issues: 20.3R3 | 343

Resolved Issues: 20.3R2 | 344

Resolved Issues: 20.3R1 | 345

Migration, Upgrade, and Downgrade Instructions | 346

Upgrading Software Packages | 347

Validating the OVA Image | 352

Upgrading Using ISSU | 353

Licensing | 353

Compliance Advisor | 353

Finding More Information | 354

Documentation Feedback | 354

Requesting Technical Support | 356

Self-Help Online Tools and Resources | 356

Creating a Service Request with JTAC | 357

Revision History | 357

Introduction

Junos OS runs on the following Juniper Networks® products: ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

These release notes accompany Junos OS Release 20.3R3 for the ACX Series, cSRX Container Firewall (cSRX), EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, virtual MX Series router (vMX), Virtual Route Reflector (vRR), and vSRX Virtual Firewall (vSRX). They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

- [In Focus guide](#)—We have a document called In Focus that provides details on the most important features for the release in one place. We hope this document will quickly get you to the latest information about Junos OS features. Let us know if you find this information useful by sending an e-mail to techpubs-comments@juniper.net.
- **Important Information:**
 - [Upgrading Using ISSU on page 353](#)
 - [Licensing on page 353](#)
 - [Compliance Advisor on page 353](#)
 - [Finding More Information on page 354](#)
 - [Documentation Feedback on page 354](#)
 - [Requesting Technical Support on page 356](#)

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 14](#)
- [What's Changed | 19](#)
- [Known Limitations | 21](#)
- [Open Issues | 23](#)
- [Resolved Issues | 26](#)

- Documentation Updates | 33
- Migration, Upgrade, and Downgrade Instructions | 34

These release notes accompany Junos OS Release 20.3R3 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 20.3R3 | 14
- What's New in Release 20.3R2 | 14
- What's New in Release 20.3R1 | 15

This section describes the new features or enhancements to existing features in Junos OS Release 20.3R3 for the ACX Series.

What's New in Release 20.3R3

There are no new features or enhancements to existing features for ACX Series routers in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for ACX Series routers in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Hardware

- We've added the following features to the ACX710 in Junos OS Release 20.3R1.

Table 1: Features Supported by the ACX710

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> • Support for hierarchical CoS. Support for up to three levels of hierarchical scheduling (physical interfaces, logical interfaces, and queues). [See Hierarchical Class of Service.]
Network management and monitoring	<ul style="list-style-type: none"> • Support for counters to display systemwide and interface-level statistics. You can configure new counters under flat-file (interface level, logical-interface-level, and ae-level statistics), interface (interface-level statistics), and Routing Engine profile (systemwide statistics) of the accounting options. [See Accounting Options Configurations.]
Timing and synchronization	<ul style="list-style-type: none"> • Support for logs, alarms, counters, and SNMP traps for Precision Time Protocol (PTP)/ Synchronous Ethernet. [See Enterprise-Specific SNMP Traps Supported by Junos OS and show chassis alarms.] • Support for the g.8275.1 profile on the ACX710. [See Assisted Partial Timing Support.] • Support for PTP G.8275.1. Use PTP profile-type g.8275.1 to enable the G.8275.1 profile. [See Profile Type.] • Support for limited images on the ACX710. [See Software Installation and Upgrade Overview.] • Support for RIP version 1, RIP version 2, and RIP next generation (RIPng) on PTX10008 routers. [See RIP and RIPng Overview.]

EVPNS

- **Multicast with IGMP or MLD snooping within VLANs for EVPN-MPLS (ACX5448, ACX5448-D, and ACX5448-M)**—Starting in Junos OS Release 20.3R1, ACX5448 routers support multicast with IGMP or Multicast Listener Discovery (MLD) snooping within VLANs on provider edge (PE) devices in an EVPN-MPLS multihoming environment. You can configure IGMP or MLD snooping with IGMPv2, IGMPv3, MLDv1, or MLDv2 in multiple routing instances of type **evpn**. Multicast receivers must be within the EVPN instance (EVI). If you have only intra-VLAN traffic, you can have the multicast sources within the EVI. (Otherwise, inter-VLAN multicast requires sources to be in an external Layer 3 PIM domain.)

With this support, PE devices:

- Process IGMPv2 and MLDv1 any-source multicast (ASM) (*,G) reports by default.

- Process IGMPv3 or MLDv2 reports in ASM mode (but only if you configure IGMPv3 or MLDv2 on all interfaces that receive multicast traffic).
- Drop IGMPv3 or MLDv2 source-specific multicast (SSM) (S,G) reports.

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

- **Multicast with IGMP or MLD snooping across VLANs for EVPN-MPLS (ACX5448, ACX5448-D, and ACX5448-M)**—Starting in Junos OS Release 20.3R1, ACX5448 routers support multicast with IGMP or MLD snooping across VLANs on provider edge (PE) devices in an EVPN-MPLS multihoming environment. You can configure IGMP or MLD snooping with IGMPv2, IGMPv3, MLDv1, or MLDv2 in multiple routing instances of type **evpn**.

Multicast receivers must be within the EVPN instance (EVI). Sources must be outside the EVI in a Layer 3 Protocol Independent Multicast (PIM) domain. All PE devices connect to a PIM gateway using Layer 3 interfaces on which they receive the multicast source traffic. Then IRB interfaces in PIM distributed designated router (DR) mode forward or route the multicast traffic locally to interested receivers.

PE devices can process ASM (*,G) reports by default, or IGMPv3 and MLDv2 SSM (S,G) reports with a configuration option.

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

- **Color-based mapping of EVPN-MPLS and EVPN services over SR-TE (ACX5448, EX9200, MX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can specify a color attribute along with an IP protocol next hop. The color attribute adds another dimension to the resolution of transport tunnels over static colored and BGP segment routing traffic-engineered (SR-TE) label-switched paths (LSPs). This type of resolution is known as the color-IP protocol next-hop resolution. To enable the color-IP protocol next-hop resolution, you must configure a resolution map and apply it to EVPN-MPLS and EVPN services, which include E-Line, E-LAN and E-Tree. With this feature, you can enable color-based traffic steering of EVPN-MPLS and EVPN services.

[See [Segment Routing LSP Configuration](#).]

Multicast

- **Support for BGP MVPN (ACX5448)**—Starting in Junos OS Release 20.3R1, ACX5448 routers support BGP MVPN (also known as “next generation,” or “NG,” MVPN) running on multipoint LDP provider tunnels, where BGP MVPN is the intra-AS and PIM-SM and multipoint LDP point-to-multipoint (P2MP) tunnels from the data plane. Other configurations and features are not supported in this release.

[See [Multiprotocol BGP MVPNs Overview](#).]

Network Management and Monitoring

- **Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can use the **probe** command to query the status of the probed interface. The proxy interface resides on the same node as the probed interface, or it can reside on a node to which the probed interface is directly connected.

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

To enable the **probe** command, configure the **extended-echo** statement under the **[edit system]** hierarchy.

[See [What is the Probe command?](#).]

Routing Protocols

- **Support for multiple MD5 for RIPv2 (ACX Series)**—Starting in Junos OS Release 20.3R1, you can define multiple MD5 authentication keys for RIPv2. This feature supports adding of MD5 keys with their **start-time**. RIPv2 packets are transmitted with MD5 authentication using the first configured key. RIPv2 authentication switches to the next key based on its configured key **start-time**. This provides automatic key switching without user intervention to change the MD5 keys as in the case of having only one MD5 key.

To enable multiple MD5 support for RIPv2, include the **authentication-selective-md5** statement at the **[edit protocols rip]** hierarchy level.

[See [Example: Configuring Route Authentication for RIP using multiple MD5 keys](#).]

- **Support for implicit filter for default EBGp route propagation behavior without policies (ACX Series, JRR200, MX204, vRR and PTX5000)**—Starting in Junos OS Release 20.3R1, we’ve introduced a new configuration hierarchy, **defaults ebgp no-policy** at the existing **[edit protocols bgp]** hierarchy level. The configuration option separates the default policy for **receive** and **advertise**, into separate clauses (**accept**, **reject**, or **reject-always**) to allow the route propagation behavior of EBGp speakers to vary independently from its default behavior.

In earlier releases, the default behavior of BGP was to receive and advertise all routes. With the introduction of this feature, the default behavior still remains to “accept” all routes for both **receive** and **advertise**, but you also have an option to reject routes by default.

With the **reject** configuration, you can reject routes of type inet unicast and inet6 unicast in instance types master, vrf, virtual-router, and non-forwarding. With the **reject-always** configuration, you can

reject all routes from being received or getting advertised, irrespective of address family or instance type. By using this feature, you can control traffic in leaf autonomous systems (AS) and thereby, prevent them from having to accidentally function as transit autonomous systems.

NOTE: The introduction of this implicit filter does not affect the existing deployments that rely on the default behavior.

[See [Implicit Filter for Default EBGp route propagation behavior without policies](#) and [defaults](#).]

- **IS-IS and OSPF support (ACX5448, ACX5448-M, and ACX5448-D)**—Starting in Release 20.3R1, Junos OS supports the following features for IS-IS and OSPF:
 - Base Segment routing (SR) support for prefix SID and Segment Routing Global Block (SRGB)
 - Anycast SID
 - BGP-LS
 - SRMS (LDP mapping server)
 - OAM
 - Topology-Independent Loop-Free Alternate (TI-LFA) link and node protection

In addition to these features, OSPF also supports Unnumbered Ethernet interface.

[See [Introduction to OSPF](#) and [IS-IS Overview](#).]

Segment Routing

- **Support for LDP Tunneling over Segment Routing Traffic Engineering (MX Series, PTX Series, and ACX5448)**—Starting in Junos OS Release 20.3R1, you can tunnel LDP LSPs over Segment Routing Traffic Engineering (SR-TE) in your network. Tunneling LDP over SR-TE provides consistency and co-existence of both LDP LSPs and SR-TE LSPs.

[See [Tunneling LDP over SR-TE](#).]

SEE ALSO

[What's Changed | 19](#)

[Known Limitations | 21](#)

[Open Issues | 23](#)

[Resolved Issues | 26](#)

[Documentation Updates | 33](#)

[Migration, Upgrade, and Downgrade Instructions | 34](#)

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 19](#)
- [What's Changed in Release 20.3R2 | 19](#)
- [What's Changed in Release 20.3R1 | 20](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.3R3 for the ACX Series routers.

What's Changed in Release 20.3R3

Junos OS XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or **set refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

What's Changed in Release 20.3R2

Junos XML API and Scripting

- **The jcs:invoke() function supports suppressing root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#)..]

MPLS

- The `show mpls lsp extensivel` and `show mpls lsp detail` commands display next-hop gateway LSPid—When you use the `show mpls lsp extensivel` and `show mpls lsp detail` commands, you'll see next-hop gateway LSPid in the output.

User Interface and Configuration

- Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—The Junos OS CLI exposes the `verbose` statement at the `edit system export-format json` hierarchy level. The default format to export configuration data in JSON changed from `verbose` format to `ietf` format starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `edit system export-format json` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See [export-format](#).]

What's Changed in Release 20.3R1

Junos OS, XML, API, and Scripting

- Changes to Junos XML RPC request tag names (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—We've updated the Junos XML request tag name for some operational RPCs to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include:
 - Most, but not all, request tag names that start with `show` replace `show` with `get` in the name.
 - Uppercase characters are converted to lowercase.

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

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 lsdist.0 and lsdist.1 routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router-ids instead of prefixes. In earlier Junos OS releases, multiple secondary loopback addresses in TED were added into lsdist.0 and lsdist.1 routing tables as part of node characteristics and advertised them as the router-id.
- **Inet6 is disabled in VT interface (ACX5448)**—Starting in this release, the **inet6** statement at the **edit interfaces vt-interface-number unit unit-number family** hierarchy level is disabled.

SEE ALSO

What's New 14
Known Limitations 21
Open Issues 23
Resolved Issues 26
Documentation Updates 33
Migration, Upgrade, and Downgrade Instructions 34

Known Limitations

IN THIS SECTION

- [General Routing | 22](#)

Learn about known limitations in this release for the ACX 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

- The time consumed on 1-Gigabit Ethernet performance is not the same compared to 10-Gigabit Ethernet. Compensation is done to bring the mean value under class A but the peak-to-peak variations are high and might go beyond 100 ns. It has a latency variation with peak-to-peak variations of around 125 ns—250 ns without any traffic. (For example, 5—10 percent of the mean latency introduced by the each phy, which is of around 2.5 microseconds). [PR1437175](#)
- With an asymmetric network connection, a 10-Gbps MACsec port connected to a 10-Gbps channelized port, high and asymmetric T1 and T4 time errors are seen. This situation introduces a high two-way time error and also different CF updates in the forward and reverse paths. [PR1440140](#)
- With the MACsec feature enabled and introduction of traffic, the peak-to-peak value varies with the percentage of traffic introduced. Finding the maximum and mean values of the time errors with different traffic rates (for example, two-router scenarios) can have the maximum value as high as 1054 ns with 95 percent traffic, 640 ns for 90 percent traffic, and 137 ns with no traffic. [PR1441388](#)
- Difference between minimum and maximum latency is very high in a latency test. [PR1483370](#)
- Transient traffic drop is seen on an aggregated Ethernet interface if a member does not carry traffic flaps. [PR1486997](#)
- The throughput test fails for the 64 bytes packet in the 100-Gigabit Ethernet line rate. [PR1489248](#)
- EVPN-VPWS, L3VPN and L2VPN FRR convergence time with aggregated Ethernet as the Active core interface is not meeting <50ms and may be 100ms to 150ms. [PR1492730](#)
- On the ACX710 router, traffic loss is beyond the tolerance limit of 200 ms during convergence. [PR1499965](#)
- On ACX710, the PTP clock recovery is re-started when the clksyncd process is re-started. This will result in the PTP lock state moving to freerun on the clksyncd process restart. [PR1502162](#)
- Not able to scale BFD to 1024 sessions with IPv4 and IPv6. [PR1502170](#)
- Satellites do not track intermittently with GPS-only constellation. [PR1505325](#)
- Do not foresee any impact on customer use cases with current implementation. The explicit VLAN configuration knob available in PTP helps to cover the possible use cases related to VLAN mapping. Will be checking with Broadcom, if any alternate and better solution exists. If yes, then it will be considered for implementation in future release. [PR1507809](#)
- Unexpected delay counter values are seen in the output of the **show ptp statistics detail** command when the upstream primary clock stops sending the PTP packets. [PR1508031](#)
- Inconsistencies in the PTP lock status behavior is observed during chassis control restart. [PR1508385](#)
- High FRR convergence is observed. [PR1515512](#)
- Sometimes PTP takes longer time to lock PTP after being locked to GPS. [PR1527346](#)

- The announce interval -1,0 sent from the upstream primary clock gets stuck in the **HOLDOVER IN** mode. [PR1529761](#)
- PTP to 1PPS noise transfer test fails for frequency 1.985 Hz. [PR1522666](#)
- SyncE to 1PPS transient test results do not meet G.8273.2 SyncE to 1PPS transient metric. [PR1522796](#)
- Virtual port and T-GM are not supported in Junos OS Release 20.3R1. Only G.8275.1 T-BC is supported. [PR1533018](#)
- The g8275.1 announcement or synchronization interval rate range is not as per FS. [PR1542516](#)
- Whenever PTP configuration is deleted, syncE goes into HOLDOVER state and comes back to locked. [PR1546681](#)
- In ACX5448 device, the ping shall fail if the MAC address of the device is modified to a static MAC address as BCM supports only one base MAC address. [PR1553472](#)
- When ACX5448 or ACX710 is configured as Two-Way Active Measurement Protocol (TWAMP) server and a client sends a start session message, the response from ACX to send **start session ack** is delayed by 10 seconds which might cause the session to fail. [PR1556829](#)

SEE ALSO

What's New 14
What's Changed 19
Open Issues 23
Resolved Issues 26
Documentation Updates 33
Migration, Upgrade, and Downgrade Instructions 34

Open Issues

IN THIS SECTION

- [General Routing | 24](#)
- [Interfaces and Chassis | 26](#)
- [Network Management and Monitoring | 26](#)
- [Virtual Chassis | 26](#)

Learn about open issues in this release for the ACX 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

- The **inet6-precendence** under **set class of service rewrite-rules** is not supported on ACX5448. [PR1344340](#)
- 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](#)
- On ACX6360 routers, enhancement is needed for FRR BER threshold SNMP support. [PR1383303](#)
- On ACX6360 router, Tx power cannot be configured using + sign. [PR1383980](#)
- The ccc logs are not compressed after rotation. [PR1398511](#)
- A jnxIfOtnOperState trap notification is sent for all ot-interfaces. [PR1406758](#)
- DHCP clients are not able to scale to 96,000. [PR1432849](#)
- Protocols get forwarded when using a nonexisting SSM map source address in IGMPv3 instead of pruning. [PR1435648](#)
- Memory leaks are expected in this release. [PR1438358](#)
- Drop profile, maximum threshold may not be reached when the packet size is other than 1000 bytes. [PR1448418](#)
- The CFM remote MEP is not coming up after configuration or remains in start state. [PR1460555](#)
- Multiple RMEP's are unsupported due to which remote defect indication is seen. [PR1478346](#)
- On ACX710 routers, packet drop is observed after changing ALT port cost for RSTP. [PR1482566](#)
- The syslog error messages related to **ACX_DFW_CFG_FAILED** are observed. [PR1490940](#)
- On ACX6360 routers, the port mirror will not work when the port-mirroring is configured with the firewall filter. [PR1491789](#)
- If we configure DHCP option 012 host-name in DHCP server configuration and the actual base configuration file also has the host-name in it, then overwriting of the base configuration file's host-name with the DHCP option 012 host-name is happening. [PR1503958](#)
- On ACX710 routers, when the following steps are done for PTP, chassis does not lock: 1. Use one or two ports as source for chassis synchronization and lock both PTP and SyncE locked. 2. Disable both

logical interfaces. 3. Restart clksyncd. 4. Rollback 1. As a workaround, you can avoid this issue by deleting the PTP configuration, restarting clksyncd, and then reconfiguring PTP. [PR1505405](#)

- Experimental (exp) remarking is supported only for single MPLS label packet. [PR1509627](#)
- Experimental (exp) remarking is supported only for single MPLS label packet. [PR1509635](#)
- For ACX710, if the console cable is plugged in and the terminal connection is active and sending characters to the interface, the system boot may be interrupted and the ACX710 boot will be stalled at the uboot# prompt. [PR1513553](#)
- On ACX710 routers, when working in holdover mode after locking to GNSS, will rely on the local oscillator for holdover. The clock class will hence be for stratum 3/e clock. The APTS mode of operation of holding over by using synce is not supported. [PR1525918](#)
- On the ACX710 router, T1 or T4 cTE should be tuned closer to two-way CTE. [PR1527347](#)
- Microsemi servo doesn't lose lock immediately on loss of one way delay data. Depending on the quality of time delay data available, the lock state can continue. [PR1528973](#)
- Virtual Port and T-GM is not supported in Junos OS Release 20.3R1. Only G.8275.1 T-BC is supported in Junos OS Release 20.3R2. [PR1533018](#)
- Making CoS related changes when traffic is flowing such as deactivate/activate, chassis control restart could lead to not working CoS configuration as expected. It is advisable to stop traffic and then change. [PR1538934](#)
- The issue is only seen when chassisd debug trace is enabled. [PR1539366](#)
- Though enhanced-ip is active, observed alarm REO network-service mode mismatch between configuration and kernel setting during ISSU. [PR1546002](#)
- For ISIS and STATIC routes, ARP resolution of next hop is needed for BFD to start. Workaround is to ping or have any traffic towards the next hop. [PR1561590](#)
- As per the current code, ACX would not delete a mac address from the mac table, there is - (a) traffic destined to the mac address or (b) traffic sourced from the mac address or (c) both Fix of this PR will allow ACX to only look at (b) traffic sourced from mac address before deleting the mac address entry from mac table. So, if there is no traffic sourced from the mac for an interval of mac aging timer, the mac would be deleted from the mac table at the end of mac aging timer with out taking into account the traffic destined to the mac address.. [PR1565642](#)
- The Precision Time Protocol (PTP) clock might fail to be locking and stuck in acquiring state at clock servo. [PR1570310](#)
- On ACX710, if **drop-profiles** is configured for scheduler-map which is then applied to interfaces, the FPC might crash hence there is traffic loss. [PR1588860](#)

Interfaces and Chassis

- The mc-ae option need to have **prefer-status-control-active** set to avoid flap on the split brain case (ICCP down or peer node reboot). Configure it on the status-control active PE device. [PR1505841](#)

Network Management and Monitoring

- On all Junos platforms, the SNMP polling might not be working if the IS-IS protocol is disabled under the same VRF (Virtual Routing and Forwarding) through which SNMP requests are sent. [PR1527251](#)

Virtual Chassis

- In the ACX5000 router, the following false positive parity error messages are observed: `_soc_mem_array_sbusdma_read` and might raise false alarm. [PR1276970](#)

SEE ALSO

[What's New | 14](#)

[What's Changed | 19](#)

[Known Limitations | 21](#)

[Resolved Issues | 26](#)

[Documentation Updates | 33](#)

[Migration, Upgrade, and Downgrade Instructions | 34](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 27](#)
- [Resolved Issues: 20.3R2 | 28](#)
- [Resolved Issues: 20.3R1 | 30](#)

This section lists the issues fixed in Junos OS Release 20.3R3 for the ACX 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: 20.3R3

General Routing

- The IPv6 BFD sessions flap when configured below 100 ms flaps. [PR1456237](#)
- Packets dropped might be seen after configuring PTP transparent clock. [PR1530862](#)
- The ACX710 **jnpr-clock-recovery.log** exhausts storage and affects swap memory operation. [PR1542018](#)
- On the ACX710 platforms, traffic received under a VRF might be discarded if it is resolved by the default route under the VRF for forwarding. [PR1551063](#)
- Verifying multiple PD synchronizations with relay deletes and adds configurations. [PR1554647](#)
- On the ACX5448 routers, single-rate-three-color policer does not work. [PR1559665](#)
- On ACX5048 router, the fxpc process generates core file on the analyzer configuration. [PR1559690](#)
- On ACX platforms, ODIN NPI - MC-LAG:- When LACP daemon is restarted LACP local partner system id remains 0 in mc-ae output. [PR1560820](#)
- Inline BFD stays down with ISIS and Static clients. [PR1561590](#)
- On ACX5448 router, the syslog message **ACX_DFW_CFG_FAILED: ACX Error (dfw):dnx_dfw_dyn_entry_counter_get : Entry is invalid** is reported every 30 seconds. [PR1562323](#)
- Analyzer (Port Mirroring) might not work on ports above 20. [PR1563774](#)
- On ACX710 and ACX5448 platforms, loopback0 firewall might not take effect along with error logs. [PR1566417](#)
- LLDP does not work on the management interface of ACX5448, ACX5448-D, and ACX5448-M routers. [PR1566454](#)
- On ACX710, the log file of the lcklsyncd process shows empty. [PR1567687](#)
- The DF (Designated Forwarder) might not forward traffic. [PR1567752](#)
- ACX resets tunable optics to default wavelength after upgrade or reboot. [PR1570192](#)
- On the ACX5448 routers, the untagged traffic is being incorrectly queued and marked. [PR1570899](#)
- On the ACX710, PEM Feed snmp trap support. [PR1571368](#)
- On ACX5448, RFC2544 reflector feature could not work on a higher port. [PR1571975](#)
- The Layer 2 circuit and CFM sessions might go down with asynchronous-notification configured. [PR1572722](#)
- ARP traffic exceeding the policer limit is not discarded on ACX platforms. [PR1573956](#)

- Packets might get tagged with the default VLAN-ID and dropped at the peer under Layer 2 circuits local switching scenario. [PR1574623](#)
- ACX as a LSR router fails to process RSVP Path Message. [PR1576585](#)
- On the ACX710 routers, continuous reboot due to configuration under auxiliary ports observed. [PR1580016](#)
- There might be a traffic drop between customer edge and provider edge devices in case of ARP resolution failure. [PR1580782](#)
- The process rpd might get stuck in 100% due to race condition. [PR1582226](#)
- ACX710 log jnpr-clock-recovery.log file size too small and archives rotate too quick. [PR1582350](#)
- On the ACX710, unexpected results observed while verifying channelised interface check with `snmp mib get ifHighSpeed` output. [PR1583995](#)
- The IPv4 traffic loss with packet size more than 1410 on ACX5448. [PR1584509](#)
- On the ACX5448, **ACX_ASIC_PROGRAMMING_ERROR** detection time shows the default value (6.000) instead of the configured value for single hop BFD. [PR1585382](#)
- Traffic not passing through CCC interface when vlan-id-range is configured. [PR1590969](#)

Class of Service (CoS)

- FPC crash might be observed after the `show class-of-service` command. [PR1568661](#)

Infrastructure

- The vme/me0 management interface cannot process any incoming packets. [PR1552952](#)

Platform and Infrastructure

- In Junos OS, upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart (vmcore) (CVE-2021-0283, CVE-2021-0284). [PR1557881](#)

Resolved Issues: 20.3R2

General Routing

- On the ACX5048 router, the egress queue statistics do not work for the aggregated Ethernet interfaces. [PR1472467](#)
- On the ACX710 router, high convergence is observed with the EVPN-ELAN service in a scaled scenario during FRR switchover. [PR1497251](#)
- EXP rewrite might cause incorrect EXP marking of traffic on ACX5448 and ACX710 routers. [PR1500928](#)
- ACX1100, ACX2100, ACX2200, ACX2000, and ACX4000 routers might stop forwarding transit and control traffic. [PR1508534](#)
- Transit DHCP packets drop is seen on ACX5448 routers [PR1517420](#)

- Tagged traffic matching the **vlan-id-list** under **vlan-ccc** configuration might get dropped in the ingress interface. [PR1519568](#)
- On the ACX500-I router, the show services session count does not work as expected. [PR1520305](#)
- Interface does not come up with the auto-negotiation setting between the ACX1100 router and the other ACX Series routers, MX Series routers, and QFX Series switches as the other end. [PR1523418](#)
- [Cos] [Hqos] ACX-710 :: Hqos= PIR/CIR Hqos behavior is inconsistent. [PR1525789](#)
- The aggregated Ethernet interface might not come up with Link Fault Management configured after reboot [PR1526283](#)
- On the ACX5448 router with 1000 CFM, the CCM state does not go in the Ok state after loading the configuration or restarting the Packet Forwarding Engine. [PR1526626](#)
- On ACX5448 and ACX710 routers, the VLAN-ID-list statement might not work as expected. [PR1527085](#)
- VPLS traffic loss might be observed on the ACX5448 or ACX710 routers [PR1527231](#)
- The l2cpd memory leak could be observed with aggregated Ethernet interface flap. [PR1527853](#)
- FEC field is not displayed when the interface is down. [PR1530755](#)
- Packets dropped might be seen after configuring PTP transparent clock. [PR1530862](#)
- The show class-of-service routing-instance command does not show the configured classifier. [PR1531413](#)
- Memory leak in Local OutLif in VPLS/CCC topology might be observed. [PR1532995](#)
- The clksyncd process generates core file. [PR1537107](#)
- The rpd process generates core file at l2ckt_vc_adv_rcv, l2ckt_adv_rt_flash (taskptr=0x4363b80, rtt=0x4418100, rtl=< optimized out>, data=< optimized out>, opcode=< optimized out>) at `../..../src/junos/usr.sbin/rpd/l2vpn/l2ckt.c:7982`. [PR1537546](#)
- Management Ethernet link down alarm is seen while verifying system alarms in a Virtual Chassis setup. [PR1538674](#)
- On the ACX5448 router, unexpected behavior of the show chassis network-services command is observed. [PR1538869](#)
- The following error message is observed while deleting the remote stream 0 0 0 0 0 0 along with feb core file at 0x00ae6484 in bcmdnx_queue_assert (queue=0xc599b60) at `../..../src/pfe/common/drivers/bcmdnx/bcmdnx_sdk_ukern_layer.c: Err] clksync_mimic_delete_clock_entry Unexpected error`. [PR1539953](#)
- The announcement or synchronization interval rate range is not as expected. [PR1542516](#)
- The ACX5448 router as transit for the BGP labeled unicast drops traffic. [PR1547713](#)
- The ARP packets from the CE device are added with VLAN tag if the VLAN-ID is configured in the EVPN routing instance. [PR1555679](#)
- On ACX5448 router, you cannot downgrade to Junos OS Release 18.4 code base. [PR1556377](#)

- On ACX5448 router, the unicast packets from the CE devices might be forwarded by the PE devices with additional VLAN tag if IRB is used. [PR1559084](#)
- On ACX5048 router, the fxpc process generates core file on the analyzer configuration. [PR1559690](#)
- On ACX5448 router, the syslog message **ACX_DFW_CFG_FAILED: ACX Error (dfw):dnx_dfw_dyn_entry_counter_get : Entry is invalid** is reported every 30 seconds. [PR1562323](#)

Class of Service

- The explicit classifier or rewrite-rule might not work as expected for a logical interface if the wildcard configuration is also applied. [PR1556103](#)

Forwarding and Sampling

- VLAN ID-based firewall match conditions might not work for the VPLS service. [PR1542092](#)

Layer 2 Features

- On ACX5448, VPLS traffic statistics information is not displayed when executing the **show vpls statistics** command. [PR1506981](#)

Resolved Issues: 20.3R1

General Routing

- Policer discarded count is shown incorrectly to the enq count of the interface queue, but the traffic behavior is as expected. [PR1414887](#)
- The **gether-options** command is enabled again under the interface hierarchy. [PR1430009](#)
- The statistics are accessed through Broadcom API, which is the same for both tagged and untagged packets. This cannot be changed in accordance with the MX Series routers since it is directly accessed from Broadcom without any statistics changes specific to tagging from the ACX5448 router side. This impacts other statistics if the change is made. [PR1430108](#)
- While performing repeated power-off or power-on of the device, the SMBUS transactions timeout occurs. [PR1463745](#)
- Unable to get shared buffer count as expected. [PR1468618](#)
- The router might become nonresponsive and bring the traffic down when the disk space becomes full. [PR1470217](#)
- On the ACX5048 router, the egress queue statistics do not work for the aggregated Ethernet interfaces. [PR1472467](#)
- The links might not come up when the 100-Gigabit Ethernet interface is channelized into the four 25-Gigabit Ethernet interfaces. [PR1479733](#)
- On the ACX6360 router, the disk usage might keep increasing. [PR1480217](#)
- Memory utilization enhancement is needed. [PR1481151](#)

- ACX AUTHD process memory usage enhancement is needed. [PR1482598](#)
- BFD over Layer 2 VPN or Layer 2 circuit does not work because of the SDK upgrade to version 6.5.16. [PR1483014](#)
- On the ACX5048 router, traffic loss is observed during the unified ISSU upgrade. [PR1483959](#)
- On the ACX5448 router, the fpc process might crash. [PR1485315](#)
- The LSP might not come up in an LSP externally-provisioned scenario. [PR1494210](#)
- When 40-Gigabit Ethernet or 10-Gigabit Ethernet interface optics are inserted in 100-Gigabit Ethernet or 25-Gigabit Ethernet interface port with 100-Gigabit Ethernet or 25-Gigabit Ethernet interface speed configured and vice versa, the Packet Forwarding Engine log message displays a speed mismatch. [PR1494591](#)
- When 40-Gigabit Ethernet or 10-Gigabit Ethernet interface optics are inserted in 100-Gigabit Ethernet or 25-Gigabit Ethernet interface with 100-Gigabit Ethernet interface speed configured and vice versa, there is a speed mismatch. [PR1494600](#)
- Outbound SSH connection flaps or leaks memory during the push configuration to the ephemeral database with a high rate. [PR1497575](#)
- All the autonegotiation parameters are not shown in the output of the **show interface media** command. [PR1499012](#)
- The hardware FRR for EVPN-VPWS, EVPN-FXC, and Layer 3 VPN with a composite next hop are not supported in Junos OS Release 20.2R1. [PR1499483](#)
- SFP-T is unrecognized on Junos OS Release 20.3DCB after FPGA upgrade and power cycle. [PR1501332](#)
- On the ACX710 router, the BFD sessions are in the **Init** state with CFM scale of 1000 on reboot or chassis-control restart. [PR1503429](#)
- On the ACX500 router, the SFW sessions might not get updated on ms interfaces. [PR1505089](#)
- The wavelength changes from CLI but does not update the hardware for the tunable optics. [PR1506647](#)
- The PIC slot might shut down in less than 240 seconds due to the over temperature start time being handled incorrectly. [PR1506938](#)
- In the PTP environment, some vendor devices acting as slave are expecting announce messages at an interval of -3 (8pps) from the upstream primary device. [PR1507782](#)
- The BFD session flaps with the following error message after a random time interval:
ACX_OAM_CFG_FAILED: ACX Error (oam):dnx_bfd_l3_egress_create : Unable to create egress object.
[PR1513644](#)
- On the ACX710 router, the following error message is observed in the Packet Forwarding Engine while the EVPN core link flaps: **dnx_l2alm_add_mac_table_entry_in_hw**. [PR1515516](#)
- The VM process generates a core file while running stability test in a multidimensional scenario. [PR1515835](#)

- The l2ald process crashes during stability test with traffic on a scaled setup. [PR1517074](#)
- On the ACX710 router, whenever a copper optic interface is disabled and enabled, the speed shows 10 Gbps rather than 1 Gbps. This issue is not seen with the fiber interface. [PR1518111](#)
- The IPv6 neighbor state change causes **Local Outlif** to leak by two values, which leads to the following error: **DNX_NH::dnx_nh_tag_ipv4_hw_install**. [PR1519372](#)
- The **Incompatible Media type** alarm is not raised when the Synchronous Ethernet source is configured over the copper SFP. [PR1519615](#)
- If the client clock candidate is configured with a virtual port, the clock class is on T-BC. [PR1520204](#)
- On the ACX710 router, the alarm port configuration is not cleared after deleting the alarm-port. [PR1520326](#)
- The **show class-of-service interface** command does not show classifier information. [PR1522941](#)
- On the ACX5448 chassis, **mac-address** and **label mac-address** might not match. [PR1489034](#)
- On the ACX5000 router, the IEEE 802.1p priority and Drop Eligibility Indicator values in the locally generated VLAN-based IP packets might be changed when sourced from the IRB interface. [PR1490966](#)
- VPLS flood groups result in IPv4 traffic drop after the core interface flaps. [PR1491261](#)
- On the ACX5048 and ACX5096 routers, the Link Aggregation Control Protocol control packets might be dropped due to high CPU utilization. [PR1493518](#)
- On the ACX710 router, high convergence is observed with the EVPN-ELAN service in a scaled scenario during FRR switchover. [PR1497251](#)
- The loopback filter cannot take more than 2 TCAM slices. [PR1513998](#)
- On the ACX5448 and ACX710 routers, the **vlan-id-list** statement might not work as expected. [PR1527085](#)
- Memory leak is observed in the local OutLif in the VPLS or CCC topology. [PR1532995](#)
- On the ACX710 router, the following error message is observed: **PFE_ERROR_FAIL_OPERATION: Failed to install in h/w, LOG: Err] dnx_nh_unilist_install: BCM L3 Egress create object failed for:Unilist nh 2097369 (0:Ok) nh 0**. [PR1495563](#)
- The following error message is observed during the MPLS route add, change, and delete operation: **mpls_extra NULL**. [PR1502385](#)
- On the ACXR6675 router, the rpd process generates core file at **l2ckt_vc_adv_recv, l2ckt_adv_rt_flash (taskptr=0x4363b80, rtt=0x4418100, rtl=< optimized out>, data=< optimized out>, opcode=< optimized out>)** at **../../../../../../../../src/junos/usr.sbin/rpd/l2vpn/l2ckt.c:7982**. [PR1537546](#)

Interfaces and Chassis

- The fpc process might crash with in an inline mode with CFM configured. [PR1500048](#)

MPLS

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

Routing Protocols

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

VPNs

- The l2circuit neighbor might become nonresponsive in the **RD** state at one end of the MG-LAG peer. [PR1498040](#)
- The rpd process might crash in certain conditions after deleting the l2circuit configuration. [PR1502003](#)

SEE ALSO

What's New 14
What's Changed 19
Known Limitations 21
Open Issues 23
Documentation Updates 33
Migration, Upgrade, and Downgrade Instructions 34

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for ACX Series routers.

SEE ALSO

What's New 14
What's Changed 19
Open Issues 23

[Known Limitations | 21](#)[Resolved Issues | 26](#)[Migration, Upgrade, and Downgrade Instructions | 34](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 34](#)

This section contains the upgrade and downgrade support policy for Junos OS for ACX Series routers. 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

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you

can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and 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 14
What's Changed 19
Known Limitations 21
Open Issues 23
Resolved Issues 26
Documentation Updates 33

Junos OS Release Notes for cSRX

IN THIS SECTION

- [What's New | 36](#)
- [What's Changed | 37](#)
- [Known Limitations | 39](#)
- [Open Issues | 39](#)
- [Resolved Issues | 40](#)

These release notes accompany Junos OS Release 20.3R3 for the cSRX Container Firewall, a containerized version of the SRX Series Services Gateway. 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 20.3R3 | 36](#)
- [What's New in Release 20.3R2 | 36](#)
- [What's New in Release 20.3R1 | 37](#)

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

What's New in Release 20.3R3

There are no new features for cSRX in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features for cSRX in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Installation and Upgrade

- **cSRX orchestration using Kubernetes**—Starting in Junos OS Release 20.3R1, you can deploy cSRX as Kubernetes Service or Pods. With Kubernetes, you can scale out and scale in cSRX in a cluster that provides an elastic firewall service to application containers.

[See [cSRX Installation using Kubernetes](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 37](#)
- [What's Changed in Release 20.3R2 | 37](#)
- [What's Changed in Release 20.3R1 | 38](#)

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

What's Changed in Release 20.3R3

There are no changes in behavior or syntax for cSRX in Junos OS Release 20.3R3.

What's Changed in Release 20.3R2

There are no changes in behavior or syntax for cSRX in Junos OS Release 20.3R2.

What's Changed in Release 20.3R1

Download Juniper Signature Pack on cSRX

- **Download of Juniper Signature Pack on cSRX—**

You can download the signature pack through a proxy server. The AppIDD and IDPD process first connect to the configured proxy server. The proxy server then communicates with the signature pack download server and provides the response to the process running on the device.

You can download the signature pack from the [Juniper Signature Repository](#) directly when the cSRX doesn't have the preinstalled signature pack.

1. To download signature pack from [Juniper Signature Repository](#):

```
root@host> request services application-identification download
```

```
root@host> request security idp security-package download
```

To download the signature pack through the proxy server:

1. Configure the proxy server so that the IP address of the proxy server is reachable from cSRX.
2. Run the following command to enter configuration mode from CLI.

```
root@host> configure
```

```
Entering configuration mode
```

```
[edit]
```

```
root@host#
```

3. Configure the proxy server profile on cSRX using the IP address and port of the proxy server.

```
root@host# set services proxy profile appid_sigpack_proxy protocol http host 4.0.0.1
```

```
root@host# set services proxy profile appid_sigpack_proxy protocol http port 3128
```

4. Attach the profile to AppID and IDP.

```
root@host# set services application-identification download proxy-profile appid_sigpack_proxy
```

```
root@host# set security idp security-package proxy-profile appid_sigpack_proxy
```

5. Commit the configuration.

```
root@host# commit and-quit
```

```
commit complete
Exiting configuration mode
```

6. Download the IDP and AppID signature pack through the proxy server.

```
root@host> request services application-identification download
```

```
root@host>request security idp security-package download
```

To verify that the download is happening through the proxy server, check the logs in proxy server.

```
[root@srxdpi-lnx39 squid]# cat /var/log/squid/access.log
```

```
1593697174.470    1168 4.0.0.254 TCP_TUNNEL/200 5994 CONNECT
signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
1593697175.704    1225 4.0.0.254 TCP_TUNNEL/200 11125 CONNECT
signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
1593697176.950    1232 4.0.0.254 TCP_TUNNEL/200 5978 CONNECT
signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
1593697178.195    1236 4.0.0.254 TCP_TUNNEL/200 11188 CONNECT
signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
1593697198.337    1243 4.0.0.254 TCP_TUNNEL/200 6125 CONNECT
signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
```

In cSRX, the TLS protocol is used and traffic through the proxy server is encrypted.

Known Limitations

There are no known behavior or limitation for cSRX in Junos OS Release 20.3R3.

Open Issues

There are no known issues for cSRX in Junos OS Release 20.3R3.

Resolved Issues

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

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

Resolved Issues: 20.3R3

There are no resolved issues for cSRX in Junos OS Release 20.3R3.

Resolved Issues: 20.3R2

There are no resolved issues for cSRX in Junos OS Release 20.3R2.

Resolved Issues: 20.3R1

There are no resolved issues for cSRX in Junos OS Release 20.3R1.

Junos OS Release Notes for EX Series

IN THIS SECTION

- [What's New | 41](#)
- [What's Changed | 53](#)
- [Known Limitations | 56](#)
- [Open Issues | 58](#)
- [Resolved Issues | 63](#)
- [Documentation Updates | 72](#)
- [Migration, Upgrade, and Downgrade Instructions | 72](#)

These release notes accompany Junos OS Release 20.3R3 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 20.3R3](#) | 41
- [What's New in Release 20.3R2](#) | 41
- [What's New in Release 20.3R1](#) | 41

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

NOTE: The following EX Series switches are supported in Release 20.3R3: EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

What's New in Release 20.3R3

There are no new features or enhancements to existing features for EX Series Switches in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for EX Series Switches in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Hardware

- **Support for the QSFP-4X10GE-SR and JNP-QSFP-4X10GE-LR transceivers (EX4650)**—Starting in Junos OS Release 20.3R1, EX4650 switches support the QSFP-4X10GE-SR and JNP-QSFP-4X10GE-LR transceivers.

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

- **New EX9200-15C fixed-configuration line card and EX9200-SF3 switch fabric module (EX9204, EX9208, and EX9214)**—In Junos OS Release 20.3R1, we introduce the EX9200-15C line card. The EX9200-15C is supported on EX9204, EX9208, and EX9214 switches. The EX9200-15C supports the following:
 - Line-rate throughput of up to 1.5 Tbps
 - Fifteen network ports that can be configured for 100-Gbps, 40-Gbps, 25-Gbps, or 10-Gbps speeds (breakout cables are used for 25-Gbps and 10-Gbps speeds)

NOTE: For the EX9200-15C line card to be operational, you must install the EX9200-SF3 Switch Fabric module (SF module) in the switch. [See [EX9200 Line Cards](#).]

The EX9200-SF3 is an enhanced Switch Fabric module supported on EX9204, EX9208, and EX9214 switches. The EX9200-SF3 supports a pluggable Routing Engine and provides a control plane and data plane interconnect to each line card slot. In a redundant configuration, the EX9200-SF3 provides fabric bandwidth of up to 1 Tbps per slot. In a non-redundant configuration, the EX9200-SF3 provides fabric bandwidth of up to 1 Tbps per slot (four fabric planes) and 1.5 Tbps per slot fabric bandwidth when all six fabric planes are used (with EX9200-15C line cards).

The following Routing Engines are supported on the EX9200-SF3: EX9200-RE2 and EX9200-RE. The EX9200-SF3 interoperates with the following existing line cards: EX9200-MPC, EX9200-12QS, EX9200-32XS, and EX9200-40XS. The EX9200-SF3 does not interoperate with any previous generation Switch Fabric modules (EX9200-SF or EX9200-SF2). The EX9200-SF3 does not interoperate with the following line cards: EX9200-2C-8XS, EX9200-4QS, EX9200-6QS, and EX9200-40 1-Gigabit line cards (EX9200-40T, EX9200-40F, and EX9200-40F-M). For the EX9200-15C line card to be operational, you must install the EX9200-SF3 Switch Fabric module (SF module) in the switch. [See [EX9200 Host Subsystem](#).]

To install the EX9200 line card and perform initial software configuration, routine maintenance, and troubleshooting, see [EX9204 Switch Hardware Guide](#), [EX9208 Switch Hardware Guide](#), and [EX9214 Switch Hardware Guide](#).

[Table 2 on page 42](#) summarizes the EX9200-15C features supported in Junos OS Release 20.3R1.

Table 2: Features Supported by the EX9200-15C

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> • Support for COS features (classifiers, rewrites, port queuing, and L3) on EX9204, EX9208, and EX9214 switches. [See Understanding Junos OS CoS Components for EX Series Switches.]

Table 2: Features Supported by the EX9200-15C (*continued*)

Feature	Description
EVPN	<ul style="list-style-type: none"> • Support for EVPN-MPLS singlehoming. This feature supports single-homed devices on an EVPN-MPLS network. [See Introduction to EVPN Multihoming.] • Support for NDP and Proxy ARP. Junos OS supports proxy Address Resolution Protocol (ARP) and Network Discovery Protocol (NDP).
Firewalls and policers	<ul style="list-style-type: none"> • Support for CCC and Layer 3 firewall forwarding. [See CCC Overview.] • Support for advanced Layer 2 features: <ul style="list-style-type: none"> • Firewall filters for Layer 2 and MAC filters. [See Layer 2 Forwarding Tables.] • Layer 2 firewall forwarding support. A firewall filter specifies required traffic and directs it to the mirror. [See Understanding Port Mirroring and Analyzers.] • Support for firewall forwarding. The following traffic policers are fully supported: GRE tunnels, including encapsulation (family any), de-encapsulation, GRE-in-UDP over IPv6, and the following sub-options: sample, forwarding class, interface group, and no-ttl-decrement. <ul style="list-style-type: none"> • Input and output filter chains • Actions, including policy-map filters, do-not-fragment, and prefix • Layer 2 policers • Policer overhead adjustment • Hierarchical policers • Shared bandwidth • Percentages • Logical interfaces <p>[See Traffic Policer Types.]</p>

Table 2: Features Supported by the EX9200-15C (continued)

Feature	Description
Junos telemetry interface	<ul style="list-style-type: none"> • JTI for FPC and optics support. Junos telemetry interface (JTI) supports streaming of Flexible PIC Concentrator (FPC) and optics statistics for the router using remote procedure calls (gRPC). gRPC is a protocol for configuration and retrieval of state information. The following base resource paths are supported: <ul style="list-style-type: none"> • <code>/junos/system/cmerror/configuration/</code> • <code>/junos/system/cmerror/counters/</code> • <code>/junos/system/linecard/environment/</code> • <code>/junos/system/linecard/optics/</code> • <code>/junos/system/linecard/optics/optics-diag[if-name =]</code> • <code>/junos/system/linecard/optics/optics-diag/if-name</code> • <code>/junos/system/linecard/optics/optics-diag/snmp-if-index</code> • <code>/junos/system/linecard/optics/lane[lane_number=]/</code> <p>[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]</p>
Layer 2 features	<ul style="list-style-type: none"> • Support for Layer 3 VPN, Layer 2 VPN and Layer 2 circuits. [See Introduction to Configuring Layer 3 VPNs, Layer 2 VPNs Configuration Overview and Layer 2 Circuits Configuration Overview.] • Support for Layer 2 forwarding services. This includes support for the following features: Layer 2 bridge and MAC learning, trunk port, and mesh groups. [See Understanding Layer 2 Bridge Domains and Learning and Forwarding.] • Support for IRB, VLAN handling, and Q-in-Q tunneling. [See Integrated Routing and Bridging, Understanding Bridging and VLANs on Switches and Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation.] • Support for VPLS. [See Introduction to Configuring VPLS.]

Table 2: Features Supported by the EX9200-15C (*continued*)

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> • Support for Layer 3 forwarding. Junos OS supports the following Layer 3 features on the EX9200-15C: <ul style="list-style-type: none"> • BGP (Multipath/v4-v6 labelled unicast) • Bidirectional Forwarding Detection (excluding micro BFD and BFD sessions with authentication) • IPv4 (forwarding and options) • IPv6 (forwarding and route accounting) • Load balancing (ECMP and FRR) • L2VPN, CCC, and L2 Circuit • MPLS (Push/Pop/Swap, LDP, RSVP-Aggregate, RSVP TE Admin Groups, RSVP-TE, OAM, LSP/VPN ping, Trace Route, Auto Bandwidth, and MPLS-FRR Link node protection. • OSPF (node-link-protection and node-link-degradation) • Protocols (ISIS, OSPF, OSPF V3 for V6, BGP + BGP-v6, BGP LU, BGP-LS, BGP optimal-route-reflection (ORR), BFD (Centralized), Micro BFD (Centralized), ICMP and ICMPv6 error handling, and LLDP) • Routing Instance Logical System VRF • Tunnel (Generic Routing Encapsulation (GRE), Logical Tunnel (LT), and Virtual Tunnel (VT))

Table 2: Features Supported by the EX9200-15C (*continued*)

Feature	Description
MPLS	<ul style="list-style-type: none"> • Support for static LSP and LDP features. The MPLS features supported are: <ul style="list-style-type: none"> • Keepalive support for GRE interfaces • LDP downstream on demand • Static, RSVP, and LDP LSPs • Layer 2 Circuit and Layer 2 VPN with or without control word • Layer 3 VPN with chain-composite-nexthop • Layer 3 VPN with vrf-table-label • MPLS link protection, node protection, and FRR • P2MP LSP traceroute • Statistics for P2MP LSPs • LSPs: statistics, ping and traceroute, TTL knobs (no-propagate-ttl and no-decrement-ttl), and point-to-multipoint LSP support for multicast VPNs. • Static LSPs: revert timer, statistics, traceoptions, support for bypass of static LSPs, support at the ingress device, and support at the transit device. <p>[See MPLS Applications User Guide.]</p>
Multicast	<ul style="list-style-type: none"> • Support for Multicast forwarding including PIM, IGMP, and MLD. [See Multicast Overview.]

Table 2: Features Supported by the EX9200-15C (continued)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> • Port mirroring support for families inet, inet6, and ethernet-switching, configured at the [edit forwarding-options port-mirroring] hierarchy level. [See Understanding Port Mirroring and Analyzers.] • Support for link fault management (LFM). You can configure IEEE 802.3ah link fault management on EX9200-15C switches. You can configure OAM LFM on point-to-point Ethernet links that are connected directly or through Ethernet repeaters, and on aggregated Ethernet interfaces. The LFM status of individual links determines the LFM status of the aggregated Ethernet interface. You can also configure the following supported LFM features: <ul style="list-style-type: none"> • Discovery and link monitoring • Distributed LFM • Remote fault detection and remote loopback [See OAM Link Fault Management.] • Support for Junos OS management and software features on the EX9200-15C: <ul style="list-style-type: none"> • Chef, Puppet, SYSLOG, Authentication, authorization, and accounting (AAA), Stylesheet Language Alternative syntaX (SLAX), SNMP, COMMIT, User Interface, Management process or daemon (MGD) Infrastructure, NETCONF, JUNOScript, Google Network Management Interface (gNMI) for Junos Telemetry Interface, YANG, and JET APIs • Support for hyper mode and non hyper mode features. [See Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches.]
Port security	<ul style="list-style-type: none"> • Provides support for MACsec on ports at these speeds 10G, 25G, 40G, and 100G. [See Understanding Media Access Control Security (MACsec).]
Services applications	<ul style="list-style-type: none"> • 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, EX9200 Switches, and NFX250 to Use Version 9 Flow Templates.]
System management	<ul style="list-style-type: none"> • Support for the Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs). [See show chassis hardware.]

To view the hardware compatibility matrix for optical interfaces, transceivers, and DACs supported across all platforms, see the [Hardware Compatibility Tool](#).

Class of Service (CoS)

- **CoS support on EVPN VXLAN (EX4300 Multigigabit)**—Starting with Junos OS Release 20.3R1, EX4300 Multigigabit switches support defining classifiers and rewrite rules on leaf (initiation and terminations) and spine nodes for EXPN VXLANs.

[See [CoS Support on EVPN VXLANs.](#)]

EVPN

- **Color-based mapping of EVPN-MPLS and EVPN services over SR-TE (ACX5448, EX9200, MX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can specify a color attribute along with an IP protocol next hop. The color attribute adds another dimension to the resolution of transport tunnels over static colored and BGP segment routing traffic-engineered (SR-TE) label-switched paths (LSPs). This type of resolution is known as the color-IP protocol next-hop resolution. With the color-IP protocol next-hop resolution, you must configure a resolution map and apply it to EVPN-MPLS and EVPN services, which includes E-Line, E-LAN and E-Tree. With this feature, you can enable color-based traffic steering of EVPN-MPLS and EVPN services.

[See [Segment Routing LSP Configuration.](#)]

Junos OS XML, API, and Scripting

- **Support for REST API over nondefault virtual routing and forwarding (VRF) instance (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can execute Junos OS operational commands using the REST API over a nondefault VRF instance. The nondefault VRF instance can be a user-defined instance or the management instance, `mgmt_junos`.

The REST API allows you to execute Junos OS operational commands over HTTP(S). If you don't specify a routing instance, REST API requests are sent over the default routing instance. Use a nondefault VRF instance to improve security and make it easier to troubleshoot.

Use the **routing-instance *routing-instance*** statement at the **[edit system services rest]** hierarchy level to specify a nondefault VRF instance for REST API requests.

[See [Management Interface in a Nondefault Instance](#) and [rest](#).]

Junos Telemetry Interface

- **EVPN statistics export using JTI (MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMXrouters, EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253 switches)**—Starting in Junos OS Release 20.3R1, you can use Junos telemetry interface (JTI) an remote procedure call (gRPC) services to export EVPN statistics from devices to an outside collector.

Use the following sensors to export EVPN statistics:

- Sensor for instance level statistics (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/`)
- Sensor for route statistics per peer (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/peer/`)
- Sensor for Ethernet segment information (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/ethernet-segment/`). This includes EVPN designated forwarder ON_CHANGE leafs `esi` and `designated-forwarder`.
- Sensor for local interface information (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/interfaces/`)
- Sensor for local IRB interface information (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/irb-interfaces/`)
- Sensor for global resource counters and current usage (resource path `/junos/evpn/evpn-smet-forwarding/`)
- Sensor for EVPN IP prefix (resource path `/junos/evpn/l3-context/`)
- Sensor for EVPN IGMP snooping database (type 6) (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/sg-db/`)
- Sensor for EVPN IGMP join sync (type 7) ad leave sync (type 8) (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/sg-db/sgdb-esi`)
- Sensor to relate selected replicator on AR leaf on QFX5100, QFX5110, QFX5120, and QFX5200 switches (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/assisted-replication/`)
- Sensor for EVPN ON_CHANGE notifications (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/ethernet-segment`)
- Sensor for overlay VX-LAN tunnel information (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/vxlan-tunnel-end-point/`). This includes VTEP information ON_CHANGE leafs `source_ip_address`, `remote_ip_address`, `status`, `mode`, `nexthop-index`, `event-type` and `source-interface`.
- EVPN MAC table information (resource path `/network-instances/network-instance[instance-name='name']/mac_db/entries/entry/`)
- Sensor for MAC-IP or ARP-ND table (resource path `/network-instances/network-instance[instance-name='name']/macip_db/entries/entry/`)
- Sensor for MAC-IP ON_CHANGE table information (resource path `/network-instances/network-instance[name='name']/macip-table-info/`). Statistics include leafs `learning`, `aging-time`, `table-size`, `proxy-macip`, and `num-local-entries`.

- Sensor for MAC-IP ON_CHANGE entry information (resource path `/network-instances/network-instance[name='name']/macip-table/entries/entry/`). Statistics include leafs `ip-address`, `mac-address`, `vlan-id` and `vni`.
- Sensor for bridge domain or VLAN information (resource path `/network-instances/network-instance[instance-name='name']/bd/`)

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

MPLS

- **Support for static LSP and LDP features (EX9200)**—Starting in Junos OS Release 20.3R1, the following MPLS features are supported:
 - Keepalive support for GRE interfaces
 - LDP downstream on demand
 - Static, RSVP, and LDP LSPs
 - Layer 2 Circuit and Layer 2 VPN with or without control word
 - Layer 3 VPN with chain-composite-nexthop
 - Layer 3 VPN with vrf-table-label
 - MPLS link protection, node protection and FRR
 - P2MP LSP traceroute
 - Statistics for P2MP LSPs
 - LSPs:
 - Statistics
 - Ping and traceroute
 - TTL knobs: `no-propagate-ttl` and `no-decrement-ttl`
 - Point-to-multipoint LSP support for multicast VPNs
 - Static LSPs:
 - Revert timer
 - Statistics
 - Traceoptions
 - Support for bypass of static LSPs

- Support at the ingress device
- Support at the transit device

Network Management and Monitoring

- **Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can use the **probe** command to query the status of the probed interface. The proxy interface resides on the same node as the probed interface, or it can reside on a node to which the probed interface is directly connected.

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

To enable the **probe** command, configure the **extended-echo** statement under the **[edit system]** hierarchy.

[See [Using the Probe command.](#)]

- **Enhancements to sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.3R1, devices running Junos OS with upgraded FreeBSD support the following enhancements to sessions over outbound HTTPS:
 - Connecting to multiple outbound HTTPS clients by configuring one or more clients at the **[edit system services outbound-https]** hierarchy level
 - Configuring multiple backup gRPC servers for a given outbound HTTPS client
 - Establishing a csh session
 - Establishing multiple, concurrent NETCONF and csh sessions between the device running Junos OS and an outbound HTTPS client
 - Configuring a shared secret that the outbound HTTPS client uses to authenticate the device running Junos OS
 - Authenticating the client using certificate chains in addition to self-signed certificates

[See [NETCONF and Shell Sessions over Outbound HTTPS.](#)]

open-config

- **OpenConfig support for Routing Policy (EX4300, EX4600, and EX9200 switches)**—Junos OS Release 20.3R1 adds support for OpenConfig Data Model Version v2.0.1, supporting all configurations at `/routing-policy/`.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [OpenConfig Data Model Version v2.0.1.](#)]

Routing Policy and Firewall Filters

- **Loopback firewall filter scale optimization (EX4650 and QFX5120-48Y)**—Starting with Junos OS Release 20.3R1, you can configure up to 768 loopback filter terms for IPv6, and up to 1152 terms for IPv4. To do so, you configure an ingress firewall filter, apply it to the loopback interface, and then enable the **loopback-firewall-optimization** statement at the `[edit chassis]` hierarchy level (this triggers the Packet Forwarding Engine to restart).

The switches do not support terms that include a reserved multicast destination, for example 224.0.0.x/24, and terms with a time-to-live (TTL) of 0/1. You need to configure a separate filter for these terms. So, for example, to count OSPF packets on the loopback interface, you would create a separate filter with terms for the protocol (OSPF) to count packets destined to a reserved multicast address (such as 224.0.0.6).

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

Software Installation and Upgrade

- **Support for phone-home client (EX4300 Virtual Chassis)**—Starting in Junos OS Release 20.3R1, the phone-home client (PHC) can securely provision a Virtual Chassis without requiring user interaction. You only need to:
 - Ensure that the Virtual Chassis members have the factory-default configuration.
 - Interconnect the member switches using dedicated or default-configured Virtual Chassis ports.
 - Connect the Virtual Chassis management port or any network port to the network.
 - Power on the Virtual Chassis members.

PHC automatically starts up on the Virtual Chassis and connects to the phone-home server (PHS). The PHS responds with bootstrapping information, including the Virtual Chassis topology, software image, and configuration. PHC upgrades each Virtual Chassis member with the new image and applies the configuration, and the Virtual Chassis is ready to go.

[See [Provision a Virtual Chassis Using the Phone-Home Client.](#)]

SEE ALSO

Known Limitations	 56
Open Issues	 58
Resolved Issues	 63
Documentation Updates	 72
Migration, Upgrade, and Downgrade Instructions	 72

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3](#) | [54](#)
- [What's Changed in Release 20.3R2](#) | [54](#)
- [What's Changed in Release 20.3R1](#) | [55](#)

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

What's Changed in Release 20.3R3

Junos OS XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or **set refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

Software Licensing

- **License key format (QFX5120-32C, QFX5120-48Y, and QFX5200)**—When you are upgrading from Junos OS release 20.3R1 to Junos OS release 20.3R2 or later releases, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS releases 20.3R2 or later.

What's Changed in Release 20.3R2

Junos XML API and Scripting

- **The jcs:invoke() function supports suppressing root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- **The jcs:invoke() function supports suppressing root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

MPLS

- **The `show mpls lsp extensive` and `show mpls lsp detail` commands display next-hop gateway LSPid**—When you use the `show mpls lsp extensive` and `show mpls lsp detail` commands, you'll see next-hop gateway LSPid in the output.

Platform and Infrastructure

- **Configure internal IPsec authentication algorithm (EX Series)**—You can configure the algorithm `hmac-sha-256-128` at the `[edit security ipsec internal security-association manual direction bidirectional authentication algorithm]` hierarchy level for internal IP security (IPsec) authentication. Earlier to this release, you can configure the algorithm `hmac-sha-256-128` for MX series devices only.

Software Licensing

- **License key format (QFX5120-32C, QFX5120-48Y, and QFX5200)**—When you are upgrading from Junos OS release 20.3R1 to Junos OS release 20.3R2 or later releases, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS releases 20.3R2 or later.

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the `verbose` statement at the `edit system export-format json` hierarchy level. The default format to export configuration data in JSON changed from `verbose` format to `ietf` format starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `edit system export-format json` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See [export-format..](#)]

What's Changed in Release 20.3R1

Class of Service (CoS)

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

Junos OS, XML, API, and Scripting

- **Changes to Junos XML RPC request tag names (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've updated the Junos XML request tag name for some operational RPCs to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include:

- Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
- Uppercase characters are converted to lowercase.

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

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.

Subscriber Management and Services

- **Command to view summary information for resource monitor (EX9200 line of Ethernet switches and MX Series routers)**—The `show system resource-monitor` command enables you to view many statistics about the use of memory resources for all line cards or for a specific line card in the device. It 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](#).

SEE ALSO

What's New 41
Known Limitations 56
Open Issues 58
Resolved Issues 63
Documentation Updates 72
Migration, Upgrade, and Downgrade Instructions 72

Known Limitations

IN THIS SECTION

- [General Routing | 57](#)
- [EVPN | 57](#)

- [Infrastructure | 57](#)
- [Platform and Infrastructure | 57](#)

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.

General Routing

- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. Device can be recovered using power-cycle of the device. [PR1385970](#)

EVPN

- Partial traffic down with breaking links between the leaf are observed. [PR1480847](#)

Infrastructure

- On an EX4300-MP device, 9000 IPv6 MC routes can be installed. If more IPv6 MC routes are added, error messages are displayed. [PR1493671](#)
- Traffic load balancing with static ECMP hashing are observed. [PR1516883](#)

Platform and Infrastructure

- On the EX9208 device, the status of the channels are displayed as up even though the peer end is down with different speed being configured. The LED light also turns green in color. [PR1530061](#)
- On the EX9208 device, the interface does not come up with the DAC BO cables. [PR1530465](#)
- On the EX9208 device, the LED behavior are not consistent across AOC, DAC, LX4, and 4x10G IR when the port is in the **admin-down** state. [PR1532930](#)
- On the EX4300 device, complete traffic drop is observed when the MSTP edge port is configured over the access and QinQ ports. [PR1532992](#)

SEE ALSO

[What's New | 41](#)

[What's Changed | 53](#)[Open Issues | 58](#)[Resolved Issues | 63](#)[Documentation Updates | 72](#)[Migration, Upgrade, and Downgrade Instructions | 72](#)

Open Issues

IN THIS SECTION

- [General Routing | 58](#)
- [Authentication and Access Control | 60](#)
- [EVPN | 60](#)
- [Forwarding and Sampling | 61](#)
- [Infrastructure | 61](#)
- [Layer 2 Features | 61](#)
- [Network Management and Monitoring | 62](#)
- [Platform and Infrastructure | 62](#)
- [User Interface and Configuration | 62](#)

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

General Routing

- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)
- 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](#)
- On BCM Packet Forwarding Engine based EX platforms frame higher than MTU+4 and lesser then MTU+8 bytes, with invalid FCS/code error/IEEE length check error, is treated as Jabber frame. [PR1487709](#)

- On the MPC10 line card, the following error message is observed on the Routing Engine 1 after GRES from the Routing Engine 0 to Routing Engine 1: **user.err aftd-trio: ([Error] L2ALIPC : L2AL IPC client failed to connect to l2ald)**. [PR1491384](#)
- On the EX4300, when running the command, **show pfe filter hw filter-name** the command fails to retrieve the Packet Forwarding Engine programming details of the filter. [PR1495712](#)
- SNMP POE MIB walk produce wither no results or some times result from the master Virtual Chassis whenever one of the Virtual Chassis is renamed. [PR1503985](#)
- When a VLAN member is specified as a string, the 'IF_MSG_IFL_VADDR' TLV is not generated with the VLAN infoformation, and the TRIO afttriostream is not updated with the nativevlanld and nativevlanenable flags. Thus, the packet is are still treated as untagged, and when it reaches the trunk egress interface, it is dropped because the trunk interface does not allow untagged traffic to pass through. The issue is specific to platforms with ZT line cards, including EX9200-SF3 and EX9200-15C. Workaround: Functionally will work if the interface-vlan-members statement contains only numeral value for VLANs. The VLAN members with input as a string is not supported in this release. [PR1506403](#)
- On the EX4300-48MP device, 35 second delay is added in reboot time. [PR1514364](#)
- On EX4650 platform, the phc daemon might crash while committing the **phone-home client** configuration. [PR1522862](#)
- On the Legacy EX series platform, when adding or removing Micro BFD LAG configuration, a kernel crash might happen. The kernel crash might cause unexpected Routing Engine reboot or switchover, and even result in traffic loss until the Routing Engine is restored. [PR1524490](#)
- In rare cases of power related failures on the FPC, Fabric Healing will detect and try to heal this fault condition by performing an offline or online FPC event. If the same FPC fails again within a 10-minute period, fabric auto-healing attempt is considered failing and the FPC will get off-lined to avoid further operational impact. If during the power offline event, the faulty FPC gets disconnected ungracefully due to the hardware power fault, the FPC might attempt an on-lined request again after 5 minutes. There might be traffic impact due to this issue. [PR1556558](#)
- When **dot1x server-fail-voip vlan-name** is configured, ensure that both **server-fail-voip vlan-name** and **voip vlan** are configured using vlan name and not by using vlan-id. [PR1561323](#)
- If a license key has ONLY features that are not applicable on the platform (unknown features), the license key is rejected. If key has one or more platform applicable features (known features) along with unknown features, license key addition is successful with the following warning for the unknown features: warning: JUNOS322716389: Ignoring unknown feature note that this has been fixed to allow all valid legacy licenses (even keys with only unknown features) in all other releases as per scopes. [PR1562700](#)
- On EX and EX-VC platforms, if post routing engine switchover, MAC address is configured to IRB interface (for ex: set interface irb.500 mac 00:11:22:33:44:55) on new primary Routing Engine, then the new primary Routing Engine might crash or go into DB mode. [PR1565213](#)

- EX2300 shows high FPC CPU usage, however, the system processes and kernel CPU usage doesn't add up to the overall FPC usage. This is due to a cosmetic issue with calculation of FPC CPU usage that has been resolved. [PR1567438](#)
- Observing traffic drop during ISSU due to LAG interface flap. [PR1569578](#)
- BUM traffic replication over VTEP is sending out more packets than expected and there seems to be a loop also in the topology. [PR1570689](#)
- A new virtual chassis (VC) member might join into VC as unexpected, if service image version on the VC member is different from the ones running on VC primary device. When this issue happens, the new VC member cannot forward packets. This issue affects EX platforms only. [PR1576774](#)
- On EX2300 and EX4650, if the system is upgraded from 20.2 or earlier release to 20.3 or later release, either using phone-home feature or when the system is in factory default state, the upgrade will fail with phone-home crash. [PR1601722](#)
- There is a remote possibility that during many reboots, the Junos VM goes into a state where NMI is needed to continue the reboot. There is no workaround for this and a subsequent reboot does not seem to hit this issue. [PR1601867](#)
- On the EX4300MP, DCPFE core will be seen with mac based vlan scale configuration after interface flap. [PR1578859](#)

Authentication and Access Control

- In all Junos platforms with LLDP and VOIP feature enabled on access VLAN port scenario, VOIP phone might fail to work if the alphabetical order of the data VLAN name is greater than the alphabetical order of the VOIP VLAN name. For example, the Data VLAN name is **C123** while the VOIP VLAN name is **A123**, so port vlan-id TLV will be populated as VOIP VLAN ID in LLDP packets instead of the expected Data VLAN ID. [PR1482275](#)

EVPN

- After a reboot during recovery process the ESI LAGs come up before the BGP sessions and routes/ARP entries are not synced [PR1487112](#)
- On all Junos and Junos Evolved platforms with EVPN-VxLAN scenario, the number of MAC-IP binding counters might reach the limit when MAC-IP is moved between interfaces. Since MAC-IP counters are not decremented when entry is deleted due to this defect, repeated moves will result in a limit (default value is 1024) that will be reached even though there are fewer entries. Meanwhile, traffic loss could be seen. [PR1591264](#)
- On all Junos platforms traffic loss might be seen if aggregated Ethernet bundle interface with ESI is disabled on primary Routing Engine followed by a Routing Engine switchover. [PR1597300](#)

Forwarding and Sampling

- **fast-lookup-filter** with match not supported in FLT Hardware might cause traffic drop. [PR1573350](#)

Infrastructure

- On EX Series switches except EX4300, EX4600, and EX9200, an interface is configured for single vlan or multiple vlans, if all these vlans of this interface have igmp-snooping enabled, then this interface will drop HSRPv2 (Hot Standby Router Protocol for IPv6) packets. But if some vlans do not have igmp-snooping enabled, then this interface is working fine. [PR1232403](#)
- On EX Series switches, if you are configuring a large-scale number of firewall filters on some interfaces, the FPC might crash and generate core files. [PR1434927](#)
- The following error message is observed continuously in AD with base configurations: **IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) failed.** [PR1485038](#)
- A double free vulnerability in the software forwarding interface daemon (sfid) process of Juniper Networks Junos OS allows an adjacently-connected attacker to cause a Denial of Service (DoS) by sending a crafted ARP packet to the device. Please refer <https://kb.juniper.net/JSA11162> for more information. [PR1497768](#)
- On EX platforms, traffic drop might be observed after a restart of the pfem process due to the stale route entry in TCAM (Ternary Content Addressable Memory). [PR1517497](#)
- Unable to verify jais-7.0R3-THIN.0.tgz in the EX4600 device due to space issue. [PR1548668](#)
- User while loading the kernel would see the message **GEOM: mmcsd0s.enh: corrupt or invalid GPT detected.** [PR1549754](#)
- When receives a unicast EAPOL (0x888e) with vlan588 tag at ae1 in this example, the packet is forwarded to ae0 without changing the vlanID to 3054. set vlans vlan588 vlan-id 588. set vlans vlan588 interface ae1.0. set vlans vlan588 interface ae0.0 mapping 3054 swap. [PR1580129](#)

Layer 2 Features

- GARPs were being sent whenever there was a MAC (fdb) operation (add or delete). This is now updated to send GARP when the interface is up and I3 interface attached to the VLAN. [PR1192520](#)
- On EX4600 platforms, if a change related to TPID is made in the Device Control Daemon, traffic might be dropped in Packet Forwarding Engine due to failure on Layer 2 learning or interfaces flapping. [PR1477156](#)

Network Management and Monitoring

- hrProcessorLoad is not supported on EX4300 and still shows up in the snmp walk. [PR1508364](#)

Platform and Infrastructure

- On the EX9200 device, 33 percent degradation with MAC learning rate is observed in Junos OS Release 19.3R1 compared to Junos OS Release 18.4R1. [PR1450729](#)
- When the DHCP relay mode is configured as no-snoop, we are observing the offer gets dropped due to incorrect ASIC programing. [PR1530160](#)
- This issue might be seen only in back-to-back GRES in about more than 40 to 50 iterations. No workaround available and FPC gets restarted. [PR1579182](#)
- On EX4300 platforms, when a firewall filter for broadcast traffic with discard action policer is applied to the loopback interface, all broadcast packets (including Layer 2 forwarding packets, such as DHCP discover packets) that match this filter rule might be dropped. [PR1597548](#)
- On EX4300 platforms with both enterprise style and service provider style configurations, an interface with enterprise style IFL and flexible-vlan-tagging configured, VLAN tagged traffic might be dropped due to incorrect programming in the system. [PR1598251](#)

User Interface and Configuration

- This PR removes Adobe Flash dependent elements from J-Web for EX and MX Series devices. For MX Series Routers, J-Web previously contained Flash elements on the following pages: Monitor-Interfaces Monitor-System view-Process Details Monitor-Routing-OSPF Information For EX Series Switches, J-Web previously contained Flash elements on the following pages: Monitor-Interfaces Monitor-System view-Process Details Monitor-Switching-IGMP Snooping Monitor-Virtual Chassis Monitor-POE Monitor-Security-Port Security Monitor-Routing-OSPF Information Monitor-Service-DHCP -Server Monitor-Service-DHCP -Relay. [PR1553176](#)
- The issue is seen on EX-series VC only which can be avoided with a simple workaround as to providing a valid package during upgrade command. [PR1557628](#)

SEE ALSO

[What's New | 41](#)

[What's Changed | 53](#)

[Known Limitations | 56](#)

[Resolved Issues | 63](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 63](#)
- [Resolved Issues: 20.3R2 | 67](#)
- [Resolved Issues: 20.3R1 | 69](#)

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

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

Resolved Issues: 20.3R3

General Routing

- **re_tvp_builtin_fwinfo_update: Unable to get firmware version** message is seen in chassisd. [PR1471938](#)
- The OSPF neighborhood gets stuck in the **Start** state after configuring the EVPN-VXLAN. [PR1519244](#)
- Traffic loss might be observed on interfaces in a VXLAN environment. [PR1524955](#)
- The JNH memory might leak on the Trio-based line cards. [PR1542882](#)
- On EX4300-48MP line of switches with Linux TVP architecture and Junos OS as VM, the Junos CLI outputs do not confirm if the Junos OS and the host kernel are compatible with each other. [PR1543901](#)
- In EX2300, high EVENTD CPU utilization upon receiving LLMNR and MDNS traffic. [PR1544549](#)
- The device might be out of service after configuring the em1 or em2 interface. [PR1544864](#)
- FPC(s) might not boot-up on EX9214 in a certain condition. [PR1545838](#)
- Classifier is not programmed in the hardware and error logs might be seen in syslog. [PR1548159](#)
- Two Routing Engines might lose communication if they have different Junos OS versions on EX series. [PR1550594](#)
- OIR of CBs might result in major errors and the Packet Forwarding Engine disable action halted traffic forwarding on the FPCs. [PR1554145](#)

- Traffic might be dropped when a firewall filter rule uses the then VLAN action. [PR1556198](#)
- On the EX4300 device, script fails while committing the IPSEC authentication configuration due to the missing algorithm statement. [PR1557216](#)
- The MAC addresses learned in a Virtual Chassis might fail, aging out in MAC scaling environment. [PR1558128](#)
- BGP NSR, RPD core seen after Routing Engine switchover. [PR1558814](#)
- Some transmitting packets might get dropped due to the **disable-pfe** action is not invoked when the fabric self-ping failure is detected. [PR1558899](#)
- The tunable optics **SFP+-10G-T-DWDM-ZR** doesn't work on EX devices. [PR1561181](#)
- EX3400VC - SMARTD pollutes syslog every 5 seconds after upgrade or system reboot. [PR1562396](#)
- On EX3400VC line of switches, the **DAEMON-7-PVIDB** throws syslog messages for every 12 to 14 minutes after you upgrade to Junos OS Release 19.1R3-S3. [PR1563192](#)
- The client authentication fails after GRES. [PR1563431](#)
- The JWeb upgrade might fail on EX2300 and EX3400. [PR1563906](#)
- On EX4650 platforms, **storm control** with IRB interface might not work correctly. [PR1564020](#)
- The Packet Forwarding Engine telemetry data might not be streamed out. [PR1566528](#)
- On the EX4600 device, the following internal comment is displayed **Placeholder for QFX platform configuration**. [PR1567037](#)
- The DF (Designated Forwarder) might not forward traffic. [PR1567752](#)
- Packet loss might be observed when sample based action is used in firewall filter. [PR1571399](#)
- Port-mirroring might not work when the analyzer output is a trunk interface. [PR1575129](#)
- Protocol convergence between end nodes might fail when L2PT is enabled on transit switch. [PR1576715](#)
- MVR configuration cannot be configured on EX2300-C switches. [PR1577905](#)
- The fxpc process might crash on EX platforms. [PR1578421](#)
- Random/silent reboot might be seen on EX2300-24MP and EX2300-48MP platforms. [PR1579576](#)
- The voice VLAN might not get assigned to the access interface. [PR1582115](#)
- The l2ald crash if a specific naming format is applied between a vlan-range and a single vlan. [PR1583092](#)
- DSCP Rewriting might fail to work on EX2300. [PR1586341](#)
- The SNMP trap for MAC notifications might not be generated when an interface is added explicitly under switch-options. [PR1587610](#)
- The rpd crash might be observed on the router running a scaled setup. [PR1588439](#)
- Packet loss could be observed on dynamically assigning VoIP vlan. [PR1589678](#)

- Traffic loss might be observed for interface configured in subnet 137.63.0.0/16. [PR1590040](#)
- The LLDP packet might loss on the EX-4300MP platform if configuring LLDP on the management interface. [PR1591387](#)
- Storm control profile might not be applied on EX2300 platforms. [PR1594353](#)

Class of Service (CoS)

- The buffer allocation for VCP ports might not get released in Packet Forwarding Engine after physically moving the port location. [PR1581187](#)

Forwarding and Sampling

- The configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. [PR1563641](#)

High Availability (HA) and Resiliency

- The ksyncd core might be observed while applying the configuration to Interface Logical. [PR1551777](#)

Infrastructure

- On observing the EX 4300 VC/VCF setup HEAP malloc(0) is detected. [PR1546036](#)
- The vme/me0 management interface cannot process any incoming packets. [PR1552952](#)
- Traffic related to IRB interface might be dropped when **mac-persistence-timer** expires. [PR1557229](#)
- Some MAC addresses might not be aged out on EX4300 platforms. [PR1579293](#)

Interfaces and Chassis

- MC-AE interfaces might go down if the same VRRP group-id is configured on multiple IRB units. [PR1575779](#)
- On EX2300, EX3400, and EX4300 platforms, the aggregated Ethernet interface might flap if **gether-options** is configured under a member interface of an aggregated Ethernet interface. [PR1576533](#)
- VRRP incorrect advertisement threshold values are seen on vrrp groups when VRRP is configured on EX2300 boxes. [PR1584499](#)

Layer 2 Features

- MAC addresses learnt from MC-LAG client device might flap between the ICL interface and MC-AE interface after one child link in MC-AE interface is disabled. [PR1582473](#)

Layer 2 Ethernet Services

- aggregated Ethernet interface flap might be seen during Nonstop software upgrade. [PR1551925](#)
- The DHCP client will be offline for 120 seconds after sending the DHCPINFORM message in the DHCP relay scenario. [PR1575740](#)

MPLS

- Incorrect EXP bit change might be seen in certain conditions under MPLS scenario. [PR1555797](#)

Platform and Infrastructure

- On Ex3400 VC, console access on backup VC member is not allowed. [PR1530106](#)
- Packets transiting via multicast-based VXLAN VTEP interface might be dropped post FPC restart. [PR1536364](#)
- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)
- The targeted-broadcast feature might send out duplicate packets. [PR1553070](#)
- The traffic might be dropped on Layer-3 LAG after rebooting or halting any member of EX4300 VC. [PR1556124](#)
- Junos OS: Upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart (vmcore) (CVE-2021-0283, CVE-2021-0284). [PR1557881](#)
- The LLDP neighbor advertisement on EX4300 might send the wrong 802.3 power format with TLV length 7 instead of length 12. [PR1563105](#)
- Last flapped, timestamp for interface fxp0 gets reset every time **monitor traffic interface fxp0** is executed. [PR1564323](#)
- PFEX might crash when soft error recovery feature is enabled on Packet Forwarding Engine. [PR1567515](#)
- On all EX9200 platforms with EVPN-VXLAN configured, the next-hop memory leak in MX Series ASIC happens whenever there is a route churn for remote MAC-IP entries learned bound to the IRB interface in EVPN-VXLAN routing instance. When the ASIC's next-hop memory partition exhausted the FPC might reboot. [PR1571439](#)
- Introduce two new major CMERRORs for XM chip-based line card to stabilize the running device. [PR1574631](#)
- DHCP packets with source IP as link-local address are dropped in EX4300. [PR1576022](#)
- Firewall filter is not programmed correctly and traffic would drop without warning. [PR1586433](#)
- On EX4300 platforms, the Egress RACL firewall filter might not get programmed correctly. [PR1595797](#)

Routing Protocols

- The ppmmd memory leak might cause traffic loss. [PR1561850](#)
- The rpd process might crash if there are more routes changed during the **commit-sync processing** window. [PR1565814](#)
- The untagged packets might not work on EX Series platforms. [PR1568533](#)

Virtual Chassis

- EX4600 and EX4300 mixed VC, Error message **ex_bcm_pic_eth_uint8_set** is seen when changing configuration related to interface. [PR1573173](#)
- EX4300 VCP might not come up after upgrade when 40GE-LX4 optics is used. [PR1579430](#)

Resolved Issues: 20.3R2

Forwarding and Sampling

- Configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. [PR1563641](#)

General Routing

- The DHCP discover packet might be dropped if DHCP inform packet is received first. [PR1542400](#)

Infrastructure

- On the EX4600 and EX4300 Virtual Chassis or Virtual Chassis Fabric, the VSTP configurations device becomes unreachable and nonresponsive after commit. [PR1520351](#)
- Traffic related to IRB interface might be dropped when **mac-persistence-timer** expires. [PR1557229](#)

Layer 2 Features

- The MAC address in the hardware table might become out of synchronization between the primary device and member in the Virtual Chassis after the MAC flaps. [PR1521324](#)

Platform and Infrastructure

- IRB MAC is not programmed in hardware when the MAC persistence timer expires. [PR1484440](#)
- While verifying the **Last-change op-state** value through XML, the **rpc-reply** message is inappropriate. [PR1492449](#)
- The mge interface might still stay up while the far end of the link goes down. [PR1502467](#)
- The output VLAN push might not work. [PR1510629](#)
- On the EX9200 devices, the Trio-based MPC memory leaks when an IRB interface is mapped to a VPLS instance or a Bridge-Domain. [PR1525226](#)
- On the EX4300 device, script fails while committing the IPSEC authentication configuration due to the missing **algorithm** statement. [PR1557216](#)
- The DHCP traffic might not be forwarded correctly when DHCP sends unicast packets. [PR1512175](#)
- The EX4300-48MP device might go out of service during a software upgrade operation. [PR1526493](#)
- On the EX2300 device, the following PoE message is observed: **poe_get_dev_class: Failed to get PD class info**. [PR1536408](#)
- On the EX3400 and EX2300 switches, the upgrade fails due to the lack of available storage. [PR1539293](#)

- The Slaac-Snoopd child process generates core file upon multiple switchovers on the Routing Engine. [PR1543181](#)
- On the EX9200 device, SF3 Fabric OIR issues is observed with Junos OS Release 23.1R1.8. [PR1555727](#)
- Traffic might be dropped when a firewall filter rule uses the **then VLAN** action. [PR1556198](#)
- The client authentication fails after GRES. [PR1563431](#)
- DHCP binding does not happen after GRES. [PR1515234](#)
- The FBF functionality on the EX4300 Virtual Chassis might be broken if the Virtual Chassis reboots or the IRB configuration is modified. [PR1531838](#)
- On the EX4300 device, the LLDP neighborship might not come up with the non-aggregated Ethernet interfaces. [PR1538401](#)
- The targeted-broadcast feature might not work after a reboot. [PR1548858](#)
- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)
- The l2cpd process might crash if the ERP is deleted after the switchover. [PR1517458](#)
- Traffic loss might be observed on interfaces in a VXLAN environment. [PR1524955](#)
- The **lldp-receive-packet-count** are not exchanged properly in the l2pt operation for LLDP after configuring protocols. [PR1532721](#)
- In every software upgrade, host must be upgraded. [PR1543890](#)
- The Broadcom chip FPC might crash during the system booting. [PR1545455](#)
- The output of the **show pfe route summary hardware** command displays random high free and used column for the IPv6 LPM(< 64' routes. [PR1552623](#)
- The **action-shutdown** statement of storm control does not work for the ARP broadcast packets. [PR1552815](#)
- The targeted-broadcast feature might send out duplicate packets. [PR1553070](#)
- FPC might not be recognized after power cycle (hard reboot). [PR1540107](#)
- The JNH memory might leak on the Trio-based line cards. [PR1542882](#)

Routing Protocols

- The OSPF neighborship gets stuck in the **Start** state after configuring the EVPN-VXLAN. [PR1519244](#)
- The OSPFv3 adjacency should not be established when IPsec authentication is enabled. [PR1525870](#)
- Sending multicast traffic to downstream receiver on a Trio based Virtual Chassis platforms might fail. [PR1555518](#)
- The dcpfe process might crash while updating VRF for multicast routes during IRB uninit. [PR1546745](#)

User Interface and Configuration

- The license errors might be returned on the backup Routing Engine when you try to commit configuration. [PR1543037](#)

Resolved Issues: 20.3R1

Authentication and Access Control

- The 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](#)
- The dot1x client will not be moved to the hold state when the authenticated P-VLAN is deleted. [PR1516341](#)

EVPN

- The VXLAN function might be broken because of a timing issue. [PR1502357](#)
- Unable to create a new VTEP interface. [PR1520078](#)

General Routing

- Constant memory leak might lead to FPC memory exhaustion. [PR1381527](#)
- Virtual Chassis split after network topology changed. [PR1427075](#)
- On the EX4600 device, traffic loss might be seen with framing errors or runts if MACsec is configured. [PR1469663](#)
- On the EX4600 switches, the DSCP marking might not work as expected if the fixed classifiers are applied to interfaces. [PR1472771](#)
- On EX4300, the output of "show security macsec statistics" shows high values incorrectly. [PR1476719](#)
- DHCP binding fails when the P-VLAN is configured with a firewall to block or allow certain IPv4 packets. [PR1490689](#)
- Traffic loss might be observed in a mixed-Virtual Chassis setup of QFX5100 and EX4300. [PR1493258](#)
- On the EX4650 switch, traffic loss might be seen under an MC-LAG scenario. [PR1494507](#)
- Authentication session might be terminated if the PEAP request is retransmitted by the authenticator. [PR1494712](#)
- Outbound SSH connection flap or memory leak issue might be observed during the high rate of pushing configuration to the ephemeral database. [PR1497575](#)
- Traffic might get dropped if the aggregated Ethernet member interface is deleted and then added, or an SFP transceiver of the aggregated Ethernet member interface is unplugged or plugged in. [PR1497993](#)

- In some cases, if we have an OSPF session on the IRB over LAG interface with a 40-Gigabit Ethernet port as member, the session gets stuck when restarted. [PR1498903](#)
- Firewall filter might not get applied on EX4600. [PR1499647](#)
- On the EX4300 Virtual Chassis with NSB and xSTP enabled, continuous traffic loss might be observed while performing GRES. [PR1500783](#)
- LLDP is not acquired when native VLAN-ID and tagged VLAN-ID are the same on a port. [PR1504354](#)
- The isolated VLAN from RADIUS is not deleted when the interface flaps. [PR1506427](#)
- LLDP might not work when P-VLAN is configured on EX Series Virtual Chassis. [PR1511073](#)
- Traffic might not flow according to the configured policer parameters. [PR1512433](#)
- 802.1X memory leak is observed. [PR1515972](#)
- MPPE-Send/Recv-key attribute is not extracted correctly by dot1xd. [PR1522469](#)
- "Drops" and "Dropped packets" counters in the output of "show interface extensive" command are double counting. [PR1525373](#)
- EX4300-MP device might go out-of-service during a software upgrade operation. [PR1526493](#)

Infrastructure

- The fxpc might crash when configuring scaled configuration with 4093 VLANs. [PR1493121](#)
- The IP communication between directly connected interfaces on EX4600 might fail. [PR1515689](#)
- OID ifOutDiscards reports zero and sometime shows a valid value. [PR1522561](#)

Interfaces and Chassis

- A stale IP address might be seen after a specific order of configuration changes under logical-systems scenario. [PR1477084](#)
- Traffic might drop because the next hop points to ICL even when the local MC-LAG is up. [PR1486919](#)

Layer 2 Ethernet Services

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

Layer 2 Features

- On EX4650, the third VLAN tag is not pushed onto the stack and SWAP is being done instead. [PR1469149](#)
- Traffic imbalance might be observed on EX4600 and QFX5000 switches when "hash-params" is not configured. [PR1514793](#)
- MAC address in the hardware table might not synhronize between the master and the member in Virtual Chassis after MAC flap. [PR1521324](#)

MPLS

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

Platform and Infrastructure

- IPv6 neighbor solicitation packets might be dropped in a transit device. [PR1493212](#)
- Packets get dropped when the next hop is IRB over the LT interface. [PR1494594](#)
- NSSU might fail on the EX4300 switches, because of a storage issue in the `/var/tmp` directory. [PR1494963](#)
- Traffic loss might be seen with framing errors or runts if MACsec is configured on the EX4300 switch. [PR1502726](#)

Routing Protocols

- The FPC process goes into the “NotPrsnt” state after upgrading the QFX5100 VC/VCF setup. [PR1485612](#)
- The BGP route-target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- Firewall filter could not work in certain conditions under a Virtual Chassis setup. [PR1497133](#)
- Packet loss might be observed for stream bLock:irb_lacp_tr_ospf while verifying traffic from access to core network for IPv4 or IPv6 interfaces. [PR1520059](#)

User Interface and Configuration

- J-Web does not display the correct flow-control status on EX Series devices. [PR1520246](#)

Virtual Chassis

- On the EX4650 device, a kldload error is observed while loading the module during booting. [PR1527170](#)

SEE ALSO

[What's New | 41](#)

[What's Changed | 53](#)

[Known Limitations | 56](#)

[Open Issues | 58](#)

[Documentation Updates | 72](#)

[Migration, Upgrade, and Downgrade Instructions | 72](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for EX Series switches.

SEE ALSO

What's New 41
What's Changed 53
Known Limitations 56
Open Issues 58
Resolved Issues 63
Migration, Upgrade, and Downgrade Instructions 72

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 72](#)

This section contains the upgrade and downgrade support policy for Junos OS for EX Series switches. 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

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and 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 | 41](#)

[What's Changed | 53](#)

[Known Limitations | 56](#)

[Open Issues | 58](#)

[Resolved Issues | 63](#)

[Documentation Updates | 72](#)

Junos OS Release Notes for JRR Series

IN THIS SECTION

- What's New | 74
- What's Changed | 76
- Known Limitations | 76
- Open Issues | 76
- Resolved Issues | 77
- Documentation Updates | 78
- Migration, Upgrade, and Downgrade Instructions | 79

These release notes accompany Junos OS Release 20.3R3 for the JRR 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 20.3R3 | 75
- What's New in Release 20.3R2 | 75
- What's New in Release 20.3R1 | 75

Learn about new features introduced in Junos OS Release main and maintenance releases for JRR Series Route Reflectors.

What's New in Release 20.3R3

There are no new features or enhancements to existing features for JRR Series in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for JRR Series in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Routing Protocols

- **Support for Implicit filter for default EBGp route propagation behavior without policies (ACX Series, JRR200, MX204, vRR and PTX5000)**—Starting in Junos OS Release 20.3R1, we've introduced a new configuration hierarchy, **defaults ebgp no-policy** at the existing **[edit protocols bgp]** hierarchy level. The configuration option separates the default policy for **receive** and **advertise**, into separate clauses (**accept**, **reject**, or **reject-always**) to allow the route propagation behavior of EBGp speakers to vary independently from its default behavior.

In earlier releases, the default behavior of BGP was to receive and advertise all routes. With the introduction of this feature, the default behavior still remains to “accept” all routes for both **receive** and **advertise**, but you also have an option to reject routes by default.

With the **reject** configuration, you can reject routes of type inet unicast and inet6 unicast in instance types master, vrf, virtual-router, and non-forwarding. With the **reject-always** configuration, you can reject all routes from being received or getting advertised, irrespective of address family or instance type. By using this feature, you can control traffic in leaf autonomous systems (AS) and thereby, prevent them from having to accidentally function as transit autonomous systems.

NOTE: The introduction of this implicit filter does not affect the existing deployments that rely on the default behavior.

[See [Implicit Filter for Default EBGp route propagation behavior without policies](#) and [defaults](#).]

SEE ALSO

[What's Changed | 76](#)

[Known Limitations | 76](#)

[Open Issues | 76](#)

[Resolved Issues | 77](#)

[Documentation Updates | 78](#)

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

What's Changed

There are no changes in behavior and syntax in Junos OS Release 20.3R3 for JRR Series Route Reflectors.

SEE ALSO

[What's New | 74](#)

[Known Limitations | 76](#)

[Open Issues | 76](#)

[Resolved Issues | 77](#)

[Documentation Updates | 78](#)

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

Known Limitations

There are no known limitations JRR Series in Junos OS Release 20.3R3 for JRR Series Route Reflectors.

SEE ALSO

[What's New | 74](#)

[What's Changed | 76](#)

[Open Issues | 76](#)

[Resolved Issues | 77](#)

[Documentation Updates | 78](#)

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

Open Issues

There are no open issues in Junos OS 20.3R3 Release for JRR Series Route Reflectors.

SEE ALSO

[What's New | 74](#)[What's Changed | 76](#)[Known Limitations | 76](#)[Resolved Issues | 77](#)[Documentation Updates | 78](#)[Migration, Upgrade, and Downgrade Instructions | 79](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 77](#)
- [Resolved Issues: 20.3R2 | 77](#)
- [Resolved Issues: 20.3R1 | 78](#)

Learn about resolved issues for JRR Series in Junos OS 20.3R3 Release.

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

Resolved Issues: 20.3R3

General Routing

- Alarm is generated when the **enhanced-ip** statement is configured. [PR1554292](#)
- Option-60 (Vendor-Class-Identifier) is not sent during ZTP. [PR1582038](#)

Resolved Issues: 20.3R2

General Routing

- JRR200:firewall filter with non-zero TTL value can cause commit error. [PR1531034](#)
- The tcp_timer_keep logs flooding on JRR200. [PR1533168](#)
- On JRR200 devices, four out of eight fans might not work. [PR1534706](#)

- On SRX4100 and SRX4200 devices, four out of eight fans might not work. [PR1534706](#)
- The **request system power-off** and the **request system halt** commands are not working as expected on JRR200. [PR1534795](#)

Resolved Issues: 20.3R1

There are no fixed issues in Junos OS Release 20.3R1 for JRR Series Route Reflectors.

SEE ALSO

What's New 74
What's Changed 76
Known Limitations 76
Open Issues 76
Documentation Updates 78
Migration, Upgrade, and Downgrade Instructions 79

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for JRR Series Route Reflectors.

SEE ALSO

What's New 74
What's Changed 76
Known Limitations 76
Open Issues 76
Resolved Issues 77
Migration, Upgrade, and Downgrade Instructions 79

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | 79

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

For information about software installation and upgrade, see the [JRR200 Route Reflector Quick Start](#) and the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No

Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes
-----------------------------	-----------	-------------------------------	-----	-----

For more information about EOL and 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 74
What's Changed 76
Known Limitations 76
Open Issues 76
Resolved Issues 77
Documentation Updates 78

Junos OS Release Notes for Juniper Secure Connect

IN THIS SECTION

- What's New | 81
- What's Changed | 82
- Known Limitations | 83
- Open Issues | 83
- Resolved Issues | 83

These release notes accompany Junos OS Release 20.3R3 for Juniper Secure Connect. 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 20.3R3 | 81](#)
- [What's New in Release 20.3R2 | 81](#)
- [What's New in Release 20.3R1 | 81](#)

Learn about new features introduced in the Junos OS main and maintenance releases for Juniper Secure Connect.

What's New in Release 20.3R3

There are no new features for Juniper Secure Connect in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features for Juniper Secure Connect in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Juniper Secure Connect

- **Juniper Secure Connect for SRX Series and vSRX next-generation firewalls**—Juniper Secure Connect is a client-based SSL-VPN application that allows you to securely connect and access protected resources on your network. This application, when combined with SRX Series Services Gateways, helps organizations quickly achieve dynamic, flexible, and adaptable connectivity from devices anywhere across the globe. Juniper Secure Connect extends visibility and enforcement from client to cloud using secure VPN connections.

The Juniper Secure Connect solution includes:

- **SRX Series firewall**—Serves as an entry and exit point for communication between users with Juniper Secure Connect and the protected resources on the corporate network or in the cloud.
- **Juniper Secure Connect application**—Secures connectivity between the protected resources and the host clients running Microsoft Windows, Apple macOS, Google Android, and iOS operating systems. The Juniper Secure Connect application connects through a VPN tunnel to the SRX Series firewall to gain access to the protected resources in the network.

Table 3: Feature Support for Juniper Secure Connect

Feature	Description
Multiplatform support	Supports Windows, macOS, Android, and iOS platforms.
Windows pre-domain login	Allows users to log on to the local Windows system through an already established VPN tunnel (using Windows Pre-Logon), so that user is authenticated to the central Windows domain or Active Directory.
Configuration support	Automatically validates that the most current policy is available before establishing the connection.
Biometric user authentication	Allows the user to protect their credentials using the operating system's built-in biometric authentication support.
Multifactor authentication (MFA)	Allows you to use multifactor authentication to extend the authentication.
Juniper Secure Connect license	Licenses are available in 1-year and 3-year subscription models.

[See [Juniper Secure Connect Administrator Guide](#), [Juniper Secure Connect User Guide](#), [Licenses for Juniper Secure Connect](#), and [Managing Licenses](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 82](#)
- [What's Changed in Release 20.3R2 | 83](#)
- [What's Changed in Release 20.3R1 | 83](#)

Learn about what changed in the Junos OS main and maintenance releases for Juniper Secure Connect.

What's Changed in Release 20.3R3

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 20.3R3.

What's Changed in Release 20.3R2

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 20.3R2.

What's Changed in Release 20.3R1

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 20.3R1.

Known Limitations

There are no known behavior or limitation for Juniper Secure Connect in Junos OS Release 20.3R3.

Open Issues

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

There are no known issues for Juniper Secure Connect in Junos OS Release 20.3R3.

Resolved Issues

Learn which issues were resolved in the Junos OS main and maintenance releases for Juniper Secure Connect.

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

Resolved Issues: 20.3R3

There are no resolved issues for Juniper Secure Connect in Junos OS Release 20.3R3.

Resolved Issues: 20.3R2

There are no resolved issues for Juniper Secure Connect in Junos OS Release 20.3R2.

Resolved Issues: 20.3R1

There are no resolved issues for Juniper Secure Connect in Junos OS Release 20.3R1.

Junos OS Release Notes for Junos Fusion for Enterprise

IN THIS SECTION

- What's New | 84
- What's Changed | 86
- Known Limitations | 87
- Open Issues | 87
- Resolved Issues | 88
- Documentation Updates | 89
- Migration, Upgrade, and Downgrade Instructions | 89

These release notes accompany Junos OS Release 20.3R3 for the Junos fusion for enterprise. 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 20.3R3 | 85
- What's New in Release 20.3R2 | 85
- What's New in Release 20.3R1 | 85

Learn about new features introduced in this release for Junos fusion for enterprise.

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

What's New in Release 20.3R3

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 20.3R2.

What's New in Release 20.3R1

There are no new features or enhancements to existing features for Junos fusion for enterprise in Junos OS Release 20.3R1.

SEE ALSO

What's Changed	86
Known Limitations	87
Open Issues	87
Resolved Issues	88
Documentation Updates	89
Migration, Upgrade, and Downgrade Instructions	89

What's Changed

IN THIS SECTION

- [What's Changed in Release Junos OS 20.3R3 | 86](#)
- [What's Changed in Release Junos OS 20.3R2 | 86](#)
- [What's Changed in Release Junos OS 20.3R1 | 86](#)

Learn about what changed in this release for Junos fusion for enterprise.

What's Changed in Release Junos OS 20.3R3

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 20.3R3 for Junos fusion for enterprise.

What's Changed in Release Junos OS 20.3R2

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 20.3R2 for Junos fusion for enterprise.

What's Changed in Release Junos OS 20.3R1

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 20.3R1 for Junos fusion for enterprise.

SEE ALSO

[What's New | 84](#)

[Known Limitations | 87](#)

[Open Issues | 87](#)

[Resolved Issues | 88](#)

[Documentation Updates | 89](#)

[Migration, Upgrade, and Downgrade Instructions | 89](#)

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.3R3 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	 84
What's Changed	 86
Open Issues	 87
Resolved Issues	 88
Documentation Updates	 89
Migration, Upgrade, and Downgrade Instructions	 89

Open Issues

There are no known issues in hardware and software in Junos OS Release for 20.3R3 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.

SEE ALSO

What's New	 84
What's Changed	 86
Known Limitations	 87
Resolved Issues	 88
Documentation Updates	 89
Migration, Upgrade, and Downgrade Instructions	 89

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 20.3R3 | 88](#)
- [Resolved Issues: Release 20.3R2 | 88](#)
- [Resolved Issues: Release 20.3R1 | 88](#)

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

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

Resolved Issues: Release 20.3R3

There are no fixed issues in Junos OS Release 20.3R3 for Junos fusion for enterprise.

Resolved Issues: Release 20.3R2

There are no fixed issues in Junos OS Release 20.3R2 for Junos fusion for enterprise.

Resolved Issues: Release 20.3R1

There are no fixed issues in Junos OS Release 20.3R1 for Junos fusion for enterprise.

SEE ALSO

What's New 84
What's Changed 86
Known Limitations 87
Open Issues 87
Documentation Updates 89
Migration, Upgrade, and Downgrade Instructions 89

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 for documentation for Junos fusion for enterprise.

SEE ALSO

[What's New | 84](#)[What's Changed | 86](#)[Known Limitations | 87](#)[Open Issues | 87](#)[Resolved Issues | 88](#)[Migration, Upgrade, and Downgrade Instructions | 89](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 89](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 91](#)
- [Preparing the Switch for Satellite Device Conversion | 92](#)
- [Converting a Satellite Device to a Standalone Switch | 93](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 93](#)
- [Downgrading Junos OS | 94](#)

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

Basic Procedure for Upgrading Junos OS on an Aggregation Device

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

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

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

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

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

To download and install Junos OS:

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

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

Customers in the United States and Canada, use the following commands, where *n* is the spin number.

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Switch

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

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases - 19.3 and 19.4 or downgrade to the previous two EEOL releases - 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

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

Downgrading 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 20.2, follow the procedure for upgrading, but replace the 20.2 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

[What's New | 84](#)

[What's Changed | 86](#)

[Known Limitations | 87](#)

[Open Issues | 87](#)

[Resolved Issues | 88](#)[Documentation Updates | 89](#)

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- [What's New | 95](#)
- [What's Changed | 96](#)
- [Known Limitations | 96](#)
- [Open Issues | 97](#)
- [Resolved Issues | 97](#)
- [Documentation Updates | 98](#)
- [Migration, Upgrade, and Downgrade Instructions | 99](#)

These release notes accompany Junos OS Release 20.3R3 for Junos fusion for provider edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features introduced in this release for Junos fusion for provider edge.

SEE ALSO

[What's Changed | 96](#)[Known Limitations | 96](#)[Open Issues | 97](#)

[Resolved Issues | 97](#)

[Documentation Updates | 98](#)

[Migration, Upgrade, and Downgrade Instructions | 99](#)

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 this release for Junos fusion for provider edge.

SEE ALSO

[What's New | 95](#)

[Known Limitations | 96](#)

[Open Issues | 97](#)

[Resolved Issues | 97](#)

[Documentation Updates | 98](#)

[Migration, Upgrade, and Downgrade Instructions | 99](#)

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.3R3 for Junos fusion for provider edge.

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

SEE ALSO

[What's New | 95](#)

[What's Changed | 96](#)

[Open Issues | 97](#)

[Resolved Issues | 97](#)

[Documentation Updates | 98](#)

[Migration, Upgrade, and Downgrade Instructions | 99](#)

Open Issues

There are no open issues in the Junos OS Release 20.3R3 for Junos fusion for provider edge.

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

SEE ALSO

[What's New | 95](#)

[What's Changed | 96](#)

[Known Limitations | 96](#)

[Resolved Issues | 97](#)

[Documentation Updates | 98](#)

[Migration, Upgrade, and Downgrade Instructions | 99](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 97](#)
- [Resolved Issues: 20.3R2 | 98](#)
- [Resolved Issues: 20.3R1 | 98](#)

This section lists the issues fixed in the Junos OS Release 20.3R3 for Junos fusion for provider edge.

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

Resolved Issues: 20.3R3

There are no fixed issues in the Junos OS Release 20.3R3 for Junos fusion for provider edge.

Resolved Issues: 20.3R2

There are no fixed issues in the Junos OS Release 20.3R2 for Junos fusion for provider edge.

Resolved Issues: 20.3R1

Junos Fusion Provider Edge

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

SEE ALSO

What's New 95
What's Changed 96
Known Limitations 96
Open Issues 97
Documentation Updates 98
Migration, Upgrade, and Downgrade Instructions 99

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for Junos fusion for provider edge.

SEE ALSO

What's New 95
What's Changed 96
Known Limitations 96
Open Issues 97
Resolved Issues 97
Migration, Upgrade, and Downgrade Instructions 99

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 99
- Upgrading an Aggregation Device with Redundant Routing Engines | 102
- Preparing the Switch for Satellite Device Conversion | 102
- Converting a Satellite Device to a Standalone Device | 104
- Upgrading an Aggregation Device | 106
- Upgrade and Downgrade Support Policy for Junos OS Releases | 106
- Downgrading from Junos OS Release 20.1 | 107

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

The download and installation process for Junos OS Release 20.3R3 is different from that for earlier Junos OS releases.

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

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

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

- For 64-bit software:

NOTE: We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot  
source/jinstall64-20.3R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.3R3.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

NOTE: We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot
source/jinstall64-20.3R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.3R3.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 20.3R3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

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

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

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

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

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

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

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

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

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

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

```
[edit]
```

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

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

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

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

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 20.3R3, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you

can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

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

Downgrading from Junos OS Release 20.1

To downgrade from Release 20.1 to another supported release, follow the procedure for upgrading, but replace the 20.1 **jinstall** package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

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

SEE ALSO

[What's New | 95](#)

[What's Changed | 96](#)

[Known Limitations | 96](#)

[Open Issues | 97](#)

[Resolved Issues | 97](#)

[Documentation Updates | 98](#)

Junos OS Release Notes for MX Series

IN THIS SECTION

- [What's New | 108](#)
- [What's Changed | 131](#)
- [Known Limitations | 137](#)
- [Open Issues | 143](#)
- [Resolved Issues | 156](#)
- [Documentation Updates | 188](#)
- [Migration, Upgrade, and Downgrade Instructions | 188](#)

These release notes accompany Junos OS Release 20.3R3 for the MX Series 5G Universal Routing Platforms. 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: 20.3R3 | 109](#)
- [New and Changed Features: 20.3R2 | 109](#)
- [New and Changed Features: 20.3R1 | 109](#)

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

New and Changed Features: 20.3R3

There are no new features or enhancements to existing features for MX Series in Junos OS Release 20.3R3.

New and Changed Features: 20.3R2

There are no new features or enhancements to existing features for MX Series in Junos OS Release 20.3R2.

New and Changed Features: 20.3R1

Hardware

- We've added the following features to the MX Series routers in Junos OS Release 20.3R1.

Table 4: Features Supported by MPC10E and MPC11E Line Cards on MX Series Routers

Feature	Description
Interfaces and chassis	<ul style="list-style-type: none"> • Support for MS-MPC on the MX2000-SFB3 Switch Fabric Board (SFB). The MS-MPC interoperates with MX2K-MPC11E, MPC9E, MPC8E, and MPC6E Modular Port Concentrators on MX2020 and MX2010 routers. • On MX2K-MPC11E line cards, you can configure Port 0 of every PIC as 400GbE ports or 200GbE ports using either QSFP56-DD optics or QSFP28-DD optics. You can channelize each of the 400GbE-capable ports either as four 100GbE interfaces or as two 100GbE interfaces. [See Port Speed on MX2K-MPC11E Overview.]
General routing	<ul style="list-style-type: none"> • Support for IP reassembly on GRE tunnel interfaces on: <ul style="list-style-type: none"> • MPC10E-15C-MRATE and MPC10E-10C-MRATE on MX240, MX480, and MX960 routers. • MX2K-MPC11E on MX2010 and MX2020 routers. [See Configuring Unicast Tunnels.] • Support for Mapping of Address and Port with Encapsulation (MAP-E) and IPv6 rapid deployment (inline 6rd) on: <ul style="list-style-type: none"> • MPC10E-15C-MRATE and MPC10E-10C-MRATE on MX240, MX480, and MX960 routers. • MX2K-MPC11E on MX2010 and MX2020 routers. [See Configuring Mapping of Address and Port with Encapsulation (MAP-E) and Configuring Inline 6rd.]

Table 4: Features Supported by MPC10E and MPC11E Line Cards on MX Series Routers (*continued*)

Feature	Description
Juniper telemetry interface	<ul style="list-style-type: none"> Support for resource paths to export traffic statistics from LDP and multipoint LDP sensors with gRPC. [See Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).] Support for LDP and multipoint LDP native sensors. [See sensor (Junos Telemetry Interface).]
Layer 3 features	<ul style="list-style-type: none"> Support for Layer 3 features. The MX2K-MPC11E interoperates with MS-MPC and MS-MIC-16G on MX2020 and MX2010 routers to support the following Layer 3 features: stateful firewall, NAT, IPsec, real-time performance monitoring (RPM), and MS MPC/MS-MIC-based inline flow monitoring services. [See Adaptive Services Overview.]
Multicast	<ul style="list-style-type: none"> Support for bidirectional Protocol Independent Multicast (PIM) on MPC10E and MX2K-MPC11E line cards running on MX240, MX480, MX960, MX2010 and MX2020 routers. These routers support GRES with NSR. [See Understanding Bidirectional PIM.] <p>NOTE: Junos OS Release 20.3R1 does not support anycast rendezvous point (RP) functionality and bidirectional PIM over next-generation multicast VPN (MVPN).</p> <ul style="list-style-type: none"> Support for Automatic Multicast Tunneling (AMT) relay on MPC10E and MX2K-MPC11E line cards running on MX240, MX480, MX960, MX2010, and MX2020 routers for IPv4 traffic. To identify a gateway, AMT relay uses a combination of the device IP address and port. [See Understanding AMT.] <p>NOTE: Junos OS Release 20.3R1 does not support AMT gateway.</p>
Network management and monitoring	<ul style="list-style-type: none"> Support for monitoring link degradation. You can monitor link degradation of the 10GbE, 40GbE, 100GbE, and 400GbE interfaces on the MX2K-MPC11E line cards. [See Link Degrade Monitoring Overview.] Support for inline continuity check messages (CCM) on MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards. You can configure inline CCM for up MEPs, down MEPs, and MIPs for all current supported topologies. [See Inline Transmission Mode.]
Security	<ul style="list-style-type: none"> Support for Media Access Control Security (MACsec) on logical interfaces (MPC10E only). VLAN tags are transmitted in cleartext, which allows intermediate switches that are MACsec-unaware to switch the packets based on the VLAN tags. [See Media Access Control Security (MACsec) over WAN.]

Table 4: Features Supported by MPC10E and MPC11E Line Cards on MX Series Routers (*continued*)

Feature	Description
Services applications	<ul style="list-style-type: none"> Support for inline video monitoring using media delivery index (MDI) criteria. [See Understanding Inline Video Monitoring on MX Series Routers.]
SNMP	<ul style="list-style-type: none"> Support for Junos OS SNMP on MPC10E-15C-MRATE, MPC10E-10C-MRATE, and MX2K-MPC11E line cards for the following multicast LDP MIB tables and objects: <ul style="list-style-type: none"> mplsMldpInterfaceStatsTable mplsMldpFecUpstreamSessPackets mplsMldpFecUpstreamSessBytes mplsMldpFecUpstreamSessDiscontinuityTime [See Standard SNMP MIBs Supported by Junos OS and SNMP MIB Explorer .]
Subscriber management and services	<ul style="list-style-type: none"> Support for resource monitoring for broadband edge subscriber management and services. [See Resource Monitoring for Subscriber Management and Services.]

- **Support for the JNP-SFP-10G-BX10D and JNP-SFP-10G-BX10U bidirectional transceivers (MX240, MX480, MX960, MX2008, MX2010 and MX2020)**—Starting in Junos OS Release 20.3R1, the MPC3E-3D-NG (with the MIC3-3D-10XGE-SFPP) and MPC5EQ-100G10G line cards on the MX240, MX480, MX960, MX2008, MX2010 and MX2020 routers support the JNP-SFP-10G-BX10D and JNP-SFP-10G-BX10U bidirectional transceivers.

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

- **Support for the JNP-SFP-10G-BX40D and JNP-SFP-10G-BX40U bidirectional transceivers (MX240, MX480, MX960, MX2008, MX2010 and MX2020)**—Starting in Junos OS Release 20.3R1, the MPC3E-3D-NG (with the MIC3-3D-10XGE-SFPP) and MPC5EQ-100G10G line cards on the MX240, MX480, MX960, MX2008, MX2010 and MX2020 routers support the JNP-SFP-10G-BX40D and JNP-SFP-10G-BX40U bidirectional transceivers.

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

Authentication, Authorization, and Accounting

- **Support for TCP authentication option (TCP-AO) for BGP and LDP connections (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can use TCP-AO to authenticate TCP segments exchanged during BGP and LDP sessions. It supports both IPv4 and IPv6 traffic. TCP-AO provides a framework to support multiple stronger algorithms, such as HMAC-SHA1 and AES-128, to create its message digest. TCP-AO supports up to 64 keys that can be used for a BGP or an LDP session. You can configure a new key for a BGP or LDP session during its lifetime without causing any session flap. Each key becomes active based on its configured start time.

In earlier releases, you could use only the TCP MD5 authentication method. It supports only MD5 algorithm to create its message digest.

[See [TCP Authentication Option \(TCP-AO\) for BGP and LDP Sessions](#) and [authentication-key-chains \(TCP-AO\)](#).]

Class of Service (CoS)

- **Support for MPLS EXP bits rewrite to all segment labels in segment routing stack (MX Series)**—Starting in Junos OS 20.3R1, on segment routing LSPs, creating an EXP rewrite rule for the egress interface on the ingress (provider edge) router imposes the rewrite rule to all transport labels in the stack. As a result, you don't need to configure rewrite rules on every segment in the LSP.

[See [exp](#).]

EVPN

- **Color-based mapping of EVPN-MPLS and EVPN services over SR-TE (ACX5448, EX9200, MX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can specify a color attribute along with an IP protocol next hop. The color attribute adds another dimension to the resolution of transport tunnels over static colored and BGP segment routing traffic-engineered (SR-TE) label-switched paths (LSPs). This type of resolution is known as the color-IP protocol next-hop resolution. With the color-IP protocol next-hop resolution, you must configure a resolution map and apply it to EVPN-MPLS and EVPN services, which includes E-Line, E-LAN and E-Tree. With this feature, you can enable color-based traffic steering of EVPN-MPLS and EVPN services.

[See [Segment Routing LSP Configuration](#).]

- **Tunnel endpoint in the PMSI tunnel attribute field for EVPN Type 3 routes (MX Series)**—Starting in Junos OS Release 20.3R1, you can set the tunnel endpoint in the Provider Multicast Service Interface (PMSI) tunnel attribute field to use the ingress router's secondary loopback address. When you configure multiple loopback IP addresses on the local provider edge (PE) router and the primary router ID is not part of the MPLS network, the remote PE router cannot set up a PMSI tunnel route back to the ingress router. To configure the router to use a secondary IP address that is part of the MPLS network, include the `pmsi-tunnel-endpoint pmsi-tunnel-endpoint` statement at the `[edit routing-instances routing-instance-name protocols evpn]` hierarchy level for both EVPN and virtual-switch instance types.

[See [evpn](#).]

High Availability (HA) and Resiliency

- **Higher scale and performance in RIFT (MX240, MX480, MX960, vMX, QFX5100, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-24YM, QFX5120-48YM, QFX5130-48C, QFX5200, QFX5210, and QFX10008)**—Starting in Junos OS Release 20.3R1, we've made the following improvements to increase the scalability and performance in Routing in Fat Tree (RIFT):
 - Prefixes in RIFT
 - Peers in RIFT
 - Convergence improvement with RIFT
 - BFD sessions with RIFT

[See [RIFT Overview](#).]

Interfaces and Chassis

- **Support for local preference when selecting forwarding next hops for load balancing (MX Series)**—Starting in Junos OS Release 20.3R1, we've expanded support for traffic to prefer local forwarding next hops rather than remote forwarding next hops for equal-cost multipath (ECMP) traffic flows and on aggregated Ethernet and logical tunnel interfaces for the following devices:
 - MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE)
 - MX2010 and MX2020 routers with MX2K-MPC11E

To configure local preference:

- For ECMP traffic flows, include the **ecmp-local-bias** statement at the **[edit forwarding-options load-balance]** hierarchy level.
- For aggregated Ethernet interfaces, include the **local-bias** statement at the **[edit interfaces aex aggregated-ether-options]** hierarchy level.
- For logical tunnel interfaces, include the **local-bias** statement at the **[edit interfaces rlt x logical-tunnel-options load-balance]** hierarchy level.

[See [ecmp-local-bias](#), [local-bias \(aggregated Ethernet\)](#), and [local-bias \(logical tunnel\)](#).]

- **Support for QSFP-100G-FR optical transceivers (MX204 and MX10003)**—Starting in Junos OS Release 20.3R1, you can use the QSFP-100G-FR optical transceivers in the MX10003 (installed with the JNP-MIC1 or JNP-MIC1-MACSEC MICs) and MX204 routers. You can use the **show chassis pic fpc-slot slot pic-slot slot** and **show chassis hardware** commands to view the details of the transceiver.

NOTE: The MX10003 routers with JNP-MIC1-MACSEC do not support unified in-service software upgrade (ISSU). However, the MX10003 routers with JNP-MIC1 support ISSU.

[See [Hardware Compatibility Tool](#).]

IP Tunneling

- **Support for IP-over-IP next-hop-based tunneling (MX Series, PTX1000, PTX10000, QFX10000, and QFX10002)**—Starting in Junos OS Release 20.3R1, we support an IP-over-IP encapsulation to facilitate IP overlay construction over an IP transport network. An IP network contains edge devices and core devices. To achieve higher scale and reliability among these devices, you need to use an overlay encapsulation to logically isolate the core network from the external network that the edge devices interact with. Among other supported encapsulation methods, only IP-over-IP allows transit devices to parse the inner payload and use inner packet fields for hash computation and customer edge devices to route traffic into and out of the tunnel without any throughput reduction. IP-over-IP relies on a next-hop-based infrastructure to support higher scale.

On MX Series routers, the routing protocol daemon (rpd) sends the encapsulation header with tunnel composite next hop and the Packet Forwarding Engine finds the tunnel destination address and forwards the packet. On PTX Series routers and QFX10000 switches, rpd sends the fully resolved next-hop-based tunnel to the Packet Forwarding Engine. You can either use static configuration or a BGP protocol configuration to distribute routes and signal dynamic tunnels. You can also configure Interface based firewall filters on any transit or egress device with an action to decapsulate IP-IP packets and forward it to the main instance or to a routing-instance as required.

[See [Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation](#).]

- **Support for filter-based decapsulation of IPv4 and IPv6 unicast traffic encapsulated in IPv4 IP-in-IP tunnels (MX Series, PTX1000, PTX10002, and QFX10002)**—Junos OS supports decapsulating IPv4 and IPv6 unicast traffic that has been encapsulated in IPv4 IP-in-IP tunnels using firewall filters. If the outer IPv4 header address matches the firewall configuration and the packet has **ipip** set as the protocol type, then the outer IPv4 header is removed and the packet is routed based on the inner IPv4 or IPv6 address. If the packet does not have the expected **ipip** header, the packet is dropped.

Configure this feature using the following CLI statements at the **[edit firewall family inet filter *filter-name* term *term-name*]** hierarchy:

- **from protocol *ipip***: Set the protocol type as IP-IP.
- **then decapsulate *ipip***: Decapsulate the IP-IP packet. The inner IP destination address is routed using the inet.0 routing table by default.
- **then decapsulate *ipip* routing-instance *routing-instance-name***: Decapsulate the IP-IP packet and route the inner destination address using the specified routing instance.

Use **show firewall** to view the configuration.

[See [filter \(Firewall Filters\)](#) and [Configuring IP Tunnel Interfaces](#).]

- **IPIP encapsulation for flexible tunnel interfaces (FTIs) (MX Series, PTX Series, and QFX10002)**—We've extended flexible tunnel interfaces (FTIs) and existing forwarding constructs to support configuring static

IPv4 IP-in-IP tunnels and RIB APIs. To configure an IP-in-IP tunnel on a FTI, use the **ipip** option at the **[edit interfaces interface-name unit logical-unit-number tunnel encapsulation]** hierarchy level.

[See [Configuring Flexible Tunnel Interfaces](#) and [ipip](#).]

Juniper Extension Toolkit

- **Juniper Extension Toolkit (JET) supports BFD Service APIs for routing protocol process (rpd) programmability (MX Series, PTX Series, QFX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can use programmable rpd (prpd) BFD APIs to add, update, and delete BFD sessions and subscribe to BFD events from outside applications. These APIs enable the integration of rpd with software-defined networking (SDN) controllers and increase the flexibility of your network. The prpd BFD APIs support BFD Echo-Lite sessions in single-hop IPv4 and IPv6 modes.

The following BFD Service APIs are supported:

- Initialize
- SessionAdd
- SessionUpdate
- SessionDelete
- SessionDeleteAll
- Subscribe
- Unsubscribe

Use the **show bfd session extensive** command to view BFD sessions. BFD sessions added through prpd BFD APIs are labeled with **PRPD:<session-id>** in the client field. The **<session-id>** is 1 for the first BFD session that is added, 2 for the second, and so on.

[See [show bfd session extensive](#) and [JET APIs on Juniper EngNet](#).]

- **Support for static backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (MX Series, PTX Series, and QFX10002)**—We've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to set up underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation. JET APIs notify the controller of active paths, interfaces, and changes to the interface state. The loop-free backup paths help quickly restore failed core transport networks built with only IP protocols.

[See [JET APIs on Juniper EngNet](#).]

- **Support for policy match condition to match programmed routes (MX Series, PTX Series, and QFX10002)**—We've introduced a new option programmed that allows policy matches for routes injected by JET APIs. To allow policy matches for routes injected by JET APIs, use the **programmed** option at the **[edit policy-options policy-statement policy-name term term-name from]** hierarchy level. To view details about programmed routes, use the **show route programmed (detail | extensive)** command.

[See [policy-statement](#) and [show route](#).]

- **RIB service API option to control route distribution (MX Series, PTX Series, and QFX10002)**—We've added a no-advertise flag to the RIB service API per-route RouteAttributes object to limit re-advertisement of the provisioned route. You can set this flag to TRUE to prevent the route from being redistributed to routing protocols and advertised to peers.

[See [JET APIs on Juniper EngNet](#).]

Junos OS XML, API, and Scripting

- **Support for REST API over nondefault virtual routing and forwarding (VRF) instance (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can execute Junos OS operational commands using the REST API over a nondefault VRF instance. The nondefault VRF instance can be a user-defined instance or the management instance, mgmt_junos.

The REST API allows you to execute Junos OS operational commands over HTTP(S). If you don't specify a routing instance, REST API requests are sent over the default routing instance. Use a nondefault VRF instance to improve security and make it easier to troubleshoot.

Use the **routing-instance** *routing-instance* statement at the **[edit system services rest]** hierarchy level to specify a nondefault VRF instance for REST API requests.

[See [Management Interface in a Nondefault Instance](#) and [rest](#).]

Junos Telemetry Interface

- **EVPN statistics export using JTI (MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016 and vMX routers, EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253 switches)**—Starting in Junos OS Release 20.3R1, you can use Junos telemetry interface (JTI) an remote procedure call (gRPC) services to export EVPN statistics from devices to an outside collector.

Use the following sensors to export EVPN statistics:

- Sensor for instance level statistics (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/**)
- Sensor for route statistics per peer (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/peer/**)
- Sensor for Ethernet segment information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/ethernet-segment/**). This includes EVPN designated forwarder ON_CHANGE leafs **esi** and **designated-forwarder**.
- Sensor for local interface information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/interfaces/**)
- Sensor for local IRB interface information (resource path **/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/irb-interfaces/**)

- Sensor for global resource counters and current usage (resource path `/junos/evpn/evpn-smet-forwarding/`)
- Sensor for EVPN IP prefix (resource path `/junos/evpn/l3-context/`)
- Sensor for EVPN IGMP snooping database (type 6) (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/sg-db/`)
- Sensor for EVPN IGMP join sync (type 7) and leave sync (type 8) (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/sg-db/sgdb-esi/`)
- Sensor to relate selected replicator on AR leaf on QFX5100, QFX5110, QFX5120, and QFX5200 switches (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/assisted-replication/`)
- Sensor for EVPN ON_CHANGE notifications (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/ethernet-segment/`)
- Sensor for overlay VX-LAN tunnel information (resource path `/network-instances/network-instance[instance-name='name']/protocols/protocol/evpn/vxlan-tunnel-end-point/`). This includes VTEP information ON_CHANGE leafs `source_ip_address`, `remote_ip_address`, `status`, `mode`, `nexthop-index`, `event-type` and `source-interface`.
- EVPN MAC table information (resource path `/network-instances/network-instance[instance-name='name']/mac_db/entries/entry/`)
- Sensor for MAC-IP or ARP-ND table (resource path `/network-instances/network-instance[instance-name='name']/macip_db/entries/entry/`)
- Sensor for MAC-IP ON_CHANGE table information (resource path `/network-instances/network-instance[name='name']/macip-table-info/`). Statistics include leafs `learning`, `aging-time`, `table-size`, `proxy-macip`, and `num-local-entries`.
- Sensor for MAC-IP ON_CHANGE entry information (resource path `/network-instances/network-instance[name='name']/macip-table/entries/entry/`). Statistics include leafs `ip-address`, `mac-address`, `vlan-id` and `vni`.
- Sensor for bridge domain or VLAN information (resource path `/network-instances/network-instance[instance-name='name']/bd/`)

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

- **Chassis management configuration and counters support on JTI (MX Series with MPC11E)**—Starting in Junos OS Release 20.3R1, Junos telemetry interface (JTI) supports streaming chassis management error (cmerror) configuration and counters to an outside collector using remote procedure calls (gRPC).

The following base resource paths are supported:

- `/junos/chassis/cmerror/configuration`
- `/junos/chassis/cmerror/counters`

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

- **Forwarding information base (FIB) sensor support on JTI (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can use the Junos telemetry interface (JTI) and remote procedure calls (gRPC) services to stream or export ON_CHANGE FIB, also known as forwarding table, statistics to outside collectors. This feature supports the OpenConfig YANG model OC-AFT.

To enable and manage FIB streaming, include the following statements on the client device:

- **set system fib-streaming** and **delete system fib-streaming** statements at the **[edit]** hierarchy level to launch or terminate the process.
- **set system fib-streaming traceoptions file *file-name*** statement at the **[edit]** hierarchy level to configure a logging file.
- **set system fib-streaming traceoptions flag *flag-name*** statement at the **[edit]** hierarchy level to configure various trace parameters.
- **set system fib-streaming traceoptions level *level-name*** statement at the **[edit]** hierarchy level to configure log levels.

Use the **restart fib-streaming** command to restart the process.

To show information about FIB streaming, use the following operational mode commands on the client device:

- **show fib-streaming**
- **show fib-streaming next-hop-groups**
- **show fib-streaming next-hops**
- **show fib-streaming routes ipv4-unicast**
- **show fib-streaming routes ipv6-unicast**
- **show fib-streaming routes mpls**

The following table shows supported sensors:

Table 5: Supported Sensors

Supported Sensors
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/id

Table 5: Supported Sensors (*continued*)

Supported Sensors
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/dscp[]
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/next-hop-group
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/interface
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/subinterface
/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/prefix
/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/prefix
/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/next-hop-group
/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/prefix
/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/prefix
/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/next-hop-group
/network-instances/network-instance/afts/mpls/label-entry/label
/network-instances/network-instance/afts/mpls/label-entry/state/label
/network-instances/network-instance/afts/mpls/label-entry/state/next-hop-group
/network-instances/network-instance/afts/mpls/label-entry/state/popped-mpls-label-stack
This leaf reports the same label value in case of pop or swap.

Table 5: Supported Sensors (*continued*)

Supported Sensors
<code>/network-instances/network-instance/afts/next-hop-groups/next-hop-group/id</code>
<code>/network-instances/network-instance/afts/next-hop-groups/next-hop-group/next-hops/nexthop/index</code>
<code>/network-instances/network-instance/afts/next-hop-groups/next-hop-group/next-hops/nexthop/state/weight</code>
<code>/network-instances/network-instance/afts/nexthops/nexthop/index</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/juniper/state/lsp-id</code>
This leaf is a new augmentation.
<code>/network-instances/network-instance/afts/next-hops/next-hop/state/ip-address</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/state/mac-address</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/state/pushed-mps-label-stack</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/interface-ref/state/interface</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/interface-ref/state/subinterface</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/juniper/state/mapped-next-hop-index</code>
This leaf is a new augmentation.

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

- **Support for policy forwarding table sensor on JTI (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can use Junos telemetry interface (JTI) and remote procedure calls (gRPC) services to stream policy forwarding table statistics on MX Series and PTX Series routers to outside collectors. The following resource paths are supported:
 - `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/`
 - `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/id`
 - `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/id`

- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/dscp[]`
- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/next-hop-group`
- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface`
- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/id`
- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/id`
- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/interface`
- `/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/subinterface`

The Junos OS class-of-service (CoS) classifiers do the code-point (CP) to forwarding-class (FC) and loss-priority (LP) mapping. The classifier used depends on the family configured on the logical interface. Devices running Junos OS support the following classifier types:

- Differentiated Services code point classifier (DSCP)
- DSCP IPv6
- MPLS EXP classifier inet-precedence
- IPv4 precedence classifier

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

- **Support for aggregated Ethernet interface ON_CHANGE with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX3000, PTX5000, PTX10001-36MR, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.3R1, Junos telemetry interface (JTI) supports ON-CHANGE statistics for aggregated Ethernet interfaces for minimum links and member interfaces.

To export these statistics to an outside collector using remote procedure call (gRPC) services and JTI, include the following resource paths in a subscription:

- `/interfaces/interface/aggregation/state/min-links/`
- `/interfaces/interface/aggregation/state/member/`

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

- **Increase the speed of telemetry sensor subscription installation (MX Series routers)**—Starting in Junos OS Release 20.3R1, Junos telemetry interface (JTI) supports enhancements to increase the sensor

subscription installation speed for collectors. Whether a dynamic sensor subscribe or unsubscribe request from a collector uses remote procedure calls (gRPC) services or gRPC Network Management Interface (gNMI) services to make the request, resource paths (sensors) in the request are individually validated and committed. The following enhancements shorten the subscription installation process and time:

- Validation is no longer done using the ephemeral database's configuration load operation.
- Network Agent instead uses information from sensor YANGs and the Packet Forwarding Engine's internal sensor table to validate the paths in a subscribe or unsubscribe request. Using these sources, Network Agent responds back to the collector with system-accepted paths and completes basic checks before proceeding to commit the request.
- Network Agent performs a single commit per subscribe or unsubscribe request instead of doing commits for each resource path in a request.

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

- **Support for fabric, optical, and FPC environment sensor on JTI (MX-2010 and MX-2020 routers with MPC11E)**—Starting in Junos OS Release 20.3R1, Junos telemetry interface (JTI) supports streaming fabric, optical, and Flexible PIC Concentrator (FPC) environment statistics to an outside collector using remote procedure calls (gRPC).

The following base resource paths are supported:

- `/junos/system/linecard/optics/`
- `/junos/system/linecard/environment/`
- `/junos/system/linecard/fabric/`

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

Layer 2 Features

- **Support for FEC 128 and FEC 129 VPLS with source packet routing (MX Series)**—Starting in Junos OS Release 20.3R1, Junos OS supports forwarding equivalence class (FEC) 128 and FEC 129 VPLS with Source Packet Routing in Networking (SPRING) with IS-IS, OSPF, and non-colored segment routing–traffic-engineering (SR-TE). Source packet routing or segment routing is applied in an MPLS network. You can use FEC 128 and FEC 129 VPLS with SPRING over MPLS as an alternative to LDP VPLS over MPLS.

[See [Example: Configuring a Multihomed VPLS \(FEC 128\)](#), [Example: Configuring VPLS Multihoming \(FEC 129\)](#), and [Understanding Source Packet Routing in Networking \(SPRING\)](#).]

Layer 2 VPN

- **Enable or disable control-word for static pseudowire in LDP VPLS instance and BGP VPLS mesh-group (MX Series)**—Starting in Junos OS Release 20.3R1, we've introduced the **control-word** and **no-control-word** options at the `[edit routing-instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor address static]` and `[edit routing-instances routing-instance-name protocols vpls neighbor address static]` hierarchy levels. The **control-word** configuration requests the other routers to insert a control word between the label stack and the MPLS payload.

[See [control-word](#) and [no-control-word](#).]

Layer 3 Features

- **Support for BGP Layer 3 VPN over IP-IP Tunnel (MX Series, PTX1000, QFX10002, and QFX10008)**—Starting in Junos OS Release 20.3R1, we support BGP Layer 3 VPN over IP over IP (IP-IP) tunnels to create a new transport service. IP-IP tunnels terminate into service-layer VRF, so you do not need to use a service label. This feature allows interoperability between the new VRF and traditional VRF, so both types of overlays can coexist in your network. You can use this feature to transition from an MPLS network to an IP fabric core network and to protect your network from distributed denial-of-service (DDoS) attacks.

To use VPN over an IP-IP tunnel, configure the **tunnel-attribute** statement at the `[edit policy-options policy-statement policy-name term term-name then]` or `[edit policy-options policy-statement policy-name then]` hierarchy level.

To configure the receiver to program the dynamic tunnel using the tunnel attribute, use the **extended-nexthop-tunnel** statement at the `[edit routing-instances routing-instance-name protocols bgp group group-name family (inet-vpn | inet6-vpn) unicast]` hierarchy level.

[See [BGP Layer 3 VPN over IP-IP Tunnels Overview](#), [family \(Protocols BGP\)](#), [policy-statement](#), [vrf-export](#), and [Configuring IP Tunnel Interfaces](#).]

MPLS

- **New output fields added in the `show path-computation-client lsp` extensive command (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.3R1, you'll see association details such as **Association type**, **ID**, and **source** in the output of the `show path-computation-client lsp` command when you use the command with the **extensive** option.

[See [show path-computation-client lsp](#).]

Multicast

- **Support for virtual tunnels in MVPN (MX240, MX480, and MX960)**—Starting in Release 20.3R1, Junos OS supports redundant virtual tunnels (VTs) and fast re-route (FRR) for both active/backup and active/active redundancy models.

VT interfaces are used in Layer 3 multicast VPNs (MVPN) to facilitate virtual routing and forwarding (VRF) table lookup based on MPLS labels and to provide resiliency.

[See [Resiliency in Multicast L3 VPNs with Redundant Virtual Tunnels](#).]

Network Management and Monitoring

- **Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can use the **probe** command to query the status of the probed interface. The proxy interface resides on the same node as the probed interface, or it can reside on a node to which the probed interface is directly connected.

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

To enable the **probe** command, configure the **extended-echo** statement under the **[edit system]** hierarchy.

[See [What is the Probe command?](#).]

- **SNMP support for RIB sharding (MX Series)**—Starting in Junos OS Release 20.3R1, you can enable RIB sharding to get network information from BGP MIB-4 and Layer 3 VPN MIB. To enable this feature, configure **rib-sharding** at the **[edit system processes routing bgp]** hierarchy level.

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

- **SNMP MIB support for Traffic Load Balancer (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.3R1, a new MIB and a few new MIB traps export the statistics of the Traffic Load Balancer application. The new MIB is `jnxTLBMIB` and the MIB traps are `juniperMIB(2636)`, `jnxTraps (4)`, and `jnxTLBNotifications (32)`.

[See [Enterprise-Specific SNMP MIBs Supported by Junos OS](#).]

- **Enhancements to sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.3R1, devices running Junos OS with upgraded FreeBSD support the following enhancements to sessions over outbound HTTPS:

- Connecting to multiple outbound HTTPS clients by configuring one or more clients at the **[edit system services outbound-https]** hierarchy level
- Configuring multiple backup gRPC servers for a given outbound HTTPS client
- Establishing a csh session
- Establishing multiple, concurrent NETCONF and csh sessions between the device running Junos OS and an outbound HTTPS client
- Configuring a shared secret that the outbound HTTPS client uses to authenticate the device running Junos OS
- Authenticating the client using certificate chains in addition to self-signed certificates

[See [NETCONF and Shell Sessions over Outbound HTTPS](#).]

Next Gen Services

- **GNFs support subscriber services (MX480 and MX960 with MX-SPC3)**—Starting in Junos OS Release 20.3R1, guest network functions (GNFs) running Next Gen Services with the MX-SPC3 card support the following subscriber services:
 - Captive portal content delivery (CPCD)
 - Logging and reporting function (LRF)
 - Deep packet inspection (DPI)
 - Junos Subscriber Aware policy and charging enforcement function (PCEF)
 - HTTP content management (HCM)

NOTE: To support the services traffic over abstracted fabric interfaces, a GNF that has an MX-SPC3 card assigned to it must also have a line card linked to it.

[See [MX-SPC3 Services Card](#).]

- **Support for flow tracing of service sets for Next Gen Services (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.3R1, you can perform flow tracing at the service-set level, which reduces file size and avoids having to sift through large files for information about a single service set.

[See [traceoptions \(Next Gen Services Service-Set Flow\)](#).]

- **Support for port block allocation for Next Gen Services (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.3R1, we support port block allocation (PBA) for Next Gen Services. PBA reduces logging in the system by allocating blocks of ports to a subscriber instead of a single port at a time. Subscribers are tracked based on their private IP address and this information is logged in the system logs. However, ports are reused at a high rate, making tracking of subscribers' usage and activity difficult. PBA enables you to easily track subscribers' usage and activity.

[See [block-allocation](#).]

Port Security

- **MACsec on logical interfaces (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.3R1, you can configure Media Access Control Security (MACsec) at the logical interface level on the MPC7E-10G line card. This configuration enables multiple MACsec Key Agreement (MKA) sessions on a single physical port. VLAN tags are transmitted in cleartext, which allows intermediate switches that are MACsec-unaware to switch the packets based on the VLAN tags.

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

- **Timer-based MACsec SAK refresh (MX10003, PTX10001, PTX10003, PTX10008, and PTX10016)**—Starting in Junos OS Release 20.3R1, you can configure a timer-based refresh of the secure association key (SAK) on a Media Access Control Security (MACsec)-secured link. The key server generates the SAK and refreshes it periodically. The key server also sets a refresh interval, by default, based on packet counter movement. If the refresh does not occur frequently, this can leave the SAK vulnerable to attack. You can enhance security of the SAK by configuring a shorter timer-based refresh interval.

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

Routing Protocols

- **Support for Implicit filter for default EBGp route propagation behavior without policies (ACX Series, JRR200, MX204, vRR and PTX5000)**—Starting in Junos OS Release 20.3R1, we've introduced a new configuration hierarchy, **defaults ebgp no-policy** at the existing **[edit protocols bgp]** hierarchy level. The configuration option separates the default policy for **receive** and **advertise**, into separate clauses (**accept**, **reject**, or **reject-always**) to allow the route propagation behavior of EBGp speakers to vary independently from its default behavior.

In earlier releases, the default behavior of BGP was to receive and advertise all routes. With the introduction of this feature, the default behavior still remains to “accept” all routes for both **receive** and **advertise**, but you also have an option to reject routes by default.

With the **reject** configuration, you can reject routes of type inet unicast and inet6 unicast in instance types master, vrf, virtual-router, and non-forwarding. With the **reject-always** configuration, you can reject all routes from being received or getting advertised, irrespective of address family or instance type. By using this feature, you can control traffic in leaf autonomous systems (AS) and thereby, prevent them from having to accidentally function as transit autonomous systems.

NOTE: The introduction of this implicit filter does not affect the existing deployments that rely on the default behavior.

[See [Implicit Filter for Default EBGp route propagation behavior without policies](#) and [defaults](#).]

- **TI-LFA SRLG protection and fate-sharing protection for OSPFv2 (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can configure Shared Risk Link Group (SRLG) protection and fate-sharing

protection for segment routing to choose a fast reroute path that does not include SRLG links and fate-sharing groups in the topology-independent loop-free alternate (TI-LFA) backup paths to avoid fate-sharing and SRLG failures. This is in addition to existing fast reroute options such as **link-protection** and **node protection** for segment routing.

To enable TI-LFA SRLG protection and fate-sharing protection with segment routing for OSPFv2, include the **srlg-protection** statement and the **fate-sharing-protection** statement respectively at the **[edit protocols ospf area *area-id* interface *name* post-convergence-lfa]** hierarchy level.

[See [Topology-Independent Loop-Free Alternate with Segment Routing for OSPF](#).]

- **BGP sharding for IPv4 and Ipv6 L3VPN, BGP-LU (MX Series, PTX-Series and vRR)**—Starting in Release 20.3R1, Junos OS supports BGP sharding and update IO features for these IPv4 and Ipv6 address families:
 - inet-vpn unicast
 - inet-vpn multicast (vrf.inet.2)
 - inet6-vpn unicast
 - inet6-vpn multicast (vrf.inet.2)
 - inet labeled-unicast
 - inet6 labeled-unicast

To enable BGP sharding, configure **rib-sharding** at the **[edit system processes routing bgp]** hierarchy level. Sharding is dependent on the update I/O thread feature. To enable update I/O, configure **update-threading** at the **[edit system processes routing bgp]** hierarchy level.

BGP Sharding is supported only on 64-bit routing protocol process (rpd) where the Routing Engine has at least 4 CPU cores and 16 GB of memory. To enable your device to always use 64-bit mode, use **set force-64-bit** at **[edit system processes routing]** hierarchy level. If you configure rib-sharding on a routing engine, RPD creates sharding threads. By default, the number of sharding threads created is the same as the number of CPU cores on the routing engine. Optionally, you can specify the number-of-shards you want to create. To set the number of sharding threads, use **set number-of-shards <number-of-shards>** at **[edit system processes routing bgp rib-sharding]** hierarchy level. To set the number of update threads, use **set number-of-threads <number-of-threads>** at the **[edit system processes routing bgp update-threading]** hierarchy level. To enable your device to always use 64-bit mode, use **set force-64-bit** at **[edit system processes routing]** hierarchy level.

[See [rib-sharding](#) and [update-threading](#).]

- **ECMP next-hop update rate throttling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.3R1, 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.

[See [pause-computation-during-churn](#).]

- **Support for Faster PFE Acks (MX Series Virtual Chassis)**—Starting in Junos OS Release 20.3R1, we support Faster PFE Acks to release Routing Engine kernel resources quicker. This support ensures that resource exhaustion scenarios are avoided

[See [virtual-chassis \(MX Series Virtual Chassis\)](#).]

- **Enabling Ifstate, peer infra, and TCP/IP stack parallelization on Virtual chassis (MX240, MX480, MX960, and MX2020)**—Starting in Junos OS Release 20.3R1, Virtual Chassis involving the listed MX Series devices support the following BFD features:
 - Ifstate parallelization
 - Peer infra parallelization
 - TCP and IP stack parallelization

These features are preserved on failover of any chassis when using Virtual Chassis.

[See [Understanding Bidirectional Forwarding Detection \(BFD\)](#).]

Segment Routing

- **SRv6 network programming in IS-IS (MX Series with MPC7E, MPC8E and MPC9E line cards)**—Starting in Junos OS Release 20.3R1, you can configure segment routing in a core IPv6 network without an MPLS data plane. This feature is useful for service providers whose networks are predominantly IPv6 and have not deployed MPLS. Such networks depend only on the IPv6 headers and header extensions for transmitting data. This feature also benefits networks that need to deploy segment routing traffic through transit routers that do not have segment routing capability yet. In such networks, the SRv6 network programming feature can provide flexibility to leverage segment routing without deploying MPLS.

To enable SRv6 network programming in an IPv6 domain, include the **srv6** statement at the **[edit routing-options source-packet-routing]** hierarchy level.

To advertise the Segment Routing Header (SRH) locator with a mapped flexible algorithm, include the **algorithm** statement at the **[edit protocols isis source-packet-routing srv6 locator]** hierarchy level.

To configure a topology-independent loop-free alternate backup path for SRv6 in an IS-IS network, include the **transit-srh-insert** statement at the **[edit protocols isis source-packet-routing srv6]** hierarchy level.

[See [How to Enable SRv6 Network Programming in IS-IS Networks](#).]

- **Support for LDP Tunneling over Segment Routing Traffic Engineering (MX Series, PTX Series, and ACX5448)**—Starting in Junos OS Release 20.3R1, you can tunnel LDP LSPs over Segment Routing Traffic Engineering (SR-TE) in your network. Tunneling LDP over SR-TE provides consistency and co-existence of both LDP LSPs and SR-TE LSPs.

[See [Tunneling LDP over SR-TE](#).]

Services Applications

- **Enhancements to the RFC 2544-based benchmarking tests (MX Series)**—Starting in Junos OS Release 20.3R1, we've extended support for these tests onto the following devices:
 - MX240, MX480, and MX960 routers with the MPC7E-MRATE or MPC7E-10G line card
 - MX2008, MX2010, and MX2020 routers with the MX2K-MPC8E or MX2K-MPC9E line card
 - MX204 and MX10003 (with the LC2103 line card) routers

You can use the RFC 2544 tests to measure and demonstrate the service-level agreement (SLA) parameters before service activation. The tests measure throughput, latency, frame loss rate, and link bursts. This enhancement supports the Layer 2 reflector (ingress direction) for family types **bridge** and **vpls**. To set the ingress direction of a test, configure the **family bridge** or **family vpls** statement and the **direction ingress** statement at the `[edit services rpm rfc2544-benchmarking tests test-name name]` hierarchy level.

To run the tests, you must configure the reflector function on the corresponding MPC. To configure the reflector function, include the **fpc fpc-slot-number slamon-services rfc2544** statement at the `[edit chassis]` hierarchy level.

[See [Understanding RFC2544-Based Benchmarking Tests on MX Series Routers](#).]

- **Support for sampling and tunneling performance improvement (MX204)**—Starting in release 20.3R1, Junos OS allows fabric-bound packets to take a new fabric loopback path, freeing up the WAN bandwidth and thus improving the sampling and tunneling performance of the router. You can configure fabric-side loopback by using the **fabric loopback wan off** statement or switch to WAN side by using the **fabric loopback wan on** statement at the `[edit chassis fpc slot-number]` hierarchy level. By default, Junos OS uses fabric loopback for the loopback packets.

[See [Tunnel Services Overview](#) and [Understanding Inline Active Flow Monitoring](#).]

- **Support for hardware timestamping of Two-Way Active Measurement Protocol (TWAMP) and real-time performance monitoring (RPM) probe messages (MX10008, MX10016, PTX10008, and PTX10016)**—Starting in Junos OS Release 20.3R1, we've extended support for hardware timestamping of TWAMP and RPM probe messages. Hardware timestamping is enabled by default for TWAMP, but you must configure it for RPM. You use TWAMP and RPM to measure IP performance between two devices in a network. By configuring hardware timestamping for RPM, you can account for the latency in the communication of probe messages and generate more accurate timers in the Packet Forwarding Engine. To configure hardware timestamping for RPM, include the **hardware-timestamping** statement at the `[edit services rpm probe probe-owner test test-name]` hierarchy level.

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

- **New configuration option for displaying descriptive information of session logs (MX Series)**—Starting in Junos OS Release 20.3R1, you can configure an option to display more descriptive information of session logs. You can configure the **enable-descriptive-session-syslog** statement at the `[edit services`

service-set service-set-name service-set-options] hierarchy level to enable syslog to display information related to inside and outside packets, byte count, and the session IDs for both open and close sessions.

[See[service-set-options](#).]

Software Defined Networking (SDN)

- **Programmable flexible VXLAN tunnels (MX960 with MPC10E; MX2010 and MX2020 with MPC11E)**—Starting in Junos OS Release 20.3R1, we support flexible VXLAN tunnels in a data center environment that includes one or more controllers. In this environment, one or more of the supported MX Series routers can function as data center edge gateways that exchange Layer 2 traffic with hosts in a data center. Through the use of static routes and tunnel encapsulation and de-encapsulation profiles, the Layer 2 traffic is dynamically tunneled over an intervening IPv4 or IPv6 network.

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

[See [Understanding Programmable Flexible VXLAN Tunnels](#) and [JET APIs on Juniper EngNet](#).]

System Management

- **Clock synchronization support (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS release 20.3R1, we’ve enhanced the clock synchronization (clksync) module. When the CBO clock failure alarm is raised, automatic Routing Engine switchover occurs. The new primary Routing engine Engine connection is made, the clksync module gets the notification.

[See [Understanding Clock Synchronization](#).]

SEE ALSO

What's Changed 131
Known Limitations 137
Open Issues 143
Resolved Issues 156
Documentation Updates 188
Migration, Upgrade, and Downgrade Instructions 188

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 131](#)
- [What's Changed in Release 20.3R2 | 132](#)
- [What's Changed in Release 20.3R1 | 134](#)

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

What's Changed in Release 20.3R3

Junos OS XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or **set refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

Layer 2 Ethernet Services

- **Active leasequery-based bulk leasequery (MX Series)**—The **overrides always-write-option-82** and **relay-option-82 circuit-id** configurations at the [**edit forwarding-options dhcp-relay**] hierarchy level are not mandatory for active leasequery-based bulk leasequery. For earlier releases, the **overrides always-write-option-82** and **circuit-id** configurations are mandatory for active leasequery-based bulk leasequery.

For regular bulk leasequery between relay and server without any active leasequery, the **overrides always-write-option-82** and **relay-option-82 circuit-id** configurations are mandatory.

[See [bulk-leasequery \(DHCP Relay Agent\)](#).]

- **Link selection support for DHCP (MX Series)**—We've introduced **link-selection** statement at the [**edit forwarding-options dhcp-relay relay-option-82**] hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Earlier to this release, the DHCP relay drops packets during the renewal DHCP process as the DHCP Server uses the leaf's address as a destination to acknowledge DHCP renewal message.

[See [relay-option-82](#).]

What's Changed in Release 20.3R2

General Routing

- **Round-trip time load throttling for pseudowire interfaces (MX Series)**—The Routing Engine supports round-trip time load throttling for pseudowire (ps) interfaces. In earlier releases, only Ethernet and aggregated Ethernet interfaces are supported.

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

- **Updates to ON-CHANGE and periodic dynamic subscriber interface metadata sensors (MX Series routers and EX9200 line of switches)**—
 - Notifications are sent when subscribers log in on either IP demux or VLAN demux interfaces. In earlier releases, login notifications are sent only for IP demux logins.
 - The **interface-set** end path has been added to the logical interface metadata. The interface-set field appears in both ON-CHANGE and periodic notifications. In earlier releases, this field is not included in the sensor metadata or notifications.

[See [gRPC Sensors for Subscriber Statistics and Queue Statistics for Dynamic Interfaces and Interface-Sets \(Junos Telemetry Interface\)](#).]

- **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](#).]

- **Changes to Junos XML operational RPC request tag names (MX480)**—Starting in Junos OS Release, we've updated the Junos XML request tag name for the below operational RPCs. The changes include:
get-security-associations-information is changed to **get-re-security-associations-information**
get-ike-security-associations-information is changed to **get-re-ike-security-associations-information**

[See [Junos XML API Operational Developer Reference](#).]

Junos XML API and Scripting

- The `jcs:invoke()` function supports suppressing root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—The `jcs:invoke()` extension function supports the **no-login-logout** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- The `jcs:invoke()` function supports suppressing root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—The `jcs:invoke()` extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

MPLS

- **The `show mpls lsp extensivel` and `show mpls lsp detail` commands display next-hop gateway LSPid**—When you use the `show mpls lsp extensivel` and `show mpls lsp detail` commands, you'll see next-hop gateway LSPid in the output.

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the `verbose` statement at the `edit system export-format json` hierarchy level. The default format to export configuration data in JSON changed from `verbose` format to `ietf` format starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `edit system export-format json` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See [export-format](#).]

What's Changed in Release 20.3R1

EVPN

- **New output flag for the `show bridge mac-ip table` command 9MX series**—The Layer 2 address learning daemon (l2ald) does not send updated MAC and IP address advertisements to the routing protocol daemon (rpd) when an IRB interface is disabled in an EVPN-VXLAN network. We've added the NAD flag in the output of the `show bridge mac-ip-table` command to identify the disabled IRB entries in which the MAC and IP address advertisement will not be sent.

[See [show bridge mac-ip-table](#).]

General Routing

- **Change in `show oam ethernet connectivity-fault-management mep-statistics` command (MX Series)**—You can now view the real-time statistics for continuity check messages (CCM) inline sessions for MPC10E (MPC10E-10C-MRATE and MPC10E-15C-MRATE) and MPC11E (MX2K-MPC11E) line cards only when you execute the `show oam connectivity-fault-management mep-statistics local-mep local-mep-id maintenance-association name` twice in immediate succession. If you execute the command once, the values are incorrectly displayed.

[See [show oam ethernet connectivity-fault-management mep-statistics](#).]

- **MS-MPC and MS-MIC service package (MX240, MX480, MX960, MX2020, MX2010, and MX2008)**—PICs of Multiservices MPCs (MS-MPCs) and Multiservices MICs (MS-MICs) do not support any service package than other extension-provider. These PICs always come up with the extension-provider service-package, irrespective of the configuration. If you try to configure any other service package, for these PICs by using the command `set chassis fpc slot-number pic pic-number adaptive-services service-package`, an error is logged. Use the `show chassis pic fpc-slot slot pic-slot slot` command to view the service package details of the PICs of MS-MPC and MS-MIC.

[See [extension-provider](#).]

High Availability (HA) and Resiliency

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the `show rift tie` output.

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](#).]

Interfaces and Chassis

- **Change in support for interface-transmit-statistics statement**—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. In Junos OS Release 20.3R1, 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.

Junos OS, XML, API, and Scripting

- **Changes to Junos XML RPC request tag names (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've updated the Junos XML request tag name for some operational RPCs to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include:
 - Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
 - Uppercase characters are converted to lowercase.

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

J-Web

- **Adobe Flash Player support (MX Series)**—Adobe Flash Player support will end on December 31, 2020. Due to this, the Flash dependent J-Web monitor pages will not load correctly for Junos OS Release 20.3R1 and earlier releases.

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 20.3R1, you can configure the `natt-install-interval seconds` option under the `[edit services ipsec-vpn rule rule-name term term-name then dynamic]` hierarchy to specify the duration of delay in installing IPsec security association (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 20.3R1, 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 displayed 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.](#)]

- **Command to view summary information for resource monitor (EX9200 line of Ethernet switches and MX Series routers)**—The `show system resource-monitor` command enables you to view many statistics about the use of memory resources for all line cards or for a specific line card in the device. It 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.](#)]

SEE ALSO

What's New	108
Known Limitations	137
Open Issues	143
Resolved Issues	156
Documentation Updates	188
Migration, Upgrade, and Downgrade Instructions	188

Known Limitations

IN THIS SECTION

- [EVPN | 138](#)
- [General Routing | 138](#)
- [Interfaces and Chassis | 139](#)
- [MPLS | 139](#)
- [Platform and Infrastructure | 140](#)
- [Routing Protocols | 140](#)
- [Network Management and Monitoring | 140](#)
- [Platform and Infrastructure | 141](#)
- [Routing Protocols | 142](#)
- [Subscriber Management and Services | 142](#)

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.

EVPN

- EVPN service over scaled SR-TE uncolored is unsupported. [PR1499719](#)

General Routing

- Access facing FPC's CPU stays at 100% for 5-6 minutes after configuration change. [PR1447003](#)
- Upon MPC11 boot up, errors such as following could be seen, but these are harmless and does not have functional impact. **Aug 27 22:29:15 sherekhan kernel: i2c i2c-100: (11/1:0x41) i2c transaction error (0x00000002)** **Aug 28 02:46:06 sherekhan kernel: i2c i2c-64: (7/1:0x41) i2c transaction error (0x00000002).** [PR1457655](#)
- The traffic stops when volume quota is reached but resumed incorrectly after APFE failover. Threshold and quota values are not updated to the secondary APFE and if quota is hit on primary APFE and traffic starts dropping due to quota and switchover happens, traffic will continue to flow until quota is hit; no work around. [PR1463723](#)
- If you move the MX2K-MPC11E line card from one guest network function (GNF) to another in an in-chassis Junos node slicing setup, the line card takes longer time than expected to come online. [PR1469729](#)
- During BGP convergence, eg full internet table load, the BFD and LACP protocol on the router may flap. [PR1472587](#)
- Line card crashes when there is a change in the ps interfaces which has active subscribers on it. [PR1486665](#)
- When the number of next-hop selectors to be repaired is very high, then time to repair them during FRR would go up and could increase packet losses. This would be observed specially when there are many unilist next hops with different next-hop selectors and each has a member next hop with a logical interface over the same physical interface, which goes down. [PR1490070](#)
- EVPN-VPWS, L3 VPN and L2 VPN FRR convergence time with AE as the active core interface is not meeting <50ms and might be 100ms to 150ms. [PR1492730](#)
- When you switch modes, either from PIC mode to port mode or port mode to PIC mode, the PIC restarts automatically. To change the default speed, you must select a port and configure a different port speed on it and reset both the PICs for the configuration to take effect. For example, select 40GE or 100GE on PIC 0 and 10GE on PIC 1. For this configuration to take effect, you must reset both PICs. The following logs might be seen: **chassisd[15345]: should_summit_pic_bounce_for_port_speed_change: FPC 0 PIC 1 Bounce not needed for port speed config change** But other PIC configuration change needs the bounce of both PICs, so above log is creating confusion. [PR1500429](#)
- Problem: Few packet loss during MBB with FTI tunnels Actions that trigger problem: If FTI IFL update (containing underlay NH) reaches PFE prior to underlay NH update, then some packet losses are expected. Details: Kernel sends the underlay NH as part of IFL update to PFE. Based on tunnel destination reachability, kernel sends FTI IFL updates to PFE containing underlay NH. As IFL to NH dependency is

not present, FTI IFL update can reach PFE ahead of underlay NH update. In most of the cases, underlay NH is expected to reach first than IFL update. But it is not guaranteed. Hence, as PFE is unaware of new underlay NH, there could be some packet losses during MBB scenarios. Result: MBB can be a challenge & could experience some packet losses. This design limitation exists in all types of FTI tunnels.

Workaround: None. [PR1507779](#)

- On a fully scaled system where all the slices are utilized by different families of CLI filters, if we try to call delete for one family and add/change for another family with a higher number of filter terms which requires either expansion of the filter or creation of a new filter, the Packet Forwarding Engine fails to add the new filter as we are getting messages out of sequence?that is, the add/change of filter is called earlier than the delete of another filter that will free up the slices. [PR1512242](#)
- Some memory leaks have been observed in the JET Service Daemon (JSD) process when one or more collectors are connecting and disconnecting to and from the router. These are observed in the gRPC stack code which is third party. The amount of memory leaked is relatively small. However, these leaks could increase with more frequent collector connects and disconnects. As a result of the memory leaks, the JSD process's memory size can increase to a value that is higher than normal (for example, when the gRPC connections are established and stable) but is unlikely to cause any adverse effects to the system with streaming telemetry. [PR1512296](#)
- MX10003 MPC will support a fixed port PIC (6xQSFP) and a modular TIC (12xQSFP28) which can be of two types - Ethernet TIC and MACSEC TIC. The MACsec TIC doesn't support unified ISSU and hence link flaps are expected on MACsec TIC. [PR1514694](#)
- Currently I2.5 traffic support is only for p2p and not for TS. So Routing Engine traffic doesn't work with traffic selector case. Only transit traffic will work with traffic selector based tunnels. For Routing Engine traffic needs to use p2p tunnels. [PR1534248](#)

Interfaces and Chassis

- Mcae don't update PW state on receiving remote PW state update on the first sync up. [PR1503102](#)

MPLS

- Process rpd might crash after network service configuration changed (like changing the range of MPLS labels) without rebooting all the Routing Engines (which is a system mandatory step). [PR1461468](#)
- As FPC1 is MPC 3D which has 2GB memory and when lot of routes gets injected the memory usage will increase and the below message was seen only once. FPC Resource Monitor: FPC 0 and 1 Heap Memory has crossed free memory watermark of 20. When this message was seen might be at that time the memory usage might have reached to 80% and when the memory usage reaches to 80% it prints this message. In this case maybe for a moment it might have gone to 80% and then when the routes were deleted it came down to 78% and it stays at that number. [PR1513436](#)

- With local reversion ON, there is a possibility of transit router not informing headend of RSVP disabled link when link is flapped more than once. Workaround is to remove local-reversion configuration. [PR1576979](#)

Platform and Infrastructure

- On the MX platform with Protocol Independent Multicast (PIM) implemented and the number of IGMP groups exceeding 15000, join message (S, G) might not be created after graceful Routing Engine switchover (GRES). [PR1457166](#)
- Unknown unicast filter applied in EVPN routing-instance blocks unexpected traffic. [PR1472511](#)
- With sensor being subscribed via Junos Telemetry Interface (JTI), after the interface is deleted/deactivated/disabled, the TCP connection is still established, and the CLI command of "show agent sensors" still shows the subscription. [PR1477790](#)
- EVPN Aliasing and load-balancing for L2 traffic is not working with Dynamic Link Next-Hop. EVPN Alaising with DLNH for L2 traffic is not supported for Junos OS 20.3 and previous releases. [PR1504412](#)
- RPM is a Juniper Networks proprietary feature. In the case of Junos OS, the RPM client never set the DF bit. Hence, we don't see this issue between a Junos OS RPM client and a Junos OS RPM server. Whereas in the case of Junos OS Evolved, the RPM client is setting the DF bit while sending the RPM probes to RPM server. In case of Junos OS TVP-based platforms, the RPM server is not able to decode the DF bit properly. This causes an interwork failure between Junos OS Evolved as the RPM client and TVP platforms as the RPM server (and vice versa). This issue is not applicable for non-TVP-based Junos OS platforms acting as the RPM server. [PR1508127](#)

Routing Protocols

- When enhanced ip is enabled we need to do restart routing else we might see label related issues. [PR1577451](#)

Network Management and Monitoring

- **SNMP Support for RIB Sharding and Threading (MX Series)**—In Junos OS Release 20.3R1, when you enable RIB Sharding, BGP MIB and L3VPN MIB don't support the below attributes:

Unsupported attributes for BGP MIB

- bgp4PathAttrPeer
- bgp4PathAttrIpAddrPrefixLen
- bgp4PathAttrIpAddrPrefix
- bgp4PathAttrOrigin

- bgp4PathAttrASPathSegment
- bgp4PathAttrNextHop
- bgp4PathAttrMultiExitDisc
- bgp4PathAttrLocalPref
- bgp4PathAttrAtomicAggregate
- bgp4PathAttrAggregatorAS
- bgp4PathAttrAggregatorAddr
- bgp4PathAttrCalcLocalPref
- bgp4PathAttrBest
- bgp4PathAttrUnknown

Unsupported attributes for L3VPN MIB

- mplsL3VpnVrfRtInetCidrDestType
- mplsL3VpnVrfRtInetCidrDest
- mplsL3VpnVrfRtInetCidrPfxLen
- mplsL3VpnVrfRtInetCidrPolicy
- mplsL3VpnVrfRtInetCidrNHopType
- mplsL3VpnVrfRtInetCidrNextHop
- mplsL3VpnVrfRtInetCidrIfIndex
- mplsL3VpnVrfRtInetCidrType
- mplsL3VpnVrfRtInetCidrProto
- mplsL3VpnVrfRtInetCidrAge
- mplsL3VpnVrfRtInetCidrNextHopAS
- mplsL3VpnVrfRtInetCidrMetric
- mplsL3VpnVrfRteXCPointer
- mplsL3VpnVrfRtInetCidrStatus

Platform and Infrastructure

- On the MX platform with Protocol Independent Multicast (PIM) implemented and the number of IGMP groups exceeding 15000, join message (S,G) might not be created after graceful Routing Engine switchover (GRES). [PR1457166](#)
- Unknown unicast filter applied in EVPN routing-instance blocks unexpected traffic. [PR1472511](#)

- With sensor being subscribed via Junos Telemetry Interface (JTI), after the interface is deleted/deactivated/disabled, the TCP connection is still established, and the CLI command of **show agent sensors** still shows the subscription. [PR1477790](#)
- EVPN aliasing and load-balancing for Layer 2 traffic does not work with Dynamic Link Next-Hop. EVPN alaising with DLNH for L2 traffic is not supported for Junos OS Release 20.3 and earlier releases. [PR1504412](#)
- RPM is Juniper propitiatory feature. In case of JUNOS RPM, RPM client never set the DF bit. Hence we didn't see this issue between JUNOS RPM client and JUNOS RPM server. Whereas in case of EVO, EVO RPM client is setting the DF bit while sending the RPM probes to RPM server. In case of JUNOS TVP based platforms, RPM server is not able to decode the DF bit properly. Issue is not applicable for non TVP based JUNOS platforms acting as RPM server. [PR1508127](#)

Routing Protocols

- Commit check fails when rib-sharding is configured with these statements:
 - **routing-instances <name> routing-options multipath**
 - **routing-instances <name> routing-options policy-multipath**
 - **routing-instances <name> protocols mvpn.**

Subscriber Management and Services

- Subscriber management and services are not supported on MPC10 or MPC11 line cards when you use these cards for subscriber access. MPC10 and MPC11 line cards support subscriber management and services only when you use these cards for uplink purposes to the core.

SEE ALSO

[What's New | 108](#)

[What's Changed | 131](#)

[Open Issues | 143](#)

[Resolved Issues | 156](#)

[Documentation Updates | 188](#)

[Migration, Upgrade, and Downgrade Instructions | 188](#)

Open Issues

IN THIS SECTION

- Class of Service (CoS) | 144
- EVPN | 144
- Flow-based and Packet-based Processing | 144
- Forwarding and Sampling | 144
- General Routing | 145
- High Availability (HA) and Resiliency | 150
- Infrastructure | 150
- Interfaces and Chassis | 150
- Intrusion Detection and Prevention (IDP) | 151
- Layer 2 Ethernet Services | 151
- MPLS | 151
- Network Management and Monitoring | 152
- Platform and Infrastructure | 152
- Routing Protocols | 153
- Subscriber Access Management | 155
- User Interface and Configuration | 155

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

Class of Service (CoS)

- If you have to take an interface out of AE bundle and configure it to operate in stand-alone mode, then doing this in a single commit may render the operation ineffective and could lead to connectivity issues. There is a known issue around this and this is seen due to a race condition between RE daemons (CoSD, DCD/Chassisd), Packet Forwarding Engine (PFE) and kernel. The below document link speaks of this issue when there is explicit CoS configuration to be made on the interface - https://www.juniper.net/documentation/en_US/junos/topics/concept/schedulers-cos-ae-sdh-limits-cos-config-guide.html. However, the problem can be seen without explicit CoS too as there is default CoS that is always present. In some cases, it is possible that a single shot commit will send out multiple operational messages down to kernel and might confuse the kernel to do unintended optimization that could lead to a message being consumed at kernel and not being sent to PFE. The result is the same even in this case. [PR1504287](#)

EVPN

- A few duplicate packets might be seen in an A/A EVPN scenario when the remote PE device sends a packet with an IM label due to MAC not learned on the remote PE device, but learned on the A/A local PE device. The nondesignated forwarder sends the IM-labeled encapsulated packet to the PE-CE interface after MAC lookup instead of dropping the packet, which causes duplicate packets to be seen on the CE side. [PR1245316](#)
- With Junos OS Release 19.3R1, VXLAN OAM host-bound packets are not throttled with DDoS policers. [PR1435228](#)

Flow-based and Packet-based Processing

- Use an antireplay window size of 512 for IPv4 or IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). Hence, there are no out-of-order packets with 512 antireplay window size. [PR1470637](#)

Forwarding and Sampling

- Packet length for ICMPv6 is shown as "0" in the output of the "show firewall log detail" CLI command. [PR1184624](#)
- When GRES is triggered by SSD hardware failure, the syslog error of **rpd[2191]: krt_flow_dfwd_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out** might be seen. Issue can be recovered by restarting the dfwd daemon. [PR1397171](#)
- After routing restarts, the remote mask, which indicates from which remote PE devices MAC IP addresses are learned, that the routing daemon sends might be different from the existing remote mask that the

Layer 2 learning daemon had before restart. This causes a mismatch between Layer 2 learning and routing daemons' interpretation as to where the MAC-IP address entries are learned, either local or remote, leading to the MAP-P table being out of synchronization. [PR1452990](#)

- During -ve scenario like clear MAC table, sometimes expected number of 512K MACs are not re-learnt. [PR1475205](#)
- **fast-lookup-filter** with match not supported in FLT Hardware might cause traffic drop. [PR1573350](#)

General Routing

- Some non-fatal interrupts (for example, CM cache or AQD interrupts) are logged as fatal interrupts. The following log messages will be shown on CM parity interrupt: fpc0 TQCHIP 0: CM parity Fatal interrupt, Interrupt status: 0x10 fpc0 CMSNG: Fatal ASIC error, chip TQ fpc0 TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180010 msec TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180005 msec [PR1089955](#)
- On MX series and PTX series with FPC-PTX-P1-A or FPC2-PTX-P1A, you might encounter a single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message gets reported: Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero Jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 The Junos OS chassis management error handling detects such a condition, raises an alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper support representative if the issue persists even after the FPC restarts. [PR1254415](#)
- If a vmhost snapshot is taken on an alternate disk and there is no further vmhost software image upgrade, the expectation is that if the current vmhost image gets corrupted, the system boots with the alternate disk so the user can recover the primary disk to restore the state. However, the host root file system and the node boots with the previous vmhost software instead of the alternate disk. [PR1281554](#)
- On 30-port MACsec-enabled line card (LC1101-M-30C, LC1101-M-30Q, and LC1101-M-96X) of the PTX10008 chassis, when the exclude-protocol lacp statement configured at the [edit security macsec connectivity-association connectivity-association-name] hierarchy level is deleted or deactivated, the LACP protocol's Mux State shown under the output of CLI command show lacp interface, might remain as attached or detached and might not change to distributing state. [PR1331412](#)
- Backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)

- On MX2010/MX2020 routers equipped with SFB2 (Switch Fabric Board 2), some error messages could be occasionally seen in the logs. There is no operational impact nor an indication of a real issue caused by these messages. [PR1363587](#)
- Due to transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. This change applies to all Platforms using Hybrid Memory Controller (HMC). [PR1384435](#)
- In scaled configurations an FPC can be shown as online and permits revert of a RLT when its actually not yet ready to pass traffic. This can lead to traffic loss for up to 8 minutes. [PR1394026](#)
- 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](#)
- If Hypertext Transfer Protocol (HTTP) Header Enrichment function is used, the traffic throughput decreases when traffic passes through Header Enrichment. [PR1420894](#)
- FPC might crash when Packet Forwarding Engine memory usage for a partition such as NH/DFW is high. Under low PFE memory condition log "Safety Pool below 25% Contig Free Space" or "Safety Pool below 50% Contig Free Space" might be observed. [PR1439012](#)
- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card. Because of vendor limitations, achieving subsecond hold timers is not possible. Juniper will continue to work with vendors to arrive at a solution for this. [PR1444516](#)
- IPv6 VRRP MAC address is not handled correctly by VFP (virtual forwarding plane). If the IPv6 traffic throughput is beyond the bandwidth of this slow path, the IPv6 packets might be dropped. [PR1449014](#)
- Physical interface policer are not supported in 19.3 for MPC11. [PR1452963](#)
- The CFM remote MEP is not coming up after configuration or remains in Start state. [PR1460555](#)
- Backport jemalloc profiling CLI support to all releases where jemalloc is present. [PR1463368](#)
- Either static routes (or) implicit filters should be configured for forwarding DNS traffic to service pic. It solves DNS packet looping issue. [PR1468398](#)
- For the MPC10E card line, the IS-IS and micro-BFD sessions do not come up during baseline. [PR1474146](#)
- Error messages "[Error] L2alm : l2alm_mac_process_hal_delete_msg:667 Ignoring MAC delete with ifl index 355, fwd_entry has 7888" after performing configuration removal/restore with IPMPLS configurations in MX480 box. [PR1475785](#)
- Possible out of order deletion of AftNode #012#012#012 AftNode details - AftIndirect token:230791 group:0 nodeMask:0xffffffffffffffff indirect:333988 hwInstall:1#012 "during baseline [PR1486158](#)
- NH learning knob is enabled by default in MPC10 and MPC11 irrespective of the knob configuration. The disabling will have no effect on the knob functionality. [PR1489121](#)
- High-scale login and logout (around 1M bearers) can prevent some sessions from logging in again. [PR1489665](#)

- On the MPC10 line card, the following error message is observed on RE1 after doing graceful switchover from RE0 to RE1: user.err aftd-trio: ([Error] L2ALIPC : L2AL IPC client failed to connect to l2ald). [PR1491384](#)
- This is rarely seen crash after doing multiple switchover (more than 50) with scale configuration. This is race scenario of accessing already free memory and causing the core file. [PR1491527](#)
- On MPC7E/8E/9E/10E, JNP10K-LC2101, and MX204/MX10003, an error syslog saying "unable to set line-side lane config (err 30)" will occasionally appear. This is non-service impacting and can be ignored. [PR1492162](#)
- The smart-sfp-present leaf was removed because this was redundant information. There is a leaf saying the type of smart sfp present on the interface. The present leaf was removed to avoid cluttering of the CLI output. [PR1492551](#)
- Scheduler ingress Packet Forwarding Engine VOQ drop counters do not match egress queue drop counters (diff > 100,000). [PR1494785](#)
- When running the command: show pfe filter hw filter-name <filter name> , the command fails to retrieve the Packet Forwarding Engine programming details of the filter. [PR1495712](#)
- After backup Routing Engine halt, CB1 goes offline and comes back online; this leads to the backup Routing Engine booting up, and it shows the reboot reason as "0x1:power cycle/failure." This issue is only for the RE reboot reason, and there is no other functional impact of this. [PR1497592](#)
- In routing-instance with table next-hop scenario (for example, if EVPN routing-instance is configured, the l2ald process creates a routing table and the EVPN adds a route pointing to this table as table next-hop in the rpd process), if the routing table created within the routing-instance is deleted and then re-added (e.g. deactivated and then re-activated the routing-instance) very fast before the rpd could delete the route pointing to the table next-hop, then the route in the rpd will end up using the staled table next-hop, hence resulting in traffic loss. Sampling configuration which delays the route deleting in the rpd increases the possibility of hitting the issue. [PR1498087](#)
- If MPLS is needed, the CRPD container must be instantiated with the MPLS modules already installed on the host. [PR1498632](#)
- The 'show dynamic-tunnels database statistics' show command is common for many different tunnel. to make common kind of infra, we are showing tunnel encap string. [PR1501576](#)
- SFB3 and MPC11 are not supported in Junos software version 19.4 releases. This change is to disable these components in Junos software version 19.4. [PR1503605](#)
- PR1463859 introduces a software defect that causes a 10GE interface to flap continuously when configuring with the WAN-PHY framing with the default "hold-down" timer (0). Once upgrading a router to an affected software release, the interface may flap continuously. This is not applicable to an interface with the default framing - LAN-PHY. [PR1508794](#)
- On MX platforms with MPC10E/MPC11E, traffic loss (i.e., including Unicast and Multicast traffic) could be seen if along with ECMP (Equal Cost MultiPath) scenario. This defect could only be observed in 20.1R1-S1 and 20.1R1-S2. [PR1513898](#)

- The log file to log the activities associated with the "request rift package activate" command is created with the permissions of the cli user. If multiple users run the command, it may fail due to problems with permissions writing to the log file. [PR1514046](#)
- 35 seconds delay is added in reboot time from 20.2R1 release compared to 19.4R2 release [PR1514364](#)
- Traffic drop observed when Multicast traffic on a group with 4000 egress AE ports is sent. The drop is always on the egress port that are on same PFE as ingress. PFE times-out before the multicast packet is processed and that causes the packet to drop. [PR1514646](#)
- In Junos OS 20.3R1 and 21.2R1: - It is possible that LFM may flap during MX Series Virtual Chassis unified ISSU to/from this release. [PR1516744](#)
- Using sensor path of /mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter/mpls-label "mpls-label" values are not getting reaped out when configured for SR-sid ingress sensors. [PR1516811](#)
- The **show configuration** command is not displaying the actual version info. [PR1517231](#)
- When an AMS ifd is configured for the first time or any member of the AMS bundle is removed or added, the PICs on which the members of AMS bundle are present go for a reboot. * There is a timer running in the AMS kernel which is used as a delay for the PIC reboot to complete and once that timer expires AMS assumes that the PICs might have been rebooted and it moves into next step of AMS fsm. * In scaled scenarios, this rebooting of the PIC is delayed due to DCD. This is because when a PIC goes down, DCD is supposed to delete the IFDs on that PIC and then the PIC reboot happens. * But DCD is busy processing the scaled config and the IFD deletion is delayed. This delay is much greater than the timer running in AMS kernel. * When the above timer expires, the FSM in AMS kernel wrongly assumes the PIC reboot would be completed by then, but the reboot is still pending. * By the time DCD deletes this IFD the AMS bundles are already UP. Because of this, there is a momentary flap of the bundles. [PR1521929](#)
- In MAC-OS platforms when client connects successfully, client does not get minimized to tray icon and it stays connected and needs to manually minimize it. [PR1525889](#)
- When global port-mirroring configured on DUT and queried for xml info, everything works as expected. When that global port-mirroring configuration is deactivated and xml is queried, display info is missing port-mirroring-instance info. [PR1529413](#)
- During GNF ISSU and if ISSU unsupported FRU is present, then such FRU is to be brought online manually once ISSU finishes. [PR1534225](#)
- Inconsistent core.python2.7.mpc0 core is seen with stacktrace @ea_wi_precl,@ea_macsec_receive() [PR1534568](#)
- In scaled MX2020 router, with vrf localisation enabled, 4 million nexthop scale, 800k route scale. FPCs may go offline on GRES. Post GRES, router continues to report many fabric related CM_ALARMS. FPC may continue to reboot and not come online. Rebooting master and backup Routing Engine will help recover and get router back into stable state. [PR1539305](#)

- A new alarm "network-service mode mismatch between configuration and kernel setting" was introduced by PR 1514840 commit. when unified ISSU is performed from images without PR1514840 commit to images with PR1514840 commit, then the transient false alarm will be seen. [PR1546002](#)
- In synce configuration, Config 1: ESMC transmit is configured Config 2: if deactivated chassis synchronization source configured OR no chassis synchronization source is configuring is active then commit error is given as "'esmc-transmit' requires 'chassis synchronization source' configuration". [PR1549051](#)
- In Junos OS Release 20.4, the return data from get_subscriber_info keyword contains string list instead of element list. [PR1560397](#)
- VE & CE mesh groups are default mesh groups created for a given Routing instance. On vlan/bridge-domain add, flood tokens and routes are created for both VE and CE mesh-group/flood-group. Ideally, VE mesh-group doesn't require on a CE router where IGMP is enabled on CE interfaces. Trinity based CE boxes have unlimited capacity of tokens, so this would not be a major issue. [PR1560588](#)
- On all Junos platforms, all traffic coming from remote end of a dynamic tunnel will not be processed and dropped when GRE based dynamic-tunnel is configured and tunnel preference is configured as 1. This issue is seen because of the missing programming of decapsulation mode in the internal system. [PR1561721](#)
- The problem is with L1 node (IFD) not reflecting correct BW configured for tunnel-services. When baseline configuration has 1G configuration on some fpc/pic in groups global chassis and if we override with local chassis tunnel-service 10G BW in scaled scenario. Out of 10Gbps BW configured only 1Gbps is allowed per 1G speed configured in baseline configuration. [PR1568414](#)
- In SAEGW-u mode, If volume based charging is configured with threshold and quota enabled - in some cases, the threshold report may not be sent to control peer (even after the threshold is reached). This issue is only applicable to 20.3 release, and is not applicable to later major releases starting 20.4. [PR1568563](#)
- [EVPN] flag, source and logical address are not expected in MAC address found in BD BD-3 instance. [PR1569546](#)
- Core files are found @ rpio_die ioctl_send rpio_tunnel_set_ioctl. [PR1569912](#)
- BUM traffic replication over VTEP is sending out more packets than expected and there seems to be a loop also in the topology. [PR1570689](#)
- PIM rib-group failure to add in vrf - PIM: ribgroup vrf not usable in this context; all RIBs are not in instance. [PR1574497](#)
- Native sensors does not work for ldp p2mp based sensor as protobuf definitions are not exported in junos telemetry package. Decoding of the stats data exported from router through udp will not be possible without protobuf definitions. [PR1577931](#)
- When SFP and Periodic ukernel thread fail to read SFP EEPROM data via I2C, threads will retry with 1 second sleep. But threads go to sleep state without yeilding CPU and this causes CPU hogging. Hog event will be logged to Syslog with the message: %PFE-0: fpc0 SCHED: Thread 12 (Periodic) aborted,

hogged 5030 ms %PFE-0: fpc0 SCHED: Thread 7 (SFP) aborted, hogged 5014 ms (Time following to "hogged" varies.) In case the total number of CPU hog event reaches 200 times, system assumes thread is in looping condition and reboots to repair looping condition. [PR1583281](#)

- Verification of filter counter statistics failed as received packets are doubled. [PR1590009](#)
- Expected snooping route is not observed after configuring one bridge with snooping & add interface check. [PR1590278](#)

High Availability (HA) and Resiliency

- If performing GRES with the interface em0 (or fpx0) disabled on the master RE, then enabling the interface on the new backup Routing Engine, it isn't able to access network. [PR1372087](#)

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](#)
- "IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) error message is observed continuously in AD with base configurations. [PR1485038](#)
- User while loading the kernel would see the message "GEOM: mmcsd0s.enh: corrupt or invalid GPT detected.". This message has no impact to functionality and will be resolved in a future release. [PR1549754](#)

Interfaces and Chassis

- 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](#)
- Changing framing modes on a CHE1T1 MIC between E1 and T1 on a MPC3E NG HQoS line card will cause the PIC to go offline. [PR1474449](#)
- The traffic (which is destined to the hosts behind static PPPoE subscriber's CPE device) drop is seen due to bad MPLS VPN label (which points to discard next-hop) after RE switchover without NSR. The traffic destined to the CPE device itself is not affected. [PR1488302](#)
- Input and output bytes count mismatch in the IPv6 traffic statistics while issuing the "show interface extensive" command. [PR1505100](#)
- mc-ae need to have "prefer-status-control-active" set to avoid flap on the split brain case (ICCP down or peer node reboot). Please configure it on the "status-control active" PE. [PR1505841](#)
- When configuring CFM sessions on MPC10 and MPC11 line cards, if syslog error "ppman: [Error] PPM:CTRL_CFM: PpmCtrlProtoCfm::getFcPlp: CFM interface is not found in intf table." is seen, the

CCM PDUs will not take the configured forwarding class. The CCMs will take forwarding class as "network-control", and queue as 3. [PR1527032](#)

- When configuring CFM sessions on MPC10 and MPC11 line cards, if syslog error "ppman: [Error] PPM:CTRL_CFM: PpmCtrlProtoCfm::getFcPlp: CFM interface is not found in intf table." is seen, the CCM PDUs will not take the configured forwarding class. The CCMs will take forwarding class as "network-control", and queue as 3. [PR1534239](#)

Intrusion Detection and Prevention (IDP)

- The CLI now provides helpful remarks about IDP's tunable detector parameters when executing the command "set security idp sensor-configuration detector protocol-name tunable-name." [PR1490436](#)

Layer 2 Ethernet Services

- On Juniper Networks Junos OS devices configured as a DHCP forwarder, the Juniper Networks Dynamic Host Configuration Protocol Daemon (jdhcp) process might crash when receiving a malformed DHCP packet. Refer to <https://kb.juniper.net/JSA11056> for more information. [PR1430874](#)
- The copying of files to the RCB over WAN ports is slow. This is observed across all platforms running Junos OS Evolved platforms. [PR1496895](#)

MPLS

- When some LSPs that request facility backup protection using bypass tunnels are brought up using respective Resv messages that do not contain the mandatory RECORD_ROUTE object, and when such LSPs undergo local repair, then RPD may core with the backtrace specified in this problem. If either the Resv messages originated by egress LERs contain the mandatory RECORD_ROUTE object or if such LSPs brought up with mal-formed Resv do not undergo local repair, then the core will not occur. [PR1560059](#)
- When ISIS-TE or OSPF-TE is enabled without admin-groups-extended-range/admin-groups-extended (which is configured under routing-options) or admin-group-extended configured, if receives the peer-router advertised the extended admin groups and then enable the config of admin-groups-extended-range/admin-groups-extended and admin-group-extended, some LSP with extended admin group constraints will fail to be established. [PR1575060](#)

Network Management and Monitoring

- Issue: Some linkUp traps are missing for physical interfaces. Cause: In the current scenario, one of the FPC is having a scale of 15K IFLs and restarting this FPC triggers a large number of SNMP traps to get generated in a very short time. Since snmp's trap threshold is 500 and the maximum throttle queue size by default is 20K, within a short period, this burst of traps causes the throttle queue to get filled up and snmpd starts dropping the trap packets. So, for higher scale systems, it is advisable to provide a larger threshold and queue size for throttling traps, so that they don't get dropped. Recommendation: In case of such scale, please increase the throttle-threshold size and max queue-size limits using the following config so that drops won't happen. 1. "set snmp trap-options throttle-threshold 10000" 2. "set snmp trap-options queue-size 40000" [PR1507780](#)
- If a node is a 'deviate not-supported' in a Yang model and when that module is installed on a Junos device, Junos device shows (if that statement is configured) " ## Warning: 'knob' is deprecated But this does not convey the right meaning. So as part of this PR the warning message is changed to 'statement ignored unsupported platform'. Example warning message before this PR fix:

user@router# show test:system bar-system a; ## Warning: 'bar-system' is deprecated {master}[edit]

Example warning message post this PR fix: user@router# show test:system ? Possible completions:<[Enter]> Execute this command host-name Leaf host-name test-grouping-leaf Test test-grouping| Pipe through a command [edit] user@router# show test:system#### Warning: statement ignored: unsupported platform (mx960)##bar-system a; [edit] [PR1516910](#)
- With Junos OS Release 20.3R1, "Traffic statistics" in the show interface CLI command is displaying incorrect cumulative values. [PR1539483](#)
- Issue: show snmp mib walk alarmModelTable fails Cause: Issue in re-reading the "snmp alarm-management" set of configuration. [PR1566597](#)

Platform and Infrastructure

- Sometimes OSPF flapping occurs during unified ISSU from Junos OS Release 16.2R2 to Release 17.2R3. [PR1371879](#)
- On MX Series platforms with MPC7/8/9 or MX-204/MX-10003 when the packets which exceed the MTU and whose DF-bit is set go into a tunnel (such as GRE, LT), they might be dropped in the tunnel egress queue. [PR1386350](#)
- The traps are the result of PPE commands injected from the host. One possible reason could be Layer 2 BD code, which is trying to decrement BD MAC count in the data plane. It is unlikely that there is a packet loss during this condition. This could happen during ISSU and this may be due to a problem with ISSU counter morphing used for LU-based cards, where certain counters are not disabled or disabled too late during ISSU. [PR1426438](#)
- A few OAM sessions are not established with scaled EVPN ETREE and CFM configurations. [PR1478875](#)

- Due to software implementation firewall filter is re-applied duration graceful switchover (GRES). This may lead to short during when filter is not applied provoking side effects like drop of traffic. [PR1487937](#)
- On all Junos OS platforms that support EVPN-MPLS or EVPN-VXLAN, when an existing ESI interface flaps or is newly added to the configuration, sometimes DF (Designated Forwarder) election happens before the local bias feature is enabled and during this time, existing broadcast, unknown unicast, and multicast (BUM) traffic might be looped for a short time duration (less than several seconds). [PR1493650](#)
- On MX platform running with enhance IP mode or enhanced ethernet mode, also, Operation, Administration and Maintenance (OAM) is enabled with Periodic Packet Management (PPM) mode by default, maintenance association end point (MEP) session might be failed to create. In the end, network connection failure might not be efficiently monitored. This is functional impact. [PR1506861](#)
- On vMX, the blockpointer in the ktree is getting corrupted leading to core-file generation. There is no function impact such as fpc restart or system down and the issue is not seen in hardware setups. [PR1525594](#)
- If you use the source-address NTP configuration parameter and issue the command "set ntp date" from the CLI, packets will be sent with the source address of the outgoing interface rather than the manually configured IP address. Typically the manually configured IP address would be a loopback address. The problem does not apply to automatically generated NTP poll packets. [PR1545022](#)
- This issue might be seen only in back to back GRES in about more than 40 to 50 iterations. No workaround available and FPC gets restarted. [PR1579182](#)
- The issue is due to output byte count not getting updated properly. The script logs shows that there is no packet loss, There is no functional impact and will be taken up in the upcoming releases. [PR1579797](#)

Routing Protocols

- While interoperating with other vendors in a draft-rosen multicast VPN, by default Junos OS attaches a route target to multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities are prevented from propagating if the BGP route-target filtering is enabled on the device running Junos OS. [PR993870](#)
- Starting in Junos OS Release 13.3R2 and later, if delegated BFD sessions flap continuously, packet buffer memory may be leaked. The automatic memory leak detection process reports this within the syslog once a certain threshold is reached, like "fpc7 SHEAF: possible leak, ID 8 (packet(clones)) (10242/128/1024)" on MX- MPC or "fpc4 SHEAF: possible leak, ID 9 (packet(clones)) (255/1/5)" on other platforms. Note that BFD sessions operating in centralized mode are not exposed. A complete fix is available from Junos OS Release 14.2R1 and later. Prior to Junos OS Release 14.2R1, there was only a partial fix that did not fix the memory leak completely. [PR1003991](#)

- 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 OSPF are in synchronization state because the IGP interface is down with ldp-synchronization enabled for OSPF. user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity. As per the current analysis, the IGP interface goes down because although LDP notified OSPF that LDP synchronization was achieved, OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up yet. [PR1256434](#)
- In rare cases, RIP replication might fail as a result of performing NSR Routing Engine switchovers when the system is not NSR ready. [PR1310149](#)
- On MX platforms, unexpected log message will appear if the CLI command 'show version detail' or 'request support information' is executed: test@test> show version detail *** messages *** Oct 12 12:11:48.406 re0 mcsnoopd: INFO: krt mode is 1 Oct 12 12:11:48.406 re0 mcsnoopd: JUNOS SYNC private vectors set [PR1315429](#)
- SCP command with Routing Instance (-JU) is not supported. [PR1364825](#)
- Bfd session flaps during ISSU only in mpc7e card (Bfd sessions from other cards of DUT to peer routers did not flap during ISSU). Issue is not seen frequently. [PR1453705](#)
- The virtual-router option is not supported under routing-instance in lean rpd image. [PR1494029](#)
- On all Junos platforms with scaling MVPN scenario, some PIM Join/Prune messages may not be processed for the first attempt. For instance, a dedicated PIM router receives more than 2500 PIM hello packets from the new neighbors, followed by PIM Join packets for the same multicast group in a very short period of time. [PR1500125](#)
- In a Layer 3 VPN scenario, the rpd (routing protocol process) on the backup Routing Engine might crash when BGP (standby) received a VPN route from the peer that is rejected due to invalid target community and because the BGP standby peer synchronization is not complete yet. [PR1508888](#)
- TILFA backup path fails to install in LAN scenario and also breaks SR-MPLS tilfa for lan with more than four end-x sids configured per interface. [PR1512174](#)
- On the devices with NG-RE (Next Generation Routing Engine) and SCBE2 (Enhanced Switch Control Board), when BFD authentication for BGP is enabled, the BFD may flap after the NG-RE switchover. The switchover should be GRES or NSR switchover. After the flap, the device could be self recovery. [PR1522261](#)
- Observing Packet loss when primary link is enabled. [PR1592884](#)

Subscriber Access Management

- There is configuration option to set slow interval to 1 minute. By this approach, pfe uses the slow interim which is set to 1 min. [PR1515899](#)

User Interface and Configuration

- On Juniper device running Junos OS Evolved, NETCONF Service over SSH with dedicated TCP port (It is configured with 'system services netconf ssh' and the default port is 830) might not work if in-band management is used (i.e. connection is established via network interface or loopback interface etc.). [PR1517160](#)
- This PR removes Adobe Flash dependent elements from J-Web for EX and MX Series devices. For MX Series Routers, J-Web previously contained Flash elements on the following pages: Monitor-Interfaces Monitor-System view-Process Details Monitor-Routing-OSPF Information For EX Series Switches, J-Web previously contained Flash elements on the following pages: Monitor-Interfaces Monitor-System view-Process Details Monitor-Switching-IGMP Snooping Monitor-Virtual Chassis Monitor-POE Monitor-Security-Port Security Monitor-Routing-OSPF Information Monitor-Service-DHCP -Server Monitor-Service-DHCP -Relay. [PR1553176](#)
- On vmhost platforms, if the xml output from command "request vmhost mode test | display xml rpc" is picked and used in netconf, it will fail. set vmhost mode custom test layer-3-infrastructure cpu count MIN set vmhost mode custom test layer-3-infrastructure memory size MIN set vmhost mode custom test nfv-back-plane cpu count MIN set vmhost mode custom test nfv-back-plane memory size MIN set vmhost mode custom test vnf cpu count MIN. [PR1559786](#)
- There is problem with one corner case, when Tenant_Systems:- "juniper.conf.gz" file creates with empty data when we create Tenant System. [PR1584850](#)

SEE ALSO

[What's New | 108](#)

[What's Changed | 131](#)

[Known Limitations | 137](#)

[Resolved Issues | 156](#)

[Documentation Updates | 188](#)

[Migration, Upgrade, and Downgrade Instructions | 188](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 156](#)
- [Resolved Issues: 20.3R2 | 165](#)
- [Resolved Issues: 20.3R1 | 174](#)

This section lists the issues fixed in Junos OS Release 20.3R3 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: 20.3R3

Class of Service (CoS)

- On the MPC7E line card, the BPS counter of the egress queue displays wrong BPS value when the cell mode is configured on the static interface. [PR1568192](#)

EVPN

- The rpd process might crash after adding route-target on a dual-Routing Engine system under the EVPN multihoming scenario. [PR1546992](#)
- The multicast traffic loss might be seen in EVPN-VXLAN scenario with CRB multicast snooping. [PR1570883](#)
- The rpd might crash if EVPN routing instances or BGP connections are in flapping. [PR1581674](#)
- Multicast traffic loss might be seen in EVPN setup with IGMP snooping used. [PR1582134](#)
- The BUM traffic might lose after triggering NSR in EVPN-MPLS or EVPN-ETREE scenario. [PR1586402](#)

Forwarding and Sampling

- The configuration archive transfer-on-commit fails when running Junos OS Release 18.2R3-S6.5. [PR1563641](#)
- In the VXLAN scenario, the locally originated packets have UDP source port 0. [PR1571970](#)

General Routing

- Unable to show to which shard a given route is hashed. [PR1430460](#)
- The following error message is observed after GRES: [user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, agglfName:ps1.0 memberIfName:lt-3/0/0.32767]. [PR1466531](#)
- "re_tvp_builtinfo_update: Unable to get firmware version" message is seen in chassisd. [PR1471938](#)
- The following line card errors are seen: HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invlaid IRB topo/ IRB ifl zero in l2 nh 40495 add IRB. [PR1472222](#)
- MX204 || Summit || Incorrect log message for PIC1 when changing the config from PIC mode to Port Mode. [PR1500429](#)
- MX150 might go into db mode after software upgrade or downgrade. [PR1510892](#)
- Sometimes external 1 pps cTE is slightly above Class B requirement of the ITU-T G.8273.2 specification. [PR1514066](#)
- On the MX960 routers, the show interfaces redundancy rlt0 statement shows current status as primary down as FPC is still in the Ready state after rlt failover (restart FPC). [PR1518543](#)
- The l2cpd process might crash when removing LLDP on an aggregated Ethernet interface. [PR1528856](#)
- On the MX150 routers, configuring the no-flow-control statement under gigether-options does not work. [PR1531983](#)
- Wavelength unlocked alarm is On when using SFP+-10G-T-DWDM-ZR optics. [PR1532593](#)
- The dcpfe process might crash and cause FPC to restart due to the traffic burst. [PR1534340](#)
- BGP SR-TE IPv6 routes might get hidden after the chassisd restarts. [PR1534511](#)
- The BFD neighborship fails with the EVPN_VXLAN configuration after the Layer 2 learning restarts. [PR1538600](#)
- The JNH memory leak could be observed on MPCs or MICs. [PR1542882](#)
- The kmd process might crash when the interface flaps. [PR1544800](#)
- Continuous rpd errors might be seen and new routes fails to be programmed by the rpd process. [PR1545463](#)
- FPC(s) may not boot-up on MX960/EX9214 in a certain condition. [PR1545838](#)
- The performance of PFE process on MX204/SRX4600/EX9251 might be degraded after Junos OS Release 19.3R1. [PR1545989](#)
- Backup Routing Engine vmcore might be seen due to the absence of the next-hop acknowledgement infra. [PR1547164](#)
- FPC crash may occur after flapping the multicast traffic. [PR1548972](#)
- PKI CMPv2 client certificate enrolment does not work on SRX when using root-CA. [PR1549954](#)

- The adapted sample rate might get reset to the configured sample rate without changing the sampling rate information in sFlow datagrams after enabling sFlow technology on a new interface. [PR1550603](#)
- The interface might not come up with 1G optics. [PR1554098](#)
- The subscriber sessions might be missed but stay in the authd after performing ISSU. [PR1554539](#)
- The chassisd may crash with repeated configuration commits on MX204/MX10003 platform. [PR1555271](#)
- The following message is not generated on the MPC11E line card due to no power: Chassisd SNMP trap Fru Offline. [PR1556090](#)
- The framed route installed for a demux Interface has no MAC address. [PR1556980](#)
- On the EX4300 device, script fails while committing the IPsec authentication configuration as the algorithm statement is missing. [PR1557216](#)
- Packets corruption on 100G or 40G interface are configured with protocol PTP. [PR1557758](#)
- The MAC addresses learned in a Virtual Chassis may fail aging out in MAC scaling environment. [PR1558128](#)
- RPD core files are seen after Routing Engine switchover. [PR1558814](#)
- On the MX150 routers, the following continuous license error is observed:
[licinfra_set_usage_nextgen_async:1733] Invalid input parameters. [PR1559361](#)
- The request system software validate command might corrupt installation of the junos-openconfig package. [PR1560234](#)
- Interface not able to send/receive packets after repeated link flaps on MPC10/11E. [PR1560772](#)
- The l2cpd process might generate a core file on reboot. [PR1561235](#)
- Client authentication is failing after performing GRES. [PR1563431](#)
- The following error message might be seen after ISSU: Turbotx process not running. [PR1564418](#)
- MX platforms with MX-SCBE3 might reboot continuously. [PR1564539](#)
- Commit error observed when tunnel-service is configured on a PIC without explicit bandwidth. [PR1565034](#)
- On the MX2010 or MX2020 routers, the following error message might be observed after switchover with GRES/NSR: CHASSISD_IPC_FLUSH_ERROR. [PR1565223](#)
- PPPoE service-name-tables does not correctly count active sessions that matches agent-specifier aci/ari used for delay. [PR1565258](#)
- The KRT log file might continue to grow after removing the KRT log configuration. [PR1565425](#)
- The active DHCP subscribers might not get synchronized to backup BNG. [PR1567735](#)
- On the MX204 routers, FPC might display high CPU utilization because of the JGCI background thread that runs for a long period. [PR1567797](#)

- On the MX150 routers, the request system software add command is disabled in Junos OS Release 19.4R3-S1, 20.1R2, and 20.4R1. [PR1568273](#)
- The nsd might crash after turning off the address translation for the NAT rules in the USF scenario. [PR1568997](#)
- The RPD process might crash while using BFD API to bring up the BFD sessions. [PR1569040](#)
- Traffic loss might be observed when SCU accounting is configured and logical-systems is enabled. [PR1569047](#)
- The agent sensor __default_fabric_sensor__ are partly applied to some FPCs, which causes zero payload issue AGENTD received empty payload for pfe sensor __default_fabric_sensor__. [PR1569167](#)
- Wi-Fi mPIM on SRX Series devices is reaching out to NTP and DNS servers. [PR1569680](#)
- The MPLS traffic passed through the back-to-back PE topology might match the wrong CoS queue. [PR1569715](#)
- The mspmand process might crash if the packet flow-control issue occurs on MS-MPC/MS-MIC. [PR1569894](#)
- The log message "/tmp//mpci_info: No such file or directory :error[1]" might be seen on VM Host platform. [PR1570135](#)
- Improve handling deletion of static demux interface with active subscribers. [PR1570739](#)
- The TFEB/FPC may fail to be online after rebooting the system or the FPC if interface-set is configured for CoS. [PR1572348](#)
- On the MX960 routers, the Require a Fan Tray upgrade alarm is raised when the top Fan Tray 0 is removed, even though the enhanced Fan Tray is already used. [PR1572778](#)
- CFP unplugged message is not logged in Junos OS Release 17.3 and later. [PR1573209](#)
- Fabric errors are observed and FPC processes might get offline when MPC3-NG/MPC3E cards are installed along with MPC7/MPC10 and SCBE3/SCB4 operating in increased-bandwidth fabric mode. [PR1573360](#)
- The rpd process on the transit node might crash when MPLS traceroute on the ingress node is performed. [PR1573517](#)
- Only root user is allowed to execute commands on host using vhclicent. [PR1574240](#)
- Slow FPC heap memory leak might be triggered by flapping the subscribers terminated over multiple pseudowires. [PR1574383](#)
- On the EA-based cards igmp group membership is displayed incorrectly. [PR1575031](#)
- On the MX150 routers, the interface might take a long time to power down while rebooting, powering-off, halting, or upgrading. [PR1575328](#)

- The show services service-sets statistics syslog command returns the following error message as the service-set does not have the syslog configuration: error: usp_ipc_client_rcv_ 1237: ipc_pipe_read fails! error:No error: 0(0), tries:1. [PR1576044](#)
- On the MX10016 routers, when the Fan X Failed alarm is cleared in the Fan Tray 1, the Fan/Blower OK SNMP traps are generated for the Fan Tray 0 [Fan 31 - 41] and Fan Tray 1 [Fan 11 - 41]. [PR1576521](#)
- The LLDP neighbor information displays hex string instead of chassis ID when subtype 1 is used. [PR1576721](#)
- Commit failure-error: Modified IFD "ae0" is in use by targeted BBE subscriber, commit denied - mtu config changed (1522), (1514). [PR1577007](#)
- When line card is booted on RE1 being Master, Nextgen stats failed to fetch the value of backup mac address correctly. [PR1577611](#)
- High FPC CPU usage might be seen when signal on the link is unstable. [PR1579173](#)
- On the MPC11E line cards, system resource monitor does not list some of the available Packet Forwarding Engines. [PR1579975](#)
- MX Series Virtual Chassis: gRPC based /components/ sensor output is missing lot of data. [PR1580120](#)
- When having analyzers mapped to channelized port then the mirror may not happen properly. [PR1580473](#)
- Kernel issue is observed in telemetry when the set services analytics streaming-server <> <> configuration is present and server is not reachable. [PR1581192](#)
- Hitting with vmcore.0 at 0xffffffff80443eef in kern_reboot. [PR1581260](#)
- The rpd process might crash on the new master after performing graceful switchover. [PR1581878](#)
- Changing bandwidth statement does not take affect for SNMP ifHigSpeed oid until a PSX interface is disable/enabled. [PR1582060](#)
- The process rpd may stuck in 100% due to race condition. [PR1582226](#)
- bbe-smgd process on both routing engines may crash due to a rare timing issue after logout of subscribers over pseudowire. [PR1582356](#)
- Node locked license addition fails in EVO. [PR1582704](#)
- Configuring or removing "hierarchical-scheduler" or "per-unit-scheduler" might cause traffic to stop forwarding. [PR1582724](#)
- SNMP SysObjectID.0 is empty with enabled unified-services. [PR1583534](#)
- Traffic might not get filtered properly when security-intelligence profile is configured on the MX platforms. [PR1584377](#)
- The rpd might crash due to a rare timing issue if both BGP Local-RIB and Adjacency-RIB-In route monitoring are enabled in BMP. [PR1584560](#)
- Bridge domain names information is not displayed properly in "show bridge statistics instance". [PR1584874](#)

- After changing configuration, "show bridge statistics" shows extreme larger value. [PR1584876](#)
- QFX5120-48Y (Stage 1), QFX5110-32Q and QFX5110-48S (Stage 2) | Allow default license for FBF, CFM, VRRP, QINQ, MC_LAG, TIMING, IGMP, PIM, GRE_TUNNEL, RIP, OSPF, VC, SFLOW. [PR1589920](#)
- Any mmcq based services might crash due to shared memory queues issue happens in a rare condition. [PR1592889](#)
- The TCP keepalive might not be processed by the private network host. [PR1593226](#)
- MX5/MX40/MX80 TEB stuck in present state. [PR1595107](#)
- CGNAT MX SPC3 AMS warm-standby 1:1 redundancy problem with CLI CPU statistics lost data after PIC failover. [PR1596976](#)

Interfaces and Chassis

- Block duplicate IP across different ifls inside same routing instance. [PR1555861](#)
- The input errors counter command on the monitor interface command does not work. [PR1561065](#)
- MAC address entry issue might be observed after the MC-LAG interface. [PR1562535](#)
- Traffic loss might be seen while verifying VRRP State Machine functionality. [PR1564551](#)
- JDI-RCT:M/Mx: not able to set member-id as Routing Engine is in synching mode forever when its having invalid Virtual Chassis data(error: Command aborted. VC configuration synch to backup Routing Engine in progress, try after 120 seconds.) [PR1569556](#)
- If-media-type missing from interface XML output on MX platforms. [PR1574035](#)
- The following errors are generated during GRES: VRRPMAN_PATRICIA_GROUP_ADD_FAIL:
vrrp_ifcm_send_bulk: Failed to add group to patricia tree key and VRRPMAN_ENTRY_KEY_PRESENT:
vrrp_ifcm_send_bulk: Already an entry present with the key. [PR1575689](#)
- MC-AE interfaces may go down if same VRRP group-id is configured on multiple IRB units. [PR1575779](#)
- Configuration for ppp NCP Max-failure number of retry count. [PR1584168](#)
- Unable to configure pseudowire interface on an MX10003 in virtual chassis mode. [PR1587499](#)

Layer 2 Ethernet Services

- AE interface flap might be seen during NSSU. [PR1551925](#)
- DHCP packet drop might be seen when the DHCP relay is configured on a leaf device. [PR1554992](#)
- Junos OS: Receipt of malformed DHCPv6 packets causes jdhcpd to crash and restart. (CVE-2021-0240) [PR1564434](#)
- The jnxJdhcpLocalServerMacAddress (.1.3.6.1.4.1.2636.3.61.61.1.4.3) returns incorrect format of the MAC address. [PR1565540](#)

- The Option 82 information is incorrectly cleared by the DHCP Relay agent. [PR1568344](#)
- JDHCPd doesn't response to any DISCOVER for 30min when it is in "clients waiting to be restored" state. [PR1592552](#)

Multicast

- FPC might crash in a multicast scenario. [PR1569957](#)

Network Management and Monitoring

- The mib2d process crashes and generates a core dump on backup Routing Engine. [PR1557384](#)
- SSH connection might become unresponsive and logs show "kern.maxfiles limit exceeded by uid" messages. [PR1567634](#)

Platform and Infrastructure

- Interwork failure between Junos OS Evolved as RPM client and TVP platforms as RPM server (and vice versa). [PR1508127](#)
- Packet Forwarding Engine errors or traps might be observed in the Layer 2 flooding scenarios. [PR1533767](#)
- The npc process generates the core file in igmp_process_wakeup_events,igmp_pfe_thread,thread_detach_tty. [PR1534542](#)
- CoS queue egress interface forwarding-class might not work as expected. [PR1538286](#)
- The following major error message might cause the Packet Forwarding Engine(s) to disable: XQ_CMERROR_SCHED_L3_PERR_ERR. [PR1538960](#)
- In rare occurrence Routing Engine kernel might crash while handling TCP sessions if GRES/NSR are enabled. [PR1546615](#)
- The kernel might crash if GRES is performed on either new iteration or after swapping the Routing Engine and restoring the HA configuration. [PR1549656](#)
- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)
- Traffic is not forwarded over IRB to a Layer 2 circuit on the lt interfaces. [PR1554908](#)
- The BUM frame might be duplicated on an aggregate device if the extended-port on the satellite device is an aggregated Ethernet interface. [PR1560788](#)
- The DHCPv4 request packets might be wrongly dropped when DDoS attack occurs. [PR1562474](#)
- The enforce-strict-scale-limit-license configuration enforces subscriber license incorrectly in the ESSM subscriber scenario. [PR1563975](#)
- "Last flapped" timestamp for interface fxp0 gets reset every time "monitor traffic interface fxp0" is executed. [PR1564323](#)
- PFEX might crash when soft error recovery feature is enabled on Packet Forwarding Engine. [PR1567515](#)

- pfe err-jnh_physmem_add_resvd_to_cntr(18014): PFE 0 jnh_app 0x08020860, add 0x00080000 from 0x00b00000-0x00b80000 to baMask 0x1. [PR1570631](#)
- On all EX9200 platforms with EVPN-VXLAN configured, the next-hop memory leak in MX Series ASIC happens whenever there is a route churn for remote MAC-IP entries learned bound to the IRB interface in EVPN-VXLAN routing instance. When the ASIC's next-hop memory partition exhausted the FPC might reboot. [PR1571439](#)
- Scale-subscriber license might be not updated properly on the backup RE which leads to "License grace period for feature scale-subscriber(44) is about to expire" alarm after GRES. [PR1573289](#)
- cassxr_err_addr(8593): Uninitialized Read Error @ EDMEM[0x7cb601b0. [PR1573920](#)
- Firewall filter configuration or modification might fail on MX platforms. [PR1586817](#)
- The traffic might not failover with shared-bandwidth-policer enabled on AE. [PR1588708](#)

Routing Policy and Firewall Filters

- Global variable policy_db_type is not set to the correct value on failure. [PR1561931](#)
- Generated route goes to the Hidden state when the protect core command is enabled. [PR1562867](#)
- bbe-smgd - dynamic-profile NACK due to config error reading address mask prefix-length in policy-options/policy-statement. [PR1583535](#)

Routing Protocols

- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. [PR1514966](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- Continuous rpd crash might be observed if a static group is added to protocol PIM. [PR1542573](#)
- Multipath information is displayed for BGP route even after disabling the interface for one path. [PR1557604](#)
- BGP LU session flap might be seen with the AIGP used scenario. [PR1558102](#)
- When admin-color based policy evaluation happens with the policy lfa configuration, the backup next hop chosen (among the different backup next hops possible) might not be correct. [PR1558581](#)
- All the Layer 3 VPN route resets when a VRF is added or removed. [PR1560827](#)
- Duplicate LSP nexthop is shown on inet.0, inet.3 and mpls.0 route table when ospf traffic-engineering shortcuts and mpls bgp-igp-both-ribs are enabled. [PR1561207](#)
- The rpd process might crash when there is BGP session re-establishing or flapping. [PR1567182](#)
- Traffic might be lost during mirror data transmit from the primary ppmr or bfdd. [PR1570228](#)
- There might be 10 seconds delay to upload the LSP on the point-to-point interface if rpd is restarted on its direct neighbor. [PR1571395](#)

- SNMP MIB ospfv3NbrState is returning drifted value. [PR1571473](#)
- The BFD session of DHCP subscriber does not come up on the MPC2E card and gets stuck in the "Down" state. [PR1572577](#)
- The DHCP packets might be dropped by the QFX5000 in the Static VXLAN scenario. [PR1576168](#)
- BGP session flap might be observed after the Routing Engine switchovers when the VRRP virtual address is used as the local address for the BGP session. [PR1576959](#)
- The dcpfe process might crash when any interface flaps. [PR1579736](#)
- BGP replication might be stuck in rare and timing conditions. [PR1581578](#)
- With IGMP snooping implemented, there is unexpected jitter issue that could cause traffic loss. [PR1583207](#)
- On rare occasion, RPD core may be observed on backup RE after loading a new image. [PR1583630](#)
- The rpd process might crash after committing with the configured static group 224.0.0.0. [PR1586631](#)
- The routing process may crash due to memory corruption while processing BGP multipath route. [PR1594626](#)
- The rpd process might crash when executing the SNMP get command to fetch the MPLS L3VPN MIBs. [PR1594664](#)

Services Applications

- IWF AVP value may not be reflected properly on LTS. [PR1581096](#)
- The "show services l2tp tunnel extensive", "show services l2tp session extensive" and "show subscribers accounting-statistics" commands do not work on LTS. [PR1596972](#)

User Interface and Configuration

- Commit might fail after the Routing Engine switchovers. [PR1531415](#)
- Apply-paths might cause validation failures during Junos upgrade. [PR1577626](#)

VPNs

- Type7 messages may not be sent from egress PE resulting in Type 3/5 messages not created for some S, Gs in source PEs. [PR1567584](#)
- The rpd might crash in the NG-MVPN scenario on all Junos/Evo platforms. [PR1579963](#)
- The ddos-protection reason "packets failed the multicast RPF check" may be seen in NG-MVPN scenario with GRE transport. [PR1591228](#)

Resolved Issues: 20.3R2

EVPN

- With dynamic list next hop configured, a forwarding problem occurs after performing graceful switchover. [PR1513759](#)
- no-arp-suppression is required for MAC learning across the EVPN domain on the static VTEP. [PR1517591](#)
- ARP table might not be updated after performing VMotion or a network loop. [PR1521526](#)
- The BUM traffic might get dropped in the EVPN-VXLAN setup. [PR1525888](#)
- The rpd might crash when auto-service-id is configured in EVPN VPWS scenario. [PR1530991](#)
- The route table shows additional paths for the same EVPN or VXLAN type 5 destination after upgrading from Junos OS Release 18.4R2-S3 to Junos OS Release 19.4R1-S2. [PR1534021](#)
- All the ARP reply packets toward some address are flooded across the entire fabric. [PR1535515](#)
- The GE LOS alarm logs on the change in IFF_CCCDOWN are not logged in the syslog message file. [PR1539146](#)
- Rpd memory leak might occur when changing EVPN configuration. [PR1540788](#)
- The L2ALD process might core-file when changing EVPN/VXLAN configuration. [PR1541904](#)
- The rpd crash might be seen after adding route-target on a dual-RE system under EVPN multihoming scenario. [PR1546992](#)
- VLAN ID information is missed while installing the EVPN route from the BGP Type 2 Route after modifying a routing-instance from instance-type EVPN to instance-type virtual-switch. [PR1547275](#)

Forwarding and Sampling

- The DHCP subscribers might get stuck in terminated state for around 5 minutes after disabling cascade ports. [PR1505409](#)
- The srrd process might crash in a high route churns scenario or if the process flaps. [PR1517646](#)
- The commit might fail if a filter enabled with enhanced-mode to et- interface is configured. [PR1524836](#)
- The l2ald process might crash when a device configuration flaps frequently. [PR1529706](#)
- VLAN-ID based firewall match conditions might not work for the VPLS service. [PR1542092](#)
- MAC learning issue might happen when EVPN-VXLAN is enabled. [PR1546631](#)
- All traffic would be dropped on AE bundle without VLAN configuration if bandwidth-percent policer is configured. [PR1547184](#)
- l2ald might crash due to next-hop issue in the EVPN-MPLS. [PR1548124](#)

General Routing

- Family IPv6 is not coming up for L2TP subscriber when additional attributes are not passed in the Framed-IPv6-Route VSA. [PR1526934](#)
- DHCP discover packet might be dropped if DHCP inform packet is received first. [PR1542400](#)
- The show dynamic-profile session client-id command displays only one IPv6 framed-route information. [PR1555476](#)
- In some MX Series deployments running Junos OS, the following random syslog messages are observed for FPCs: fpcx ppe_img_ucose_redistribute Failed to evict needed instr to GUMEM - xxx left. These messages might not have a service impact. These messages are addressed as INFO level messages. On a Packet Forwarding Engine, there are dedicated UMEM and shared GUMEM memory blocks. This informational message indicates some evicting events between UMEM and GUMEM and can be safely ignored. [PR1298161](#)
- The max-drop-flows statement is not available. [PR1375466](#)
- Need to be able to show which shard a given route is hashed to. [PR1430460](#)
- The MPC2E-NG or MPC3E-NG card with specific MIC might crash after a high rate of interface flaps. [PR1463859](#)
- Dynamic SR-TE tunnels do not get automatically recreated at the new master Routing Engine after the Routing Engine switchover. [PR1474397](#)
- Traffic decreases during throughput testing. [PR1483100](#)
- SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. [PR1484322](#)
- The AMS bundle might remain inactive when adding member interface to AMS bundle with scaled service sets. [PR1489607](#)
- The following error messages are observed on the MPC card in the manual mode:
clk_sync_as_evaluate_sync_ref: 362 - Failed to configure clk. [PR1490138](#)
- Some of the virtual services might not up after GRES or rpd restart. [PR1499655](#)
- Prefix is not emitted for the te-lsp-timers/state/cleanup-delay sensor path for OCST. [PR1500690](#)
- Transit v4 traffic forwarding over BGP SR-TE might not work. [PR1505592](#)
- Errors on vjunos0 Regarding TSensor related to PR 1362108. [PR1508580](#)
- Not able to forward traffic to VCP FPC after the MX Virtual Chassis reboots, FPC reboots, or adding VCP link. [PR1514583](#)
- On the MX2020 and MX2010 routers, the SPMB CPU is elevated when an SFB3 is installed. [PR1516287](#)
- The l2cpd might crash if the ERP is deleted after the switchover. [PR1517458](#)
- On the MX960 routers, the show interfaces redundancy rlt0 statement shows current status as primary down as FPC is still in the Ready state after rlt failover (restart FPC). [PR1518543](#)

- Junos OS: Command injection vulnerability in 'request system software' CLI command (CVE-2021-0219). [PR1519337](#)
- Traffic loss might happen when an Uncorrected (Fatal) AER error is detected. [PR1519530](#)
- During an upgrade, vSRX3.0 would display the following incorrect license warnings when utilizing licensable features even if the license was present on the device: such as warning: requires 'idp-sig' license. [PR1519672](#)
- The BFD session status remains down at non-anchor FPC even though bfd session is up after anchor FPC reboot/panic. [PR1523537](#)
- PSM firmware upgrade should not allow multiple PSM upgrade in parallel to avoid the firmware corruption and support multiple firmwares for different hardware Revs. [PR1524338](#)
- No response from the other routing engine for the last 2 seconds" triggers "SNMP trap generated: Fru Offline" messages. [PR1524390](#)
- Commit is successful while deactivating CB0 or CB1 interfaces with GNF. [PR1524766](#)
- Problem With static VLAN deletion with active subscribers and the FPC might be stuck at Ready state during restart. [PR1525036](#)
- The following error message is observed during GRES if an IRB interface is configured without a profile: RPD_DYN_CFG_GET_PROF_NAME_FAILED. [PR1526481](#)
- Commit error messages come twice while validating the physical-cores command. [PR1527322](#)
- The transit PTP packet might be unexpectedly modified when passing through MPC2E-NG, MPC3E-NG, and MPC5E line cards. [PR1527612](#)
- The speed command cannot be configured under the interface hierarchy on an extended port when the MX204 or MX10003 router works as an aggregation device. [PR1529028](#)
- In the subscriber management environment, the RADIUS interim accounting records does not get populated with the subscriber statistics. [PR1529602](#)
- The SFP-LX or SFP-SX optics on MIC-3D-20GE-SFP-E/EH might show as unsupported after unified ISSU. [PR1529844](#)
- BiDi 1G SFP optics giving wrong value in JVision for "optics/laser_rx_power_*_thresholds". [PR1530120](#)
- If the ECMP is set to 128 and a route is learnt from 128 peers, the unilist nexthop might have incorrect ECMP path and traffic might be routed to undesirable paths which could cause traffic drop. [PR1530803](#)
- After performing unified ISSU with a high-scale bridge-domain configuration, less than 0.0254 percent of traffic loss is observed for a single bridge-domain interface. [PR1531051](#)
- On the MX10003 router, PEM 0 always shows as Absent or Empty even if PEM 0 is present. [PR1531190](#)
- Commit may fail after Routing Engine switchover. [PR1531415](#)
- New subscribers might fail to connect due to "Filter index space exhausted" error. [PR1531580](#)

- Deleting the address of the jmgmt0 interface might fail if the shortened version of the CLI command is used. [PR1532642](#)
- The interface with the "pic-mode 10GE" configuration may not come up if upgrading to 18.4R3-S4 or later versions. [PR1534281](#)
- Some routes might get incorrectly programmed in the forwarding table in the kernel which is no longer present in rpd. [PR1534455](#)
- On vMX platform, if vFPC is not getting the required CPU resources from host server, multiple vmxt cores might be generated and vFPC gets rebooted. It is most likely to occur in lite mode while less likely to occur while in performance mode. [PR1534641](#)
- The clear ike statistics command does not work with remote gateway. [PR1535321](#)
- Certain BGP SR-TE segment lists cause the rpd process to generate the core file during tunnel attribute parsing. [PR1535632](#)
- Snmp mib walk for jnxSubscriber OIDs returns General error. [PR1535754](#)
- All SFBs might go offline due to fabric failure and fabric self-ping probes performing the disable-pfe action. [PR1535787](#)
- Junos OS: MX Series: Dynamic filter fails to match IPv6 prefix (CVE-2021-0205). [PR1536100](#)
- Multicast traffic might be observed even through unexpected interfaces with distributed IGMP is enabled. [PR1536149](#)
- Enhancements are needed for debugging l2ald. [PR1536530](#)
- The chassisd memory leak might cause traffic loss. [PR1537194](#)
- The following error message might be observed when the JAM packages for the MX204, MX10003, and MX10008 are installed: JAM: Plugin installed for summit_xxx PIC. [PR1537389](#)
- Version-alias gets missed for subscribers configured with dynamic profiles after unified ISSU. [PR1537512](#)
- Not able to get the sessions after Configure IDS, Add IDS-RULE in the SS in the next-hop style. [PR1537609](#)
- Deactivating/activating PTP/syncE in the upstream router causes the 100G links on the LC2103 to flap. [PR1538122](#)
- AFT based TRIO FPCs (MPC10, 11) PFE cli command "show jnh exceptions inst <inst-number> may cause FPC to crash. [PR1538138](#)
- Traffic drop might be seen when executing "request system reboot". [PR1538252](#)
- Junos OS: Upon receipt of a specific BGP FlowSpec message network traffic may be disrupted. (CVE-2021-0211) [PR1539109](#)
- The accounting interim-updates for subscriber does not work after GRES and subsequent reboot of FPCs in the node-slicing setup. [PR1539474](#)
- The rpd memory leak might be observed on the backup Routing Engine due to link flaps. [PR1539601](#)

- The mspmand process leaks memory in relation to the MX telemetry reporting the following error message: RLIMIT_DATA exceed. [PR1540538](#)
- With hold time configuration, the ge Interfaces remain down on reboot. [PR1541382](#)
- Subscriber might not come up on some dynamic VLAN ranges in a subscriber management environment. [PR1541796](#)
- After changing addresses in the source pool, if the carrier-grade NAT traffic does not stop, the source pool cannot perform the NAT translation from the new pool. [PR1542202](#)
- The KRT queue might get stuck after RE switchover. [PR1542280](#)
- Port mirroring with maximum-packet-length configuration does not work over the GRE interface. [PR1542500](#)
- The license errors may get returned on backup Routing Engine when trying to commit the configuration. [PR1543037](#)
- The mspmand process might generate core file on activating or deactivating the interface. [PR1544794](#)
- Traffic loss might be observed when Switch Fabric Board 3/MPC8E 3D combination is used in MX2010/MX2020. [PR1544794](#)
- In the syslog output, the sylog-local-tag name is truncated (as SYSLOG_SF) when he sylog-local-tag name is configured as SYSLOG_SFW. [PR1547505](#)
- Continuous rpd errors might be seen and new routes will fail to be programmed by rpd. [PR1545463](#)
- The nsd daemon crashes after configuring the inline NAT44 in the USF mode. [PR1547647](#)
- The verbose command unexpectedly becomes hidden after Junos OS Release 16.1 for set system export-format json. [PR1547693](#)
- SENSOR APP DWORD leak is observed during the period of churn for routes bound to the sensor group. [PR1547698](#)
- SR-TE may stay UP when the routes deleted through policy. [PR1547933](#)
- Multicast traffic drop might be seen after ISSU. [PR1548196](#)
- The rpd crash might be seen when BGP service route is resolved over color-only SRTE policy. [PR1550736](#)
- In the EVPN-VXLAN scenario, as part of fixing 1535515 (All the ARP reply packets toward some address are flooded across the entire fabric), on bd with no irb, the mac+ip ageout is adjusted by +30 seconds. This change exposed an issue in arp expiry handling. This change resulted in high cpu utilization and is fixed through this PR 1551025. [PR1551025](#)
- The PPPoE subscribers might fail to login. [PR1551207](#)
- "LCM Peer Absent" might be seen on all TVP platforms. [PR1551760](#)
- The fabric errors are observed and the FPC processes might get offlined with SCBE3, MPC3E-NG, or MPC3E and MPC7 or MPC10 in the increased-bandwidth fabric mode. [PR1553641](#)

- Configuring HFRR i.e. link-protection on an interface may cause rpd to crash. [PR1555866](#)
- ISSU may be aborted on MX devices for version 20.2R2-S1. [PR1557413](#)
- On MX Platform with any of these linecards -MPC9E/JNP10K-LC2101/JNP10003-LC2103/MX204-MPC, Packets corruption might occur with enabling PTP(Protocol Time protocol) on 100G/40G interfaces mapped to Channelized MAC. [PR1557758](#)
- The l2cpd core files might be seen on reboot. [PR1561235](#)
- The rpd crash might be observed during processing huge amount of PIM prune messages. [PR1561984](#)
- MX platforms with MX-SCBE3 may reboot continuously. [PR1564539](#)
- The ALQ session between the two routers is expecting to have a controlled source and destination address (peer config in both end). To be able to control what this address is used as source on a router with multiple routed interfaces, a good technique is to use a directly connected interface for this communication. In the case where the routers are not directly connected a tunnel interface is equally good technique. But the ALQ need to be allowed to use this. This PR fix this. [PR1567735](#)
- On MX150, "request system software add" CLI is disabled in 19.4R3-S1, 20.1R2, and 20.4R1. [PR1568273](#)
- agent sensor - "__default_fabric_sensor__" seems to be partly applied to some FPCs, which caused zero payload issue - "AGENTD received empty payload for pfe sensor __default_fabric_sensor__". [PR1569167](#)

Infrastructure

- Output drops in 'show interfaces extensive' might display 0 temporarily during a race condition when SNMP query for JnxCos is also issued. [PR1533314](#)

Interfaces and Chassis

- The configuration might not be applied after deleting all existing logical interfaces and adding a new logical interface for an IFD in a single commit. [PR1534787](#)
- Inline Y.1731 SLM or DM does not work in enhanced-cfm-mode for the EVPN UP MEP scenario. [PR1537381](#)
- Backup router generates VRRP_NEW_BACKUP syslog during bringup. [PR1539277](#)
- The following error message might occur after commit for configuration under interface hierarchy: should have at least one member link on a different FPC. [PR1539719](#)
- The following the commit error is observed while trying to delete unit 1 logical systems interfaces: ae2.1: Only unit 0 is valid for this encapsulation. [PR1547853](#)
- The startup-silent-period command might not work in Junos OS Release 20.3R1 or later. [PR1548464](#)
- The VCP port is marked as administratively down on the wrong MX Series Virtual Chassis member. [PR1552588](#)
- The dcd process might leak memory on pushing the configuration to the ephemeral database. [PR1553148](#)

MPLS

- The rpd scheduler might slip after the link flaps. [PR1516657](#)
- The inter-domain LSP with loose next hop path might get stuck in down state. [PR1524736](#)
- The ping mpls rsvp command does not take into account lower MTU in the path. [PR1530382](#)
- The rpd process might crash when the LDP route with indirect next hop is deleted on the aggregated Ethernet interface. [PR1538124](#)
- Performing commit may trigger externally provisioned LSP MBB mechanism. [PR1546824](#)
- A new LSP might not be up even if bypass LSP is up and "setup-protection" is configured. [PR1555774](#)

Network Address Translation (NAT)

- Need to improve the maximum eNode connections for one persistent NAT binding from 8 to 32. [PR1532249](#)

Network Management and Monitoring

- Commit error while deleting the routing instance when snmp trap-group also have the same routing instance referred. [PR1555563](#)

Platform and Infrastructure

- PE-CE OAM CFM might have issues in AE interface case. [PR1501656](#)
- The output of the show jnh qmon queues-sensor stats 0 command has no content. [PR1514881](#)
- The VPLS connection might be stuck in the Primary Fail status when a dynamic profile is used on the VPLS pseudowire logical interface. [PR1516418](#)
- The state of the flow detection configuration might not be displayed properly if DDOS-SCFD is configured globally. [PR1519887](#)
- Flow programming issue for lt- interface in the Packet Forwarding Engine level is observed. [PR1525188](#)
- Junos OS: MX Series: Trio-based MPC memory leak when Integrated Routing and Bridging (IRB) interface is mapped to a VPLS instance or a Bridge-Domain (CVE-2021-0202). [PR1525226](#)
- The following error message is observed when alarms after interface reset: 7836 ifl 567 chan_index 8 NOENT & jnh_ifl_topo_handler_pfe(13015): ifl=567 err=1 updating channel table nexthop. [PR1525824](#)
- The VxLAN encapsulation over IPv6 underlay might not work on MX routers. [PR1532144](#)
- There is a TWAMP interoperability issue between Junos OS releases. [PR1533025](#)
- The fpc process might crash when the next hop memory of ASIC is exhausted in the EVPN-MPLS scenario. [PR1533857](#)
- The ISSU might fail on Junos platforms with LUCHIP based line cards. [PR1535745](#)
- Subscribers are not coming up VPLS on PS interface. [PR1536043](#)

- TWAMP interoperability issue can be seen if the Junos release has only the fixes for PR-1434740, PR-1533025 but not the fix for PR-1536939. [PR1536939](#)
- Packet loss might be observed when the RFC2544 egress reflector session is configured on the non-zero Packet Forwarding Ethernet interface. [PR1538417](#)
- AUTO-CORE-PR : JDI CI ROUTING : vmxt_lnx core found @ l2_metro_bd_host_inject_del bd_platform_delete bd_handle_msg. [PR1538516](#)
- The rmopd process memory leak might be seen if TWAMP client is configured. [PR1541808](#)
- Trio-based FPC might crash when the underlying Layer 2 interface for ARP over IRB interface is changed from the physical interface to the LSI interface. [PR1542211](#)
- ARP expired timer on backup Routing Engine is not same with master Routing Engine if aging-timer is configured. [PR1544398](#)
- On all MX platforms with BNG (Broadband Network Gateway) scenario, an internal timer (re-ARP timer) on backup RE could cause an ARP storm upon GRES switchover since there are lots of arp timeout on the new master RE in 2 minutes. The re-ARP timer is one-tenth of the ARP aging timer (default ARP aging timer is 20 minutes, so 1/10 of 20 minutes is 2 minutes). The fix will automatically adjust the timer based on the scale and the configured aging time avoiding ARP storm on new master. [PR1547583](#)
- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)
- Traffic is not forwarded over IRB to l2circuit on lt interfaces. [PR1554908](#)
- IPv4 EXP rewrite might not work properly when inet6-vpn enabled. [PR1559018](#)
- On MX platform, T4000 platform and EX9200 platform, end-users or end-hosts might not get an IPv4 address from Dynamic Host Configuration Protocol (DHCP) server when Distributed Denial-of-Service (DDOS) attack is happened on DHCP rebind packets or renew packets. In the end, end-users or end-hosts could not access into network after lease time of the IPv4 address expired. [PR1562474](#)

Routing Policy and Firewall Filters

- The policy configuration might be mismatched between the rpd and mgd process when deactivate policy-options prefix-list is involved in the configuration sequence. [PR1523891](#)
- Generate route goes to hidden state when protect core knob is enabled. [PR1562867](#)

Routing Protocols

- The output of the show isis interface detail command might be incorrect if wide-metrics-only is enabled for IS-IS and the ASCII representation of the metric in decimal is more than 6 characters long. [PR1482983](#)
- On NFX-series and MX150 devices the following error messages are seen in the messages log file for the interfaces that have SFP installed in them: fpc0 FAILED(-1) read of SFP eeprom for port: 13. [PR1529939](#)
- The rpd might crash with BGP RPKI enabled in a race condition. [PR1487486](#)

- Ppmd core file generated after MS-MPC restart. [PR1490918](#)
- The rpd might crash after deleting and re-adding a BGP neighbor. [PR1517498](#)
- Tag matching in the VRF policy does not work properly when the independent-domain option is configured. [PR1518056](#)
- The BGP session with VRRP virtual address might not come up after a flap. [PR1523075](#)
- The VRF label is not assigned at ASBR when the inter AS is implemented. [PR1523896](#)
- The IS-IS LSP database synchronization issue might be seen while using the flood-group feature. [PR1526447](#)
- The rpd process generates core file at is_srv6_delete_locator_end_sid_data isis_srv6_end_sid_local_data_delete isis_srv6_locator_config_check. [PR1531830](#)
- Transit labels for Layer 3 VPN routes are pushed momentarily to the MPLS.0 table. [PR1532414](#)
- Configuring then next hop and then reject on a route policy for the same route might cause the rpd process to crash. [PR1538491](#)
- After moving peer out of the protection group, the path protection does not get removed from the PE router. Multipath routes are still present. [PR1538956](#)
- For spring with TE-shortcuts, MPLS S=0 route label is missing in the logical_system r5_lsr for label 801007 upon activating mpls label switched path just after deactivating isis TE inet shortcuts. [PR1539671](#)
- The rpd process generates the core file at gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash. [PR1541768](#)
- Traffic loss might be seen in next-hop-based dynamic tunnels of L3VPN scenario after changing the dynamic-tunnel preference. [PR1542123](#)
- The metric of prefixes in intra-area-prefix LSA might be changed to 65535 when the metric of one of the OSPFv3 p2p interfaces is set to 65535. [PR1543147](#)
- The BGP session neighbor shutdown configuration does not effect the non-established peer. [PR1554569](#)
- The changes do not get effective when the values are set under static default hierarchy. [PR1555187](#)
- The BGP session might not come up if extended-nexthop is enabled by default on the other vendor remote peer. [PR1555288](#)
- Sending multicast traffic to downstream receiver on Trio based Virtual Chassis platforms might fail. [PR1555518](#)
- 6PE prefixes may not be removed from RIB upon reception of withdrawal from a BGP neighbor when RIB sharding is enabled. [PR1556271](#)
- Multipath info still shown for BGP route even after disabling interface for one path. [PR1557604](#)
- 6PE prefixes may not be removed from RIB upon reception of withdrawal from a BGP neighbor when RIB sharding is enabled. [PR1556271](#)

- VPN routes learned from core files were not advertised to CE when bgp sharding is configured. [PR1560661](#)
- All Layer3 VPN route ages reset when adding or deleting a VRF. [PR1560827](#)
- Wrong SPF calculation might be observed for OSPF with ldp-synchronization hold-time configured after interface flap. [PR1561414](#)
- If BGP route flap damping is enabled and some routes received from a BGP peering session are hidden due to damping, the routes which are stored in the route list after the damped routes might be stuck in routing table with "Accepted DeletePending" state and not be removed when the BGP peering session goes down. [PR1562090](#)

Services Applications

- L2TP subscribers might fail to establish a session on MX if the CPE is a virtual host. [PR1527343](#)
- The following error message is observed: SPD_CONN_OPEN_FAILURE: spd_pre_fetch_query: unable to open connection to si-1/0/0. [PR1550035](#)

Subscriber Access Management

- Subscriber accounting messages retransmissions exist even after configuring accounting retry 0. [PR1405855](#)

VPNs

- The Junos image upgrade/installation with 'validate' will fail with XML errors. [PR1525862](#)
- MVPN multicast route entry might not be properly updated with the actual downstream interfaces list. [PR1546739](#)

Resolved Issues: 20.3R1

Application Layer Gateways (ALGs)

- The srxpfe or mspmand process might crash if FTPS is enabled in a specific scenario. [PR1510678](#)

Class of Service (CoS)

- The following error message is observed: **GENCFG write failed (op, minor_type) = (delete, Scheduler map definition) for tbl id 2 ifl 0 TABLE Reason: No such file or directory.** [PR1476531](#)
- The MX Series routers with MPC1 Q and MPC2 Q line cards might report memory errors. [PR1500250](#)

EVPN

- When a dynamic list next hop is referenced by more than one route, it might result in an early deletion of the next hop from the kernel, thereby assigning the NH index as 0 (Next hop type: Dynamic List, next hop index: 0" in the output of the **show route** command). This would not result in a crash, but an early delete from kernel. As a workaround, restarting the routing solves the issue and the NH index gets reassigned properly. [PR1477140](#)
- The ARP resolution to the gateway IRB address fails if **decapsulate-accept-inner-vlan** or **encapsulate-inner-vlan** is configured. [PR1526618](#)
- The rpd process might crash when **auto-service-id** is configured in the EVPN-VPWS scenario. [PR1530991](#)
- The rpd process might generate a core file when the Routing Engine switches over after disabling the BGP protocol globally. [PR1490953](#)
- VXLAN bridge domain might lose the VTEP logical interface after restarting chassisd. [PR1495098](#)
- 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](#)
- Configuring the **proxy-macip-advertisement** command for EVPN-MPLS leads to functionality breakage. [PR1506343](#)
- With the EVPN-VXLAN configurations, the IRB MAC does not get removed from the route table after disabling IRB. [PR1510954](#)
- ARP might break when multicast snooping is enabled in EVPN for the VLAN-based and VLAN-bundle service scenarios. [PR1515927](#)
- Unable to create a new VTEP interface. [PR1520078](#)
- Packets might not be sent out of the IRB interface if there is no Layer 2 interface in the associated bridge-domains. [PR1498534](#)
- IRB interface might get stuck in the **Down** state in an EVPN multihome scenario. [PR1479681](#)

Forwarding and Sampling

- UTC timestamp is used in the **flat-file-accounting** files when a profile is configured. [PR1509467](#)
- DHCP subscribers might get stuck in the **Terminated** state for around 5 minutes after disabling the cascade ports. [PR1505409](#)
- Traffic might get dropped due to not exceeding the configured bandwidth under policer. [PR1511041](#)
- The DHCP relay might not work normally under EVPN with VXLAN environment. [PR1487385](#)
- The pfd process might crash while running the **show pfe fpc x** command. [PR1509114](#)

General Routing

- The **show security group-vpn member IPsec security-associations detail | display xml** command is not in the expected format. [PR1349963](#)
- Constant memory leak might lead to FPC memory exhaustion. [PR1381527](#)
- The chassisd might crash due to hardware-database errors. [PR1383246](#)
- On the MX2000, the following error message might be observed if the MPC7 line card is offline when Routing Engine switchover occurs: **Failed to get xfchip**. [PR1388076](#)
- After an MX Series router with the JNP10K-LC2101 line card is powered on, a voltage of 1345-1348 mV is read for about 20 seconds, which gets stabilized to 1493 mV. During this period, the FPC **x Voltage Tolerance Exceeded** major alarm is raised. [PR1415671](#)
- The following Error messages are observed on the MPC card in the manual mode:
clksync_as_evaluate_synce_ref: 362 - Failed to configure clk. [PR1490138](#)
- FPC might crash after GRES when committing changes in the firewall filter with the **next term** statements in a subscriber scenario. [PR1421541](#)
- The RPD scheduler slips might be seen upon executing the **show route resolution extensive 0.0.0.0 | no-more** command if the number of routes in the system is large (several millions). [PR1425515](#)
- Layer 2 over GRE is not supported in Junos OS Release 19.3R1. Even though, the configuration gets committed, the feature does not work. [PR1435855](#)
- The MPC9E line card does not get offline due to unreachable destinations in the phase 3 stage. [PR1443803](#)
- FEC statistics are not reset after changing the FEC mode. [PR1449088](#)
- When an M-VLAN interface (OIF map) is changed, the existing multicast subscribers with membership reports in place experience loss of multicast traffic till traffic is forwarded to the new OIF map. For example, a new M-VLAN interface. [PR1452644](#)
- Interfaces shut down by the **disable-pfe** action might not come up when you use the MIC offline or online command. [PR1453433](#)
- The FPC or the Packet Forwarding Engine might crash with the ATM MIC installed in the FPC. [PR1453893](#)
- Application and removal of 1-Gbps speed results in the channel being down. [PR1456105](#)
- In the MVPN instance, the traffic drops on multicast receivers within the range of 0.1 to 0.9 percent. [PR1460471](#)
- The bbe-smgd process generates core files on the backup Routing Engine. [PR1466118](#)
- With the BGP rib-sharding and update-threading, traffic drops 100 percent in the BGP Layer 3 VPN streams, post the removal or restoration configuration. [PR1469873](#)
- The following syslog message are observed: **fpcX user.notice logrotate: ALERT exited abnormally with [1]**. [PR1471006](#)

- When you reboot the external server, the SNMP values configured within the `/etc/snmp/snmpd.conf` file at the server get overwritten with the content from the JDM SNMP configuration section. The trap configuration changes get completely removed. Restarting or stopping and starting JDM does not change the host `/etc/snmp/snmpd.conf` file. Only system reboot of the server occurs. [PR1474349](#)
- The `kmd` process might crash in a specific simultaneous rekey scenario. [PR1474797](#)
- The following error log messages are observed: **chassisd[7836]: %DAEMON-3-CHASSISD_IOCTL_FAILURE: acb_get_fpga_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device) after every commit.** [PR1477941](#)
- The `cpd` process might generate core file after upgrading to Junos OS Release 19.4 and later. [PR1527602](#)
- The `ukern-platformd` process might crash on the MX2000 router with the MPC11 line card. [PR1478243](#)
- Interface traffic statistics in the **show interface** command might display incorrect values for a LAG with the MPC10 or MPC11 line card child links. [PR1478540](#)
- All PPPoE subscribers might not log in after FPC restarts. [PR1479099](#)
- Fabric healing logic incorrectly makes all MPC line cards go offline in the MX2000 router while the hardware fault is located on one specific MPC line card slot. [PR1482124](#)
- The downstream IPv4 packet greater than BR MTU gets dropped in MAP-E. [PR1483984](#)
- The traffic rate might not be as expected on the aggregated Ethernet interface after applying a shared-bandwidth policer. [PR1484193](#)
- The peer interface does not go down after the MPC11E line card reboot. [PR1485682](#)
- The input errors on the MX150 router might be zero in the output of the **show interfaces extensive** command when there are CRC or align errors on the interface. [PR1485706](#)
- The `aftd` process might crash. [PR1487416](#)
- XML is not properly formatted. [PR1488036](#)
- Daemon might restart due to mishandling of data. [PR1489512](#)
- With the MX-SPC3 service card, NAT might not be processed on an order as setup. [PR1489581](#)
- Prolonged flow control might occur with MS-MPC or MS-MIC. [PR1489942](#)
- The ISSU is not supported on the NG-MPC line cards from Junos OS Release 19.4R1. [PR1491337](#)
- Multiple deactivation or activation of the security traceoptions along with a single NAPT44 session might crash the flowd process. [PR1491540](#)
- MS-MIC goes down after loading some Junos OS releases in an MX-VC scenario. [PR1491628](#)
- User-configured MTU might be ignored after the ISSU upgrade using the **request vmhost software in-service-upgrade** command. [PR1491970](#)

- There is a delay in the LT interfaces on the MPC11E line card coming up after configuring the scaled PS interfaces anchoring to RLT. [PR1492330](#)
- On the MX10008 router, the SNMP table **entPhysicalTable** does not match the PICs shown in the output of the **show chassis hardware** command. [PR1492996](#)
- The MPC10 or MPC11 line card might crash if the interface is configured with the firewall filter referencing a shared-bandwidth policer. [PR1493084](#)
- In an MX Series, setting or deleting a Virtual Chassis C port causes other Virtual Chassis ports on the same FPC or MIC slot to bring the link in the **Down** state for a few seconds, possibly interrupting the communication with the other member chassis. [PR1493699](#)
- **Used-Service-Unit** of the CCR-U has **Output-Bytes** counter zero. [PR1516728](#)
- The LSP might not come up in the LSP externally provisioned scenario. [PR1494210](#)
- The following error message is seen for the AF interfaces on an FPC when the peer FPC is restarted: **PFE_ERROR_FAIL_OPERATION: Unable to unbind cos scheduler from physical interface.** [PR1494452](#)
- In a node slicing setup, after GRES, the RADIUS interim updates might not carry actual statistics. [PR1494637](#)
- Group address is not programmed back post deactivation and activation of the bridge domain. [PR1495480](#)
- VPLS flood NH might not get programmed correctly. [PR1495925](#)
- B4 might not be able to establish the softwire with AFTR. [PR1496211](#)
- The following error messages are generated by Packet Forwarding Engine when the subscribers come up over a pseudowire interface: **PFEIFD: Could not decode media address with length 0.** [PR1496265](#)
- The MPC10E line card might restart with sensord crash due to a timing issue. [PR1497343](#)
- Outbound SSH connection flaps or memory leaks during the push configuration to ephemeral database with high rate. [PR1497575](#)
- Port numbers logged in the ALG syslog are incorrect. [PR1497713](#)
- Subscribers might be disconnected after one of the aggregated Ethernet participating FPCs comes online in a Junos node slicing scenario. [PR1498024](#)
- SNMP polling does not show correct **PSM jnxOperatingState** when one of the PSM inputs fails. [PR1498538](#)
- The rpd process might crash when multiple VRFs with **IFLs link-protection** are deleted at a single time. [PR1498992](#)
- The commit check might fail when adding a logical interface into a routing-instance, which has no-normalization command enabled under the routing-instances stanza. [PR1499265](#)
- Heap memory leak might be seen on the MPC10 and MPC11 line cards. [PR1499631](#)
- After disabling and enabling the ams0 interfaces, the NAT sessions do not get synchronized back to the current standby SDG. [PR1500147](#)

- The SPC3 card might crash if the SIP ALG is enabled. [PR1500355](#)
- Unexpected behavior during | **display inheritance** is observed when the foreground is deactivated. [PR1500569](#)
- The **show services alg conversations** and **show services alg sip-globals** commands are not supported in USF mode. [PR1501051](#)
- The MX2020 and MX2010 routers continuously log **pem_tiny_power_remaining:** in the chassisd log. [PR1501108](#)
- Application ID does not get displayed under the **nat/sfw** rule configured with application any rule. [PR1501109](#)
- The chassisd process might become nonresponsive. [PR1502118](#)
- On the MPC11 line card, the **show syslog** command in the Packet Forwarding Engine shell might time out. [PR1502877](#)
- The packets from a nonexistent source on the GRE or UDP designated tunnel might be accepted. [PR1503421](#)
- Configuring the **ranges** statement for autosensed VLANs might not work on the vMX platforms. [PR1503538](#)
- MIBS added as part of **jnxLicenseInstallTable: jnxLicenseStartDate jnxLicenseEndDate**. [PR1503790](#)
- The **show bridge statistics** command output does not display the statistics information for the pseudowire subscriber interfaces. [PR1504409](#)
- The gNMI stream does not follow the frequency on the subscription from the collector. [PR1504733](#)
- Fan speed might toggle between full and normal on the MX960 router with an enhanced FRU. [PR1504867](#)
- The rpd process might crash in case of a network churn when the telemetry streaming is in progress. [PR1505425](#)
- The PSM firmware upgrade must not allow multiple PSM upgrades in parallel to avoid the firmware corruption and support multiple firmwares for different hardware. [PR1524338](#)
- Addition and removal of an aggregated Ethernet interface member link might cause the PPPoE subscriber session and traffic to drop. [PR1525585](#)
- After sending the Layer 4 or Layer 7 traffic, the HTTP redirect messages are not captured as expected. [PR1505438](#)
- The l2cpd process might crash if the ERP configuration is added or removed, and the l2cpd process is restarted. [PR1505710](#)
- VRRPv6 might not work in an EVPN scenario. [PR1505976](#)
- Mapping leaks when the private and public IP addresses are from the same prefix. [PR1507477](#)
- **GnmiJuniperTelemetryHeader** incompatibility is introduced in Junos OS Release 19.3. [PR1507999](#)

- Outbound SSH connection flap or memory leak issues might be observed during push configuration to the ephemeral database with a high rate. [PR1508324](#)
- JET API RouteMonitorRegister might result in an unresponsive gRPC session. [PR1509655](#)
- The host-generated packets might be dropped if the **force-control-packets-on-transit-path** statement is configured. [PR1509790](#)
- The disabled QSFP transceiver might fail to get turned on. [PR1510994](#)
- PFCP message acknowledgment or non- acknowledgment responses are not tracked without the fix. If the CPF peer drops an acknowledged UPF response message and CPF retries the request, the reattempts do not get an acknowledgment by the response cache at UPF and get silently dropped. This causes the CPF state machine to constantly retry requests with those message being dropped at UPF, which leads to the **Established** state at both CPF and UPF. [PR1511708](#)
- Static subscribers are logged out after creating a unit under the demux0 interface. [PR1511745](#)
- The multicast traffic might be dropped if ALB is enabled on the aggregated Ethernet interface. [PR1512157](#)
- Memory leak on l2ald might be seen when adding or deleting the routing-instances or bridge-domains configuration. [PR1512802](#)
- The wavelength configured through the CLI might not be set on the **SFP+-10G-T-DWDM-ZR** optics when the optics is used on the MPC7E line card. [PR1513321](#)
- Modifying the segment list of the segment routing LSP might not work. [PR1513583](#)
- Subscribers might not be able to bind again after performing back-to-back GRES followed by an FPC restart. [PR1514154](#)
- Active sensor check fails while checking the **show agent sensors |display xml** command. [PR1516290](#)
- The MPC7E line card with QSFP installed might get rebooted when the **show mtip-chmac <1|2> registers** vty command is executed. [PR1517202](#)
- There might be memory leak in cfmd if both the CFM and inet/IPv4 interfaces are configured. [PR1518744](#)
- The vgd process might generate a core file when the OVSDB server restarts. [PR1518807](#)
- The PADI packets might be dropped when the interface encapsulation VPLS is set along with accepted protocol configured as PPPoE. [PR1523902](#)
- According to the OC data model, the **openconfig-alarms.yang** subscription path must be used as **system/alarms/alarm**. [PR1525180](#)
- WAG control route prefix length are observed. [PR1526666](#)
- Non-impacting error message is seen in the message logs: **IFP error>/..../..../..../src/pfe/usp/control/applications/interface/ifp.c@3270:(errno=1000) tunnel session add failed.** [PR1529224](#)
- On the MX960 router, the following error message might be observed: **SCHED L4NP[0] Parity errors.** [PR1464297](#)

- The vmcore process crashes sometimes along with the mspmand process on MS-MPC/MS-MIC if large-scale traffic flows are processed. [PR1482400](#)
- The heap memory utilization might increase after extensive subscriber login or logout. [PR1508291](#)
- On the MPC10 and MPC11 line cards, the heap memory leaks with the MoFRR feature. [PR1479024](#)
- Some of the virtual services might not up after GRES or rpd restart. [PR1499655](#)
- On the MX150 series of routers, the **request system halt** and **request system power-off** commands do not work as expected. [PR1468921](#)
- With MPC10 and MPC11 line cards, switchovers are slow to backup the upstream interface. [PR1497127](#)
- The MACsec session might fail to establish if 256 bit cipher suite is configured for MACsec connectivity association assigned to a logical interface. [PR1514680](#)
- The MPC10E line card might crash with the sensord process generating a core file due to a timing issue. [PR1526568](#)
- The **commit confirm** command might not rollback the previous configuration when the commit operation fails. [PR1527848](#)
- Certain BGP SRTE segment lists cause the rpd process to generate core file during tunnel attribute parsing. [PR1535632](#)
- Any change in the nested groups might not be detected on commit and does not take effect. [PR1484801](#)
- In the MX10003 routers, RCB always detect fire temperature and shutdown in a short time after downgrade. [PR1492121](#)
- Inline JFlow might report wrong value for some fields in the flow records after enabling the next hop-learning and route churn occurs. [PR1500179](#)
- The MACsec delay protection fails to drop or discard the delayed MACsec packets. [PR1503010](#)
- The transit PTP packet might be unexpectedly modified when passing through MPC2E-NG, MPC3E-NG, and MPC5E line cards. [PR1527612](#)
- Not able to get the sessions after configuring IDS, adding IDS-RULE in the SS in the next-hop style. [PR1537609](#)
- The MPC11E line card might get stuck in the **Present** state during booting in a rare condition. [PR1482105](#)
- The SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at collector. [PR1484322](#)
- The mgd process might become nonresponsive, crash the dcd process, or crash the dcd process commit check process. [PR1491363](#)
- The fpc process might crash in an inline mode with CFM configured. [PR1500048](#)
- On the MX150 router, the logical interfaces stay up during the vmhost halt or power-off senario. [PR1526855](#)

Infrastructure

- If the serial number of the PEM starts with 1F1, the following alarm might be generated: **Minor FPC PEM Temp Sensor Failed**. [PR1398128](#)
- SNMP polling might return an unexpectedly high value for the ifHCOutOctets counter for a physical interface when any jnxDom OID is processed at the same time. [PR1508442](#)
- Unknown MIB OID 1.3.6.1.2.1.47.2.0.30 are referenced in the SNMP trap after upgrading to Junos OS Release 18.4R3.3. [PR1508281](#)
- Packet counter does not work as expected when SNMP is used. [PR1422929](#)
- Kernel stack data disclosure is observed. [PR1485747](#)

Interfaces and Chassis

- Traffic might get dropped as the next hop points to ICL even though the local MC-LAG is up. [PR1486919](#)
- The **sonet-options** configuration statement is disabled for the xe interface that works in wan-phy mode. [PR1472439](#)
- The vrrpd might crash when dual VLAN on VRRP interfaces is configured. [PR1512658](#)
- Fail to configure proactive ARP detection. [PR1476199](#)
- A stale IP address might be seen after a specific order of configuration changes under the logical-systems scenario. [PR1477084](#)
- Control logical interface 32767 is not created on the VLAN-tagged IFD even after removing the VLAN 0 configuration. [PR1483395](#)
- On the MPC6 line cards, the CFM DM two way verification fails with invalid timestamp. [PR1489196](#)
- Some of the logical interfaces might not come up with the configured vlan-bridge encapsulation. [PR1501414](#)
- Unexpected dual VRRP backup state might occur after performing two subsequent Routing Engine switchovers with **track priority-hold-time** configured. [PR1506747](#)
- Commit failure is observed while deleting all the units under the ps0 interface. [PR1514319](#)
- The following error message is observed: **Request failed: OID not increasing: ieee8021CfmStackServiceSelectorType**. [PR1517046](#)
- Buffer overflow vulnerability in device control process is observed. [PR1519334](#)

Intrusion Detection and Prevention (IDP)

- When creating the custom IDP signatures that match raw bytes (hexadecimal), the commit check fails if the administrator configures the depth parameter. [PR1506706](#)

J-Web

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

Juniper Extension Toolkit (JET)

- JET application configuration must be disabled before upgrading Junos OS vmhost images. [PR1488769](#)

Junos Fusion Provider Edge

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

Layer 2 Ethernet Services

- For the MX204 router, the vendor ID is set as **MX10001** in the factory-default configuration and in the DHCP client messages. [PR1488771](#)
- The DHCP subscribers might not come up when DHCP ALQ and VRRP are configured. [PR1490907](#)
- Issues with the DHCPv6 relay processing confirm and reply packets are observed. [PR1496220](#)
- The MC-LAG might be down after disabling and then enabling the force-up configuration. [PR1500758](#)
- The aggregated Ethernet interface sometimes might not come up after switch is rebooted. [PR1505523](#)
- The DHCPv6 lease query is not as expected while verifying the DHCPv6 server statistics. [PR1506418](#)
- The **show dhcp relay** statistics display **DHCPLEASEUNASSIGNED** instead of **DHCPLEASEUNASSINGED**, which is spelling error. [PR1512239](#)
- The **show dhcpv6 relay** statistics must display **DHCPV6_LEASEQUERY_REPLY** instead of **DHCPV6_LEASEQUERY_REPL** for the messages sent. [PR1512246](#)
- The DHCP6 lease query is not as expected while verifying the DHCPV6v relay statistics. [PR1521227](#)
- The memory leak in jdhcpcd might be seen if access-profile is configured under the **dhcp-relay** or **dhcp-local-server** statement. [PR1525052](#)
- Receipt of malformed DHCPv6 packets causes jdhcpcd to crash. [PR1511782](#)
- The jdhcpcd process crashes when processing a specific DHCPDv6 packet in the DHCPv6 relay configuration. [PR1512765](#)

MPLS

- The RSVP interface bandwidth calculation rounds up. [PR1458527](#)
- The rpd process might crash in PCEP for the RSVP-TE scenario. [PR1467278](#)
- The rpd process might crash when the BGP flaps with FEC 129 VPWS enabled. [PR1490952](#)

- If there are two directly connected BGP peers established over MPLS LSP and the MTU of the IP layer is smaller than the MTU of the MPLS layer. Also, if the BGP packets from the host have the DF bit set, the BGP session might keep flapping because of the usage of the wrong TCP-MSS. [PR1493431](#)
- The rpd process might crash in a rare condition in the SR-TE scenario. [PR1493721](#)
- The rpd process saves the core file while performing ISSU from Junos OS Release 19.3R2 or later. [PR1493969](#)
- The same device responds twice for traceroute in case it goes through the MPLS network under specific conditions. [PR1494665](#)
- The rpd process might crash when the SNMP polling is done using the OID `jnxMplsTeP2mpTunnelDestTable`. [PR1497641](#)
- Traffic loss might occur if ISSU is performed when P2MP is configured for an LSP. [PR1500615](#)
- The CSPF job might get stalled for a new or an existing LSP in a high-scale LSP setup. [PR1502993](#)
- The rpd process might crash with RSVP configured in a rare timing case. [PR1505834](#)
- Activating or deactivating the LDP-sync under OSPF might cause the LDP neighborship to go down and stay down. [PR1509578](#)
- The rpd process might crash after upgrading Junos OS Release 18.1 to a later release. [PR1517018](#)
- The SNMP trap is sent with the incorrect OID `jnxSpSvcSetZoneEntered`. [PR1517667](#)
- The LDP session-group might throw a commit error and flap. [PR1521698](#)
- The rpd process generates core file on the backup Routing Engine. [PR1495746](#)
- The rpd process might crash when rpd restarts or GRES switchovers. [PR1506062](#)
- The auto-bandwidth feature might not work correctly in the MPLS scenario. [PR1504916](#)
- The inter-domain LSP with loose next-hops path might get stuck in the **Down** state. [PR1524736](#)

Network Management and Monitoring

- The SNMPv3 informs might not work properly after rebooting. [PR1497841](#)

Platform and Infrastructure

- Configured scheduler-map is not applied on ms- interface if the service PIC is in the **Offline** state during commit. [PR1523881](#)
- `core.vmx.mpc0` seen at `5 0x096327d5` in the `l2alm_sync_entry_in_pfes` (`context=0xd92e7b28`, `sync_info=0xd92e7a78`) at `../../../../src/pfe/common/applications/l2alm/l2alm_common_hw_api.c:1727`. [PR1430440](#)
- The output of the `show jnh qmon queues-sensor stats 0` command has no content. [PR1514881](#)

- On the MX204 router, GRE with sampling causes the following Packet Forwarding Engine error: **MQSS(0): MALLOC: Underflow error during reference count read - Overflow 1, Underflow 1, HMCIF 0, Address 0x8d62e0**. [PR1463718](#)
- On MX150 and vMX, the VXLAN packet might get discarded because the flow caching does not support VXLAN when flow caching is enabled. [PR1466470](#)
- CFM session malfunctions when it is configured along with the inner and outer native VLAN ID configuration. [PR1484303](#)
- In the MX104 chassis, the **show system buffer** command displays all zeros. [PR1484689](#)
- Traceroute monitor with MTR version v.69 shows a false 10 percent loss. [PR1493824](#)
- Packets get dropped when next hop is IRB over an It interface. [PR1494594](#)
- The Routing Engine might crash when a large number of next hops are quickly deleted and added again in a large ARP or ND scaled scenario. [PR1496429](#)
- The rmopd.core process generates core files when committing a configuration replacement of the ms-interface used. [PR1499230](#)
- Traffic to VRRP virtual IP or MAC addresses might be dropped when ingress queuing is enabled. [PR1501014](#)
- Python or SLAX script might not be executed. [PR1501746](#)
- Traffic originated from another subnet is sent out with 0x8100 instead of 0x88a8. [PR1502867](#)
- Traffic loss might be seen in certain conditions under an MC-LAG setup. [PR1505465](#)
- The kernel might crash causing the router or the Routing Engine to reboot when making virtual IP related change. [PR1511833](#)
- During route table object fetch failure, the FPC might crash. [PR1513509](#)
- With multiple different fixed-sized traffic streams configured at 10,000,00 fps (40-Gbps combined rate) on aggregated Ethernet0 along with another independent aggregated Ethernet interface (aggregated Ethernet1, 50 percent line rate 4 streams bidirectional => 118-Gbps combined traffic rate), both hosted on a single Packet Forwarding Engine instruction of the MPC11E line card, small varying packet drops occur for every iteration on aggregated Ethernet1 on disabling aggregated Ethernet0. [PR1464549](#)
- There is a TWAMP interoperability issue between Junos OS releases. [PR1533025](#)
- Arbitrary code execution vulnerability in the Telnet server. [PR1502386](#)

Routing Protocols

- The BGP session might become nonresponsive with high BGP OutQ value after GRES on both sides. [PR1323306](#)
- Cannot configure **set system services ssh protocol-version v1**. [PR1440476](#)

- When configuring an alternate incoming interface for a PIM RPF check using `rpf-selection`, the additional groups outside the configured range might switch to the alternate incoming interface. [PR1443056](#)
- Multicast traffic loss might be seen in certain conditions while enabling the IGMP snooping under EVPN-VXLAN ERB scenario. [PR1481987](#)
- RIPv2 might malfunction when changing the interface type from P2MP to broadcast. [PR1483181](#)
- There might be rpd process memory leak in a certain looped MSDP scenario. [PR1485206](#)
- Layer 3 VPN RR with the **family route-target** and **no-client-reflect** statements does not work as expected. [PR1485977](#)
- Traffic loss might be observed while performing GRES in an MPLS setup. [PR1486657](#)
- The BGP route-target family might prevent the RR from reflecting the Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- The rpd process generates core files at `rt_nh_resolve_add_gen` in `../../../../src/junos/usr.sbin/rpd/lib/rt/rt_resolve_ind.c` with the evpn-dhcp configurations. [PR1494005](#)
- In all platforms with IPv6 scenario, the last route entry in the inet6.0 or inet6.3 RIB might not get deleted if there is another configuration present under the RIB configuration. (For example, set routing-options rib inet6.0 static defaults active). This might cause a service to still be available that the customer no longer wants to use. [PR1495477](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes the rpd process to crash. [PR1497721](#)
- The IS-IS hello authentication does not generate the correct digest value for **hmac_sha1** algorithm. [PR1498452](#)
- The rpd process might crash if the import policy is changed to accept more routes that exceed the teardown function threshold. [PR1499977](#)
- The rpd process might crash in a multicast scenario with BGP configured. [PR1501722](#)
- The rpd process might crash while processing a specific BGP packet. [PR1502327](#)
- The mcsnoopd process generates core files during the execution of an internal script. [PR1503211](#)
- BGP might not advertise routes to peers after a peer flap. [PR1507195](#)
- The rpd process might crash due to RIP updates being sent on an interface in down state. [PR1508814](#)
- The IS-IS SR routes might not be updated to reflect the change in the SRMS advertisements. [PR1514867](#)
- The BGP link-bw of the non-multipath routes are included in an aggregation. [PR1515264](#)
- The rpd process might crash if there is a huge number of SA messages in an MSDP scenario. [PR1517910](#)
- NLRI handling improvements for BGP-LS ID TLV is needed. [PR1521258](#)

- The output of the **show isis interface detail** command might be incorrect if **wide-metrics-only** is enabled for IS-IS and the ASCII representation of the metric in decimal is more than 6 characters long. [PR1482983](#)
- The BGP RPKI ROA withdrawal might lead to an unexpected BGP route flap. [PR1483097](#)
- The rpd process might crash after deleting and then adding a BGP neighbor. [PR1517498](#)
- Core file is generated in **krt_mcnh_update_rpf_info()** when TI-LFA is used with MOFRR. [PR1493259](#)
- The route entries might be unstable after being imported into the inet6.x RIB through rib-group. [PR1498377](#)

Services Applications

- The FPC process might crash with an npc core file if the service interface is configured under a service set in USF mode. [PR1502527](#)
- The output of the **show services l2tp tunnel extensive** command does not show the configured session limit. [PR1503436](#)
- Destination lockout functionality does not work at the tunnel session level when CDN code is received. [PR1532750](#)

Subscriber Access Management

- The following syslog messages are observed: **pfe_tcp_listener_open_timeout: Peer info msg not received from addr: 0x6000080. Socket 0xfffff804ad23c2e0 closed** [PR1474687](#)
- LTS incorrectly sends the access-request with the Tunnel-Assignment-ID, which is not compliant with RFC 2868. [PR1502274](#)
- CCR-T does not contain the usage-monitoring information. [PR1517507](#)
- The **show network-access aaa subscribers statistics username "<>"** command fails to fetch the **subscriber-specific AAA** statistics information if a subscriber username contains a space. [PR1518016](#)

User Interface and Configuration

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

VPNs

- The l2circuit neighbor might become nonresponsive in the **Ready** state at one end of the MG-LAG peer. [PR1498040](#)
- The rpd process might crash in certain conditions after deleting the l2circuit configuration. [PR1502003](#)
- The MPLS label manager might allow configuration of a duplicated VPLS static label. [PR1503282](#)
- The output value of the **show mvpn c-multicast inet source-pe | display xml** command is not proper. [PR1509948](#)

- The rpd process might crash after removing the last configured interface under the l2circuit neighbor. [PR1511783](#)
- The rpd process might crash when deleting the l2circuit configuration in a specific sequence. [PR1512834](#)

SEE ALSO

What's New	 108
What's Changed	 131
Known Limitations	 137
Open Issues	 143
Documentation Updates	 188
Migration, Upgrade, and Downgrade Instructions	 188

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for MX Series routers.

SEE ALSO

What's New	 108
What's Changed	 131
Known Limitations	 137
Open Issues	 143
Resolved Issues	 156
Migration, Upgrade, and Downgrade Instructions	 188

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.3R3](#) | 189
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS](#) | 190

- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | 192](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 194](#)
- [Upgrading a Router with Redundant Routing Engines | 195](#)
- [Downgrading from Release 20.3R3 | 195](#)

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

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

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

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

Basic Procedure for Upgrading to Release 20.3R3

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-x86-32-20.3R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.3R3.9-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-x86-32-20.3R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.3R3.9-limited.tgz
```

Replace source with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 20.3R3, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following MX Series routers:
 - MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008

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

NOTE: After you install a Junos OS Release 20.3R3 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-20.3R3.9-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-20.3R3.9-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 20.3R3 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

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

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

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 20.3R3

To downgrade from Release 20.3R3 to another supported release, follow the procedure for upgrading, but replace the 20.3R3 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 | 108](#)

[What's Changed | 131](#)

[Known Limitations | 137](#)

[Open Issues | 143](#)

[Resolved Issues | 156](#)

[Documentation Updates | 188](#)

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 196](#)
- [What's Changed | 198](#)
- [Known Limitations | 199](#)
- [Open Issues | 200](#)
- [Resolved Issues | 201](#)
- [Documentation Updates | 203](#)
- [Migration, Upgrade, and Downgrade Instructions | 203](#)

These release notes accompany Junos OS Release 20.3R3 for the NFX Series Network Services Platforms. 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 20.3R3 | 197](#)
- [What's New in Release 20.3R2 | 197](#)
- [What's New in Release 20.3R1 | 197](#)

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

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

What's New in Release 20.3R3

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Application Security

- **Listing of micro-applications and non-configurable applications (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.3R1, we've introduced the following operational commands to display applications details:

- **show services application-identification application micro-applications** to display the list of micro-applications.
- **show services application-identification application non-configurable** to display the list of non-configurable applications.

[See [show services application-identification application micro-applications](#) and [show services application-identification application non-configurable](#).]

- **Application signature package rollback (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.3R1, you can roll back the current version of the application signature package to the previous version by using one of the following methods:
 - **Automatic**—The system automatically rolls back to the previous version of the application signature package when the signature package installation fails on your security device.
 - **Manual**—You can roll back the application signature package to its previous version on your security device using the **request services application-identification rollback** command.

[See [Predefined Application Signatures for Application Identification](#).]

Wireless WAN

- **LTE support in dual CPE deployments (NFX250 NextGen)**—Starting in Junos OS Release 20.3R1, you can provide a backup WAN connection by configuring LTE modules on a pair of NFX250 NextGen devices operating in cluster mode.

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

SEE ALSO

What's Changed 198
Known Limitations 199
Open Issues 200
Resolved Issues 201
Documentation Updates 203
Migration, Upgrade, and Downgrade Instructions 203

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 198](#)
- [What's Changed in Release 20.3R2 | 199](#)
- [What's Changed in Release 20.3R1 | 199](#)

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

What's Changed in Release 20.3R3

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 20.3R3 for NFX Series devices.

What's Changed in Release 20.3R2

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 20.3R2 for NFX Series devices.

What's Changed in Release 20.3R1

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 20.3R1 for NFX Series devices.

SEE ALSO

[What's New | 196](#)

[Known Limitations | 199](#)

[Open Issues | 200](#)

[Resolved Issues | 201](#)

[Documentation Updates | 203](#)

[Migration, Upgrade, and Downgrade Instructions | 203](#)

Known Limitations

There are no known limitations for NFX Series devices in Junos OS Release 20.3R3.

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 | 196](#)

[What's Changed | 198](#)

[Open Issues | 200](#)

[Resolved Issues | 201](#)

[Documentation Updates | 203](#)

[Migration, Upgrade, and Downgrade Instructions | 203](#)

Open Issues

IN THIS SECTION

- Platform and Infrastructure | 200

Learn about open issues in this release for NFX Series devices.

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

Platform and Infrastructure

- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock (No such file or directory)** messages are seen during FTP. [PR1315605](#)
- On vmhost platforms, if the XML output of the **request vmhost mode test | display xml rpc** command is used in netconf, it fails. As a workaround, use the following XML command:

```
<rpc>
  <request-vmhost-mode>
    <mode-name>flex</mode-name>
  </request-vmhost-mode>
</rpc>
```

[PR1559786](#)

SEE ALSO

- [What's New | 196](#)
- [What's Changed | 198](#)
- [Known Limitations | 199](#)
- [Resolved Issues | 201](#)
- [Documentation Updates | 203](#)
- [Migration, Upgrade, and Downgrade Instructions | 203](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 201](#)
- [Resolved Issues: 20.3R2 | 201](#)
- [Resolved Issues: 20.3R1 | 202](#)

Learn which issues were resolved in the Junos OS Release 20.3R3 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: 20.3R3

General Routing

- The l2cpd core files might be seen on reboot. [PR1561235](#)

Platform and Infrastructure

- On NFX150 devices, when J-Flow v5 is configured and the J-Flow v5 server is reachable through an IPsec tunnel, and the MTU size of this IPsec tunnel is configured as 1500, the J-Flow packets are not generated on NFX Series devices. As a workaround, use J-Flow v9 or IPFIX version, instead of J-Flow v5, to enable the J-Flow functionality on NFX Series devices. [PR1539964](#)

Resolved Issues: 20.3R2

Interfaces

- When you configure analyzers on VNF interfaces with output port as other VNF interfaces, all the incoming and outgoing packets can be mirrored on to the designated analyzer port. However, it is noticed that after a system reboot, this functionality stops working and no packets are mirrored on the output analyzer port. [PR1480290](#)

Platform and Infrastructure

- False positive TSensor errors are reported on vjunos0. [PR1508580](#)
- When you upgrade the NFX150 devices from Junos OS Release 19.4 to Junos OS Release 20.2, the upgrade fails and an error message is displayed: `"/usr/sbin/boot_mgmt_fsm: line 40: echo: write error: No space left on device issue"`. [PR1532334](#)

Resolved Issues: 20.3R1

Application Security

- AppQoS is sending active probe packets for the deleted **active-probe-params**. [PR1492208](#)

High Availability

- On an NFX250 chassis cluster, L3 interfaces are not getting created after secondary automatic reboot when control port recovery is enabled. [PR1502449](#)

Interfaces

- On NFX350 devices, the **show interfaces | no-more** command output stops appearing for around 20 seconds after displaying the d10 interface. [PR1502626](#)
- On NFX350 devices, the **clear interface statistics all** command takes a longer time to execute. [PR1475804](#)

Platform and Infrastructure

- After initiation of zeroization, the NFX250 device is going into a reboot loop. [PR1491479](#)
- The **request vmhost power-off** command reboots the NFX250 NextGen device instead of powering off the device. [PR1493062](#)
- On NFX150 devices, MAC aging does not work. You must remove aged MAC entries from the CLI. [PR1502700](#)
- After you upgrade the JDM image from Junos OS Release D497.1 to Junos OS Release 18.4R3-S2, tunnels are down in the gateway router (GWR). [PR1507165](#)
- On NFX150 devices, ZTP over LTE configuration commit fails for **operation=create** in XML operations configuration. [PR1511306](#)
- The device reads the board ID from eeprom directly using I2C upon power cycle. [PR1529667](#)

SEE ALSO

[What's New | 196](#)

[What's Changed | 198](#)

[Known Limitations | 199](#)

[Open Issues | 200](#)

[Documentation Updates | 203](#)

[Migration, Upgrade, and Downgrade Instructions | 203](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for NFX Series devices.

SEE ALSO

[What's New | 196](#)

[What's Changed | 198](#)

[Known Limitations | 199](#)

[Open Issues | 200](#)

[Resolved Issues | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 203](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 203](#)
- [Basic Procedure for Upgrading to Release 20.3 | 204](#)

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

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

Basic Procedure for Upgrading to Release 20.3

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 20.3R3:

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 | 196](#)

[What's Changed | 198](#)

[Known Limitations | 199](#)

[Open Issues | 200](#)

[Resolved Issues | 201](#)

[Documentation Updates | 203](#)

Junos OS Release Notes for PTX Series

IN THIS SECTION

- What's New | 206
- What's Changed | 218
- Known Limitations | 224
- Open Issues | 227
- Resolved Issues | 229
- Documentation Updates | 234
- Migration, Upgrade, and Downgrade Instructions | 235

These release notes accompany Junos OS Release 20.3R3 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 Release 20.3R3 | 207
- What's New in Release 20.3R2 | 207
- What's New in Release 20.3R1 | 207

Learn about new features introduced in Junos OS Release 20.3R3 for the PTX Series.

What's New in Release 20.3R3

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Hardware

- **Support for QSFP-100G-DR transceivers (PTX1000, PTX10002-60C, PTX10008, and PTX10016)**—Starting in Junos OS Release 20.3R1, we provide support for the QSFP-100G-DR transceivers. These transceivers interoperate with 400-Gbps breakout optics. For example, the QDD-400G-DR4 interconnects with up to four QSFP-100G-DR transceivers. The QSFP-100G-DR transceivers interconnect in single links (QSFP-100G-DR to QSFP-100G-DR or to QSFP-100G-FR) and interoperate at the shortest link length.

NOTE: These transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

- **Support for QSFP-100G-FR transceivers (PTX1000, PTX10002-60C, PTX10008, and PTX10016)**—Starting in Junos OS Release 20.3R1, we provide support for the QSFP-100G-FR transceivers. These transceivers interoperate with the QDD-4X100G breakout optics. For example, the QDD-4X100G-FR interconnects with up to four QSFP-100G-FR transceivers. The QSFP-100G-FR transceivers interconnect in single links (QSFP-100G-FR to QSFP-100G-FR or to QSFP-100G-DR) and interoperate at the shortest link length.

NOTE: These transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

Authentication, Authorization, and Accounting

- **Support for TCP authentication option (TCP-AO) for BGP and LDP connections (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can use TCP-AO to authenticate TCP segments exchanged during BGP and LDP sessions. It supports both IPv4 and IPv6 traffic. TCP-AO provides a framework to support multiple stronger algorithms, such as HMAC-SHA1 and AES-128, to create its message digest. TCP-AO supports up to 64 keys that can be used for a BGP or an LDP session. You can configure a new key for a BGP or LDP session during its lifetime without causing any session flap. Each key becomes active based on its configured start time.

In earlier releases, you could use only the TCP MD5 authentication method. It supports only MD5 algorithm to create its message digest.

[See [TCP Authentication Option \(TCP-AO\) for BGP and LDP Sessions](#) and [authentication-key-chains \(TCP-AO\)](#).]

IP Tunneling

- **IPIP encapsulation for flexible tunnel interfaces (FTIs) (MX Series, PTX Series, and QFX10002)**—We've extended flexible tunnel interfaces (FTIs) and existing forwarding constructs to support configuring static IPv4 IP-in-IP tunnels and RIB APIs. To configure an IP-in-IP tunnel on a FTI, use the **ipip** option at the **[edit interfaces interface-name unit logical-unit-number tunnel encapsulation]** hierarchy level.

[See [Configuring Flexible Tunnel Interfaces](#) and [ipip](#).]

- **Support for IP over IP next hop based tunneling (MX Series, PTX1000, PTX10000, QFX10000, and QFX10002)**—Starting in Junos OS Release 20.3R1, we support an IP-over-IP encapsulation to facilitate IP overlay construction over IP transport network. An IP network contains edge devices and core devices. To achieve higher scale and reliability among these devices, you need to logically isolate the core network from the external network that the edge devices interact with, by using an overlay encapsulation. Among the other overlay encapsulations supported, IP over IP encapsulation is the only kind where transit devices are able to parse the inner payload and use inner packet fields for hash computation and customer edge devices are able to route traffic into and out of the tunnel without any throughput reduction. IP over IP relies on a next hop-based infrastructure to support higher scale.

On MX Series routers, routing protocol daemon(RPD) sends the encapsulation header with tunnel composite nexthop and the Packet Forwarding Engine finds the tunnel destination address and forwards the packet. On PTX Series routers and QFX10000 switches, RPD sends fully resolved next hop-based tunnel to PFE. You can either use static configuration or a BGP protocol configuration to distribute routes and signal dynamic tunnels. You can also configure Interface based firewall filters on any transit or egress device with an action to decapsulate IP-IP packets and forward it to main instance or to a routing-instance as required.

[See [Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation](#).]

- **Support for filter-based decapsulation of IPv4 and IPv6 unicast traffic encapsulated in IPv4 IP-in-IP tunnels (MX Series, PTX1000, PTX10002, and QFX10002)**—Junos OS supports decapsulating IPv4 and IPv6 unicast traffic that has been encapsulated in IPv4 IP-in-IP tunnels using firewall filters. If the outer IPv4 header address matches the firewall configuration and the packet has **ipip** set as the protocol type,

then the outer IPv4 header is removed and the packet is routed based on the inner IPv4 or IPv6 address. If the packet does not have the expected **ipip** header, the packet is dropped.

Configure this feature using the following CLI statements at the **[edit firewall family inet filter filter-name term term-name]** hierarchy:

- **from protocol ipip**: Set the protocol type as IP-IP.
- **then decapsulate ipip**: Decapsulate the IP-IP packet. The inner IP destination address is routed using the inet.0 routing table by default.
- **then decapsulate ipip routing-instance routing-instance-name**: Decapsulate the IP-IP packet and route the inner destination address using the specified routing instance.

Use **show firewall** to view the configuration.

[See [filter \(Firewall Filters\)](#) and [Configuring IP Tunnel Interfaces](#).]

Juniper Extension Toolkit (JET)

- **Support for static backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (MX Series, PTX Series, and QFX10002)**—We've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to set up underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation. JET APIs notify the controller of active paths, interfaces, and changes to the interface state. The loop-free backup paths help quickly restore failed core transport networks built with only IP protocols.

[See [JET APIs on Juniper EngNet](#).]

- **Support for policy match condition to match programmed routes (MX Series, PTX Series, and QFX10002)**—We've introduced a new option programmed that allows policy matches for routes injected by JET APIs. To allow policy matches for routes injected by JET APIs, use the **programmed** option at the **[edit policy-options policy-statement policy-name term term-name from]** hierarchy level. To view details about programmed routes, use the **show route programmed (detail | extensive)** command.

[See [policy-statement](#) and [show route](#).]

- **RIB service API option to control route distribution (MX Series, PTX Series, and QFX10002)**—We've added a no-advertise flag to the RIB service API per-route RouteAttributes object to limit re-advertisement of the provisioned route. You can set this flag to TRUE to prevent the route from being redistributed to routing protocols and advertised to peers.

[See [JET APIs on Juniper EngNet](#).]

- **Juniper Extension Toolkit (JET) supports BFD Service APIs for routing protocol process (rpd) programmability (MX Series, PTX Series, QFX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can use programmable rpd (prpd) BFD APIs to add, update, and delete BFD sessions and subscribe to BFD events from outside applications. These APIs enable the integration of rpd with software-defined networking (SDN) controllers and increase the flexibility of your network. The prpd BFD APIs support BFD Echo-Lite sessions in single-hop IPv4 and IPv6 modes.

The following BFD Service APIs are supported:

- Initialize
- SessionAdd
- SessionUpdate
- SessionDelete
- SessionDeleteAll
- Subscribe
- Unsubscribe

Use the **show bfd session extensive** command to view BFD sessions. BFD sessions added through prpd BFD APIs are labeled with **PRPD:<session-id>** in the client field. The **<session-id>** is 1 for the first BFD session that is added, 2 for the second, and so on.

[See [show bfd session extensive](#) and [JET APIs on Juniper EngNet.](#)]

Junos OS XML, API, and Scripting

- **Support for REST API over nondefault virtual routing and forwarding (VRF) instance (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can execute Junos OS operational commands using the REST API over a nondefault VRF instance. The nondefault VRF instance can be a user-defined instance or the management instance, `mgmt_junos`.

The REST API allows you to execute Junos OS operational commands over HTTP(S). If you don't specify a routing instance, REST API requests are sent over the default routing instance. Use a nondefault VRF instance to improve security and make it easier to troubleshoot.

Use the **routing-instance *routing-instance*** statement at the **[edit system services rest]** hierarchy level to specify a nondefault VRF instance for REST API requests.

[See [Management Interface in a Nondefault Instance](#) and [rest](#).]

Junos Telemetry Interface

- **Support for forwarding information base (FIB) sensor on JTI (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can use the Junos telemetry interface (JTI) and remote procedure calls (gRPC) services to stream or export ON_CHANGE FIB, also known as forwarding table, statistics to outside collectors. This feature supports the OpenConfig YANG model OC-AFT.

To enable and manage FIB streaming, include the following statements on the client device:

- **set system fib-streaming** and **delete system fib-streaming** statements at the **[edit]** hierarchy level to launch or terminate the process.
- **set system fib-streaming traceoptions file *file-name*** statement at the **[edit]** hierarchy level to configure a logging file.

- **set system fib-streaming traceoptions flag *flag-name*** statement at the **[edit]** hierarchy level to configure various trace parameters.
- **set system fib-streaming traceoptions level *level-name*** statement at the **[edit]** hierarchy level to configure log levels.

Use the **restart fib-streaming** command to restart the process.

To show information about FIB streaming, use the following operational mode commands on the client device:

- **show fib-streaming**
- **show fib-streaming next-hop-groups**
- **show fib-streaming next-hops**
- **show fib-streaming routes ipv4-unicast**
- **show fib-streaming routes ipv6-unicast**
- **show fib-streaming routes mpls**

The following table shows supported sensors:

Table 6: Supported Sensors

Supported Sensors
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/dscp[]
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/next-hop-group
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/id

Table 6: Supported Sensors (*continued*)

Supported Sensors
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/interface
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/input-interfaces/input-interface/state/subinterface
/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/prefix
/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/prefix
/network-instances/network-instance/afts/ipv4-unicast/ipv4-entry/state/next-hop-group
/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/prefix
/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/prefix
/network-instances/network-instance/afts/ipv6-unicast/ipv6-entry/state/next-hop-group
/network-instances/network-instance/afts/mpls/label-entry/label
/network-instances/network-instance/afts/mpls/label-entry/state/label
/network-instances/network-instance/afts/mpls/label-entry/state/next-hop-group
/network-instances/network-instance/afts/mpls/label-entry/state/popped-mpls-label-stack
This leaf reports the same label value in case of pop or swap.
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/id
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/next-hops/nexthop/index
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/next-hops/nexthop/state/weight
/network-instances/network-instance/afts/nexthops/nexthop/index

Table 6: Supported Sensors (continued)

Supported Sensors
<code>/network-instances/network-instance/afts/next-hops/next-hop/juniper/state/lsp-id</code>
This leaf is a new augmentation.
<code>/network-instances/network-instance/afts/next-hops/next-hop/state/ip-address</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/state/mac-address</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/state/pushed-mpls-label-stack</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/interface-ref/state/interface</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/interface-ref/state/subinterface</code>
<code>/network-instances/network-instance/afts/next-hops/next-hop/juniper/state/mapped-next-hop-index</code>
This leaf is a new augmentation.

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

- **Support for aggregated Ethernet interface ON_CHANGE with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX3000, PTX5000, PTX10001-36MR, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.3R1, Junos telemetry interface (JTI) supports ON-CHANGE statistics for aggregated Ethernet interfaces for minimum links and member interfaces.

To export these statistics to an outside collector using remote procedure call (gRPC) services and JTI, include the following resource paths in a subscription:

- `/interfaces/interface/aggregation/state/min-links/`
- `/interfaces/interface/aggregation/state/member/`

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

Layer 3 Features

- **Support for BGP Layer 3 VPN over IP-IP Tunnel (MX Series, PTX1000, QFX10002, and QFX10008)**—Starting in Junos OS Release 20.3R1, we support BGP Layer 3 VPN over IP over IP (IP-IP) tunnels to create a new transport service. IP-IP tunnels terminate into service-layer VRF, so you do not need to use a service label. This feature allows interoperability between the new VRF and traditional VRF, so both types of overlays can coexist in your network. You can use this feature to transition from an MPLS network to an IP fabric core network and to protect your network from distributed denial-of-service (DDoS) attacks.

To use VPN over an IP-IP tunnel, configure the **tunnel-attribute** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* then]** or **[edit policy-options policy-statement *policy-name* then]** hierarchy level.

To configure the receiver to program the dynamic tunnel using the tunnel attribute, use the **extended-nexthop-tunnel** statement at the **[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family (inet-vpn | inet6-vpn) unicast]** hierarchy level.

[See [BGP Layer 3 VPN over IP-IP Tunnels Overview](#), [family \(Protocols BGP\)](#), [policy-statement](#), [vrf-export](#), and [Configuring IP Tunnel Interfaces](#).]

MPLS

- **New output fields added in the show path-computation-client lsp extensive command (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.3R1, you'll see association details such as **Association type**, **ID**, and **source** in the output of the **show path-computation-client lsp** command when you use the command with the **extensive** option.

[See [show path-computation-client lsp](#).]

Network Management and Monitoring

- **Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can use the **probe** command to query the status of the probed interface. The proxy interface resides on the same node as the probed interface, or it can reside on a node to which the probed interface is directly connected.

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

To enable the **probe** command, configure the **extended-echo** statement under the **[edit system]** hierarchy.

[See [Using the Probe command](#).]

- **Enhanced sFlow (PTX5000)**—Starting in Junos OS Release 20.3R1, you can use sFlow to detect and sample MPLS and GRE traffic flows on PTX5000 routers. sFlow technology is a monitoring technology for high-speed switched or routed networks.

[See [Overview of sFlow Technology](#).]

- **Enhancements to sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.3R1, devices running Junos OS with upgraded FreeBSD support the following enhancements to sessions over outbound HTTPS:
 - Connecting to multiple outbound HTTPS clients by configuring one or more clients at the **[edit system services outbound-https]** hierarchy level
 - Configuring multiple backup gRPC servers for a given outbound HTTPS client
 - Establishing a csh session
 - Establishing multiple, concurrent NETCONF and csh sessions between the device running Junos OS and an outbound HTTPS client
 - Configuring a shared secret that the outbound HTTPS client uses to authenticate the device running Junos OS
 - Authenticating the client using certificate chains in addition to self-signed certificates

[See [NETCONF and Shell Sessions over Outbound HTTPS](#).]

Port Security

- **MACsec preshared key hitless rollover (PTX10008 and PTX10016)**—Starting in Junos OS Release 20.3R1, we support preshared key (PSK) hitless rollover for Media Access Control Security (MACsec) on the PTX10K-LC1104 and PTX10K-LC1105 line cards. PSK hitless rollover uses a keychain of multiple security keys to prevent session drops when the connectivity association key (CAK) configuration changes.

[See [Configuring Media Access Control Security \(MACsec\) on Routers](#).]
- **Timer-based MACsec SAK refresh (MX10003, PTX10001, PTX10003, PTX10008, and PTX10016)**—Starting in Junos OS Release 20.3R1, you can configure a time-based refresh of the secure association key (SAK) on a Media Access Control Security (MACsec)-secured link. The key server generates the SAK and refreshes it periodically. The key server also sets a refresh interval, by default, based on packet counter movement. If the refresh does not occur frequently, this can leave the SAK vulnerable to attack. You can enhance security of the SAK by configuring a shorter time-based refresh interval.

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

Routing Protocols

- **Support for Implicit filter for default EBGp route propagation behavior without policies (ACX Series, JRR200, MX204, vRR and PTX5000)**—Starting in Junos OS Release 20.3R1, we've introduced a new configuration hierarchy, **defaults ebgp no-policy** at the existing **[edit protocols bgp]** hierarchy level. The configuration option separates the default policy for **receive** and **advertise**, into separate clauses (**accept**, **reject**, or **reject-always**) to allow the route propagation behavior of EBGp speakers to vary independently from its default behavior.

In earlier releases, the default behavior of BGP was to receive and advertise all routes. With the introduction of this feature, the default behavior still remains to “accept” all routes for both **receive** and **advertise**, but you also have an option to reject routes by default.

With the **reject** configuration, you can reject routes of type inet unicast and inet6 unicast in instance types master, vrf, virtual-router, and non-forwarding. With the **reject-always** configuration, you can reject all routes from being received or getting advertised, irrespective of address family or instance type. By using this feature, you can control traffic in leaf autonomous systems (AS) and thereby, prevent them from having to accidentally function as transit autonomous systems.

NOTE: The introduction of this implicit filter does not affect the existing deployments that rely on the default behavior.

[See [Implicit Filter for Default EBGp route propagation behavior without policies](#) and [defaults](#).]

- **TI-LFA SRLG protection and fate-sharing protection for OSPFv2 (MX Series and PTX Series)**—Starting in Junos OS Release 20.3R1, you can configure Shared Risk Link Group (SRLG) protection and fate-sharing protection for segment routing to choose a fast reroute path that does not include SRLG links and fate-sharing groups in the topology-independent loop-free alternate (TI-LFA) backup paths to avoid fate-sharing and SRLG failures. This is in addition to existing fast reroute options such as **link-protection** and **node protection** for segment routing.

To enable TI-LFA SRLG protection and fate-sharing protection with segment routing for OSPFv2, include the **srlg-protection** statement and the **fate-sharing-protection** statement respectively at the **[edit protocols ospf area area-id interface name post-convergence-lfa]** hierarchy level.

[See [Topology-Independent Loop-Free Alternate with Segment Routing for OSPF](#).]

- **BGP sharding for IPv4 and Ipv6 L3VPN, BGP-LU (MX Series, PTX-Series and vRR)**—Starting in Release 20.3R1, Junos OS supports BGP sharding and update IO features for these IPv4 and Ipv6 address families:
 - inet-vpn unicast
 - inet-vpn multicast (vrf.inet.2)
 - inet6-vpn unicast
 - inet6-vpn multicast (vrf.inet.2)

- `inet` labeled-unicast
- `inet6` labeled-unicast

To enable BGP sharding, configure **rib-sharding** at the `[edit system processes routing bgp]` hierarchy level. Sharding is dependent on the update I/O thread feature. To enable update I/O, configure **update-threading** at the `[edit system processes routing bgp]` hierarchy level.

BGP Sharding is supported only on 64-bit routing protocol process (rpd) where the Routing Engine has at least 4 CPU cores and 16 GB of memory. To enable your device to always use 64-bit mode, use **set force-64-bit** at `[edit system processes routing]` hierarchy level. If you configure rib-sharding on a routing engine, RPD creates sharding threads. By default, the number of sharding threads created is the same as the number of CPU cores on the routing engine. Optionally, you can specify the number-of-shards you want to create. To set the number of sharding threads, use **set number-of-shards <number-of-shards>** at `[edit system processes routing bgp rib-sharding]` hierarchy level. To set the number of update threads, use **set number-of-threads <number-of-threads>** at the `[edit system processes routing bgp update-threading]` hierarchy level. To enable your device to always use 64-bit mode, use **set force-64-bit** at `[edit system processes routing]` hierarchy level.

[See [rib-sharding](#) and [update-threading](#).]

- **ECMP nexthop update rate throttling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.3R1, 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.

[See [pause-computation-during-churn](#).]

Segment Routing

- **Support for LDP Tunneling over Segment Routing Traffic Engineering (MX Series, PTX Series, and ACX5448)**—Starting in Junos OS Release 20.3R1, you can tunnel LDP LSPs over Segment Routing Traffic Engineering (SR-TE) in your network. Tunneling LDP over SR-TE provides consistency and co-existence of both LDP LSPs and SR-TE LSPs.

[See [Tunneling LDP over SR-TE](#).]

Services Applications

- **Support for hardware timestamping of Two-Way Active Measurement Protocol (TWAMP) and real-time performance monitoring (RPM) probe messages (MX10008, MX10016, PTX10008, and PTX10016)**—Starting in Junos OS Release 20.3R1, we've extended support for hardware timestamping of TWAMP and RPM probe messages. Hardware timestamping is enabled by default for TWAMP, but you must configure it for RPM. You use TWAMP and RPM to measure IP performance between two

devices in a network. By configuring hardware timestamping for RPM, you can account for the latency in the communication of probe messages and generate more accurate timers in the Packet Forwarding Engine. To configure hardware timestamping for RPM, include the **hardware-timestamping** statement at the **[edit services rpm probe probe-owner test test-name]** hierarchy level.

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

- **IPFIX IPv4 and IPv6 template support for forwarding class and PLP (PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016)**—Starting in Junos OS Release 20.3R1, two more information elements have been added to the IPFIX IPv4 and IPv6 templates. These elements carry the packet loss priority (PLP) values and the first two characters of the configured forwarding class name that the sampled packet carries. The collector uses these elements to derive the DiffServ code point (DSCP) bits that the packet would contain while exiting the router. To use these elements, you must configure the **next-hop-learning enable** statement at the **[edit services flow-monitoring version-ipfix template name]** hierarchy level.

[See [nexthop-learning](#).]

SEE ALSO

What's Changed 218
Known Limitations 224
Open Issues 227
Resolved Issues 229
Documentation Updates 234
Migration, Upgrade, and Downgrade Instructions 235

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 219](#)
- [What's Changed in Release 20.3R2 | 220](#)
- [What's Changed in Release 20.3R1 | 222](#)

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.3R3 for the PTX Series.

What's Changed in Release 20.3R3

General Routing

- **Change in OID ifHighSpeed**—Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

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

Interfaces and Chassis

- **Warning message when taking an FPC offline**—PTX10003-80C and PTX10003-160C devices do not support the **request chassis fpc slot slot-number online** command. The only way to bring up an FPC (MPC) that is offline is by rebooting the chassis. So, when you take an FPC offline by using the **request chassis fpc slot slot-number online** command, the screen displays the following message. 'Warning : FPC <slot> cannot be made online using a CLI command. You need to perform router reboot using "request system reboot" to online the FPC <slot>. Do you wish to continue ? [yes,no] (no).

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

- **Blocking duplicate IP detection in the same routing instance (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. Refer to the table mentioned in the topic [inet \(interfaces\)](#). When you try to configure same IP on two logical interfaces inside same routing instance, the commit will be blocked with the error displayed as shown below:

```
[edit] user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24 [edit] user@host# commit commit complete [edit] user@host# set interfaces ge-0/0/2 unit 0 family inet address 2.2.2.2/24 [edit] user@host# commit [edit interfaces ge-0/0/2 unit 0 family inet] 'address 2.2.2.2/24' identical local address found on rt_inst [default], intfs [ge-0/0/2.0 and ge-0/0/1.0], family [inet]. error: configuration check-out failed
```

[See [inet\(interfaces\)](#).]

Junos OS XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or **set refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

Layer 2 Ethernet Services

- **Link selection support for DHCP**—We have introduced the **link-selection** statement at the **[edit forwarding-options dhcp-relay relay-option-82]** hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Prior to this release, the DHCP relay dropped packets during the renewal DHCP process and the DHCP server used the leaf's address as a destination to acknowledge the DHCP renewal message.

[See [relay-option-82](#).]

Network Management and Monitoring

- **Change in OID ifHighSpeed**—Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

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

What's Changed in Release 20.3R2

General Routing

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**—Starting in this release, we've renamed the **arp-snoop** packet type option in the **[edit system ddos-protection protocols] arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

- Loading of the default configurations in a RIFT package causes the following changes:

1. Output of the **show rift node status** command displays the node ID in hexadecimal number even though the node ID is configured in decimal, hexadecimal, or octal number.
 2. Some of the DDoS default configurations change because of the DDoS protection interferes with the RIFT BFD operation.
- **New commit check for MC-LAG (MX Series, PTX Series, QFX 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](#).]

Junos OS XML API and Scripting

- **The jcs:invoke() function supports suppressing root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are logged in system log files.
[See [invoke\(\) Function \(SLAX and XSLT\)](#).]
- **The jcs:invoke() function supports suppressing root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are logged in system log files.
[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

MPLS

- The **show mpls lsp extensive** and **show mpls lsp detail** commands display next-hop gateway LSPid—When you use the **show mpls lsp extensive** and **show mpls lsp detail** commands, you'll see next-hop gateway LSPid in the output.

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the **verbose** statement at the **[edit system export-format json]** hierarchy level. The default format to export configuration data in JSON changed from **verbose** format to **ietf** format starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **[edit system export-format json]** hierarchy level. Although the **verbose** statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See [export-format](#).]

What's Changed in Release 20.3R1

Class of Service (CoS)

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

General Routing

- **Trigger alarms when a PTX10008 or PTX10016 router has a mix of AC and DC power supplies**—If you insert a mix of AC and DC power supply units (PSU) into a PTX10008 or PTX10016 router, 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 the router has the same type of power supplies.

[See [Understanding Chassis Alarms](#).]

- **The show chassis power command displays the power supply state (PTX10008 and PTX10004)**—The **show chassis power** command displays the information regarding the state of the power supply (for instance, Online or Empty). This enhancement makes the **show chassis power** command output in Junos OS Evolved software consistent with that in Junos OS software.

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

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**—Starting in this release, we've renamed the **arp-snoop** packet type option in the **[edit system ddos-protection**

protocols] arp protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#)]

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.
- **Warning changed for configuration statements that correspond to deviate not-supported nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.
- **Python 3 add-on modules (PTX Series)**—Junos OS Evolved includes additional Python 3 libraries and modules, which Python scripts can import and use.

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

High Availability (HA) and Resiliency

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.

Junos OS XML, API, and Scripting

- **Changes to Junos XML RPC request tag names (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've updated the Junos XML request tag name for some operational RPCs to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include:
 - Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
 - Uppercase characters are converted to lowercase.

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

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 \(Services\)](#).

MPLS

- **Change in auto bandwidth adjustment (PTX5000)**—If auto bandwidth adjustment fails because of bandwidth unavailable error, the router tries to bring up the LSP with the same bandwidth during the subsequent reoptimization. In earlier releases, when the auto bandwidth adjustment fails, the current bandwidth is reset to the bandwidth that was already active.

See [rsvp-error-hold-time](#).

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

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

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.

SEE ALSO

What's New 206
Known Limitations 224
Open Issues 227
Resolved Issues 229
Documentation Updates 234
Migration, Upgrade, and Downgrade Instructions 235

Known Limitations

IN THIS SECTION

- [General Routing | 225](#)
- [MPLS | 226](#)
- [Routing Protocols | 226](#)

Learn about known limitations in Junos OS Release 20.3R3 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

- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error.** [PR1268678](#)
- During reconfiguration and link events at the physical interface level, the **PECHIP[4]:pe.ipw.misc_int.status:iq_disabled(0): (Count:3561)err_pkt(0)** error message can be seen. This does not impact traffic. [PR1476553](#)
- On the PTX1000 routers, the following error message is observed when the sampling MPLS+IPv4/IPv6 traffic is forwarded over the IP-IP tunnel: **dlu.ucode.jflow_not_routable pechip**. If an entropy label is present in the packet, then the packet has to be recirculated in the ASIC to do IPv4 or IPv6 lookup after stripping the outer entropy labels. If only an explicit NULL label is present, the ASIC has the capability to do the stripping of the NULL label and do IPv4 or IPv6 lookup without doing recirculation. In this case, because the packet has entropy labels, the packet gets recirculated and in the second pass processing, the inet sampling takes effect. J-Flow sampling is not designed to work after MPLS recirculation. This results in offset errors and a corrupted packet is being fed to the J-Flow pipeline in the ASIC. Enable MPLS-IPvx J-Flow on these interfaces and disable the inet filter on these logical interfaces. In this way, the packets will be sampled in first pass itself and not in the second pass. In this case, the exported flow record is MPLS-IPv4 or MPLS-IPv6 and not IPv4 or IPv6. The flow records include explicit NULL and entropy labels in addition to IP. [PR1485770](#)
- Filter based GRE tunneling is supported only in enhance-mode on PTX3000 routers. [PR1497819](#)
- On PTX Series platform with **set routing-options resolution preserve-nextthop-hierarchy** statement configured, reaching tunnel destination out-going route via BGP-over-BGP route recursive resolution is not supported. [PR1498085](#)
- In a tunnel termination scenario, packets with NULL, EL, and ELI labels followed by IPv4 or IPv6 header are treated as MPLS labeled packets and MPLS flows are created for these packets. Because these packets are treated as MPLS flows, the explicit NULL/EL labels are used for lookup, which results in failing the OIF getting reported as 0 for J-Flow records. [PR1502423](#)

- sFlow for IPoIP traffic is not supported in this release. [PR1508919](#)
- when counter sample is enabled, it attempts to fetch the physical interface statistics for sFlow enabled interfaces using rtsock messages to kernel. This blocks call and wait for the reply of earlier request and sends a new request only after receiving the reply of first one. So, FPC is occupied when this request is made and could not reply on time and hence the scheduler slip occurs. [PR1517076](#)

MPLS

- Increasing ECMP from 64 to 128 might cause the ingress LSP setup rate to be lower because of increased number of next hop changes for the IGP routes using shortcut. [PR1421976](#)
- LDP session might drop during the FRR if the **maxecmp** is configured to 128 and LDP/IGP has more than 64 RSVP LSP next hops and **ldp-tunneling** is configured on those next hops. [PR1430361](#)

Routing Protocols

- Commit check fails when rib-sharding is configured with these statements:
 - **routing-instances <name> routing-options multipath**
 - **routing-instances <name> routing-options policy-multipath**
 - **routing-instances <name> protocols mvpn.**
- Because of a race between route re-converge and the BGP-PIC version up message to the Packet Forwarding Engine, after a remote transit router reboot, certain BGP routes might reuse stale LDP next hops and cause packet discard at the transit router during the route re-convergence window. [PR1495435](#)

SEE ALSO

[What's New | 206](#)

[What's Changed | 218](#)

[Open Issues | 227](#)

[Resolved Issues | 229](#)

[Documentation Updates | 234](#)

[Migration, Upgrade, and Downgrade Instructions | 235](#)

Open Issues

IN THIS SECTION

- [General Routing | 227](#)
- [Layer 2 Ethernet Services | 229](#)
- [MPLS | 229](#)
- [Routing Protocols | 229](#)

Learn about open issues in the Junos OS Release 20.3R3 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 platforms, some non-fatal interrupts (for example, CM cache or AQD interrupts) are logged as fatal interrupts. The following log messages will be shown on CM parity interrupt: **fpc0 TQCHIP 0: CM parity Fatal interrupt,Interrupt status:0x10 fpc0 CMSNG: Fatal ASIC error, chip TQ fpc0 TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180010 msecs TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180005 msecs.** [PR1089955](#)
- On the PTX Series platforms with FPC-PTX-P1-A or FPC2-PTX-P1A line card might encounter a single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message is reported: **Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pq_t_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002.** The Junos OS chassis management error handling detects such a condition, raises an alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper support representative if the issue persists even after the FPC restart. [PR1254415](#)
- When CFP2-DCO-T-WDM-1 is plugged into a PTX Series PIC, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)
- On 30-port MACsec-enabled line card (LC1101-M-30C, LC1101-M-30Q, and LC1101-M-96X) of the PTX10008 chassis, when the **exclude-protocol lacp** statement is configured at the **[edit security macsec connectivity-association connectivity-association-name]** hierarchy level is deleted or deactivated, the

LACP protocol's Mux State shown under the output of **show lacp interface** CLI command, might remain as attached or detached and might not change to distributing state. [PR1331412](#)

- Alarm action does not work for minor errors after the threshold is changed to 1. [PR1345154](#)
- The PTX Series platform drops the wireless access point (WAP) heartbeat packets; as a result, the WAP cannot work. [PR1352805](#)
- Due to transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. This change applies to all platforms using Hybrid Memory Controller (HMC). [PR1384435](#)
- On the PTX3000 routers, the firewall counter for lo0 does not increment. [PR1420560](#)
- When the firewall filter has **Port-Mirror** as an action along with discard action, the mirrored packet will have two Layer 2 headers. The first Layer 2 header will be the original Layer 2 header and the second Layer 2 header will be egress interface Layer 2 header. This causes packet corruption and packets get discarded. [PR1437546](#)
- Memory leaks are expected in this release. [PR1438358](#)
- On PTX1000 and PTX10001 routers, the port mirror does not work when the **port-mirroring** is configured with the firewall filter. [PR1491789](#)
- The **show chassis fpc** command reports high CPU utilization in steady state. [PR1492731](#)
- IPv4 traffic verification of interface statistics at logical interfaces level aggregated Ethernet bundle is failed. [PR1531358](#)
- Flapping might be observed on channelized ports of PTX Series routers during ZTP, when one of the port is disabled on the supporting device. [PR1534614](#)
- When we run continuous sync (show interfaces aex extensive) and async (SNMP polling) queries on aggregated Ethernet interface in parallel, we might observe spikes in aggregated Ethernet interface framing error counter in between correct values. [PR1539537](#)
- IS-IS over Layer 2 circuit will not come up if the encapsulation is TCC. [PR1590387](#)

Layer 2 Ethernet Services

- The copying of files to the RCB over WAN ports is slow. [PR1496895](#)

MPLS

- At high scale, LSP setup rate will be relatively slower in IP-in-IP networks. [PR1457992](#)

Routing Protocols

- In a Layer 3 VPN scenario, the rpd process on the backup Routing Engine might crash when BGP (standby) received a VPN route from the peer that is rejected due to invalid target community, because the BGP standby peer synchronization is not complete yet. [PR1508888](#)

SEE ALSO

[What's New | 206](#)

[What's Changed | 218](#)

[Known Limitations | 224](#)

[Resolved Issues | 229](#)

[Documentation Updates | 234](#)

[Migration, Upgrade, and Downgrade Instructions | 235](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.3R3 | 230](#)
- [Resolved Issues: 20.3R2 | 231](#)
- [Resolved Issues: 20.3R1 | 232](#)

Learn which issues were resolved in the Junos OS Release 20.3R3 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: 20.3R3

General Routing

- On PTX10008 router, FPC UKERN core file is not transferred to Routing Engine in a scaled setup. [PR1500418](#)
- The following error message might be seen after link flap: **t6e_dfe_tuning_state:et-6/0/0 - Failed to dfe tuning count 10.** [PR1512919](#)
- Traffic might get dropped when the **set routing-options forwarding-table no-ecmp-fast-reroute** configuration is changed to 128 ECMP entries. [PR1547457](#)
- The interface filter with source-port 0 matches everything instead of port 0. [PR1551305](#)
- An enhancement to enable watchdog petting log on PTX10K line cards. [PR1561980](#)
- Upgrading PTX1000 with unified SSDs (2x32G SSD) might result in boot loop in certain scenario. [PR1571275](#)
- Traffic loss might be observed on the PTX5000 platform. [PR1578511](#)
- TACACS traffic might get dropped on PTX platforms. [PR1578579](#)
- When **rpf-check** is configured on PE chip based PTX platforms, the traffic might be dropped. [PR1580211](#)
- On PTX5000 and PTX3000 routers, configuring and deleting the FEC mode disables the auto-FEC91 on an interface that uses QSFP28-SR4. [PR1582200](#)
- Node locked license addition fails. [PR1582704](#)
- The following error message is seen during bring up: **Failed to get pechip handle for chip 0" and "prds_encap_sample_flood_lpbk_desc_install: Egress NH descriptor install OK for Flabel 7808.** [PR1585594](#)

Infrastructure

- The kernel crash with core file might be seen if churn happens for a flood composite next hop. [PR1548545](#)

Multicast

- FPC might crash in a multicast scenario. [PR1569957](#)

Network Management and Monitoring

- The mib2d process crashes and generates a core file on backup Routing Engine. [PR1557384](#)

Platform and Infrastructure

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)

- FPC crash might be observed in a scaled firewall configuration on PTX Series platforms. [PR1586817](#)

Routing Policy and Firewall Filters

- Generated route goes to the hidden state when the protect core command is enabled. [PR1562867](#)

Routing Protocols

- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. [PR1514966](#)
- BGP LU session flap might be seen with the AIGP used scenario. [PR1558102](#)

Resolved Issues: 20.3R2

General Routing

- On PTX10016 routers, if aggregated Ethernet member or interface flow control is in disabled state, then it does not enable its own. [PR1478715](#)
- SNMP index in the Packet Forwarding Engine reports as 0. This causes the sFlow records to have either IIF (Input interface value) or OIF (Output interface value) as 0 value in sFlow record data at collector. [PR1484322](#)
- On PTX5000 and PTX3000 routers, the FPC E might get stuck. [PR1519673](#)
- The chassisd memory leak might cause traffic loss. [PR1537194](#)
- The error message `expr_dfw_action_topo_connect_anh:1434`
`expr_dfw_action_topo_connect_anh:eda_anh_discard` is FALSE for nh-id 568 - return is observed in PTX1000 routers. [PR1540064](#)
- The Packet Forwarding Engine might crash in MPLS IPv6-tunneling scenario when the next hop changes. [PR1540793](#)
- Traffic might be silently dropped and discarded after swapping an FPC type 3 card with an FPC type 1 card in the same slot on a PTX3000 router. [PR1547790](#)
- RPD crash might be seen when BGP service route is resolved over color-only SR-TE policy. [PR1550736](#)
- Interface filter with source-port 0 is matching everything instead of just port 0. [PR1551305](#)
- An enhancement to enable watchdog petting log on the PTX10000 line cards. [PR1561980](#)

Forwarding and Sampling

- The l2ald process might crash due to next hop issue in the EVPN-MPLS. [PR1548124](#)

Infrastructure

- Output might display 0 temporarily during a race condition for `show interfaces extensive` command when SNMP query for JnxCos is issued. [PR1533314](#)

- Interface drop counters might display 0 during a race condition and voq statistics are also polled simultaneously. [PR1537960](#)
- The kernel crashes and core file generates if churn happens for a flood composite next hop. [PR1548545](#)

Interfaces and Chassis

- EOAM IEEE802.3ah link discovery state is **Down** instead of **Active Send Local** after deactivating interfaces on routers. [PR1532979](#)
- Logs not being written in `/var/log/messages` on certain PTX platforms. [PR1551374](#)

Network Management and Monitoring

- The syslog messages might not be sent with correct port. [PR1545829](#)

Platform and Infrastructure

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)

Routing Policy and Firewall Filters

- Generate route goes to hidden state when **protect core** statement is enabled. [PR1562867](#)

Routing Protocols

- The rpd process generates the core file at `gp_rtargt_tsi_update,bgp_rtargt_flash_rt,bgp_rtargt_flash`. [PR1541768](#)
- BGP LU session might flap with AIGP scenario. [PR1558102](#)

Resolved Issues: 20.3R1

General Routing

- On PTX5000 and PTX10008 routers, the `show filter index < number> counter` vty command displays values as zero at `28-02-HOSTBOUND_NDP_DISCARD_TERM`. [PR1420057](#)
- The `show snmp mib walk jnxContentsDescr` command does not show fan controllers. [PR1455640](#)
- PHP device has NH mis-programming for members of ECMP for SR label route used for reaching IPv6 destinations. [PR1457230](#)
- The PTX1000 and PTX10002 routers might discard traffic silently after the transient SIB or FPC voltage alarms. [PR1460406](#)
- Optics-options syslog and link-down do not work as expected on PTX5000 with FPC3. [PR1461404](#)
- The router might become nonresponsive and bring traffic down when the disk space becomes full. [PR1470217](#)

- On PTX10016 routers, after device reboot, the FPC takes a long time to come up and hence MKA sessions establishment is delayed. The error message **Frame 08: sp = 0x48d222b8, pc = 0x10fad3bc , blaze fpc2 SCHED: Thread 59 (PFE Manager) ran for 2177 ms without yielding** is observed. [PR1477585](#)
- Disk usage might keep increasing on PTX1000 platforms. [PR1480217](#)
- LSP auto-bandwidth adjust-interval change does not get detected on commit in some cases. [PR1484801](#)
- The Layer 2 VPN might flap and the CE-facing interface cannot restore the TX optical laser power even if the Layer 2 VPN is up under asynchronous-notification. [PR1486181](#)
- Dynamic tunnels trace options do not offer state tracing and cause JTASK_SCHED_SLIP with single underlay route bounce. [PR1493236](#)
- Kernel routing table queue become nonresponsive after J-Flow sampling of a malformed packet. [PR1495788](#)
- Outbound SSH connection flaps or a memory leak issue during push configuration to ephemeral database with high rate. [PR1497575](#)
- Packet drop is observed following an RSVP load-balance configuration on PTX10003 routers. [PR1500711](#)
- Routes are being installed in the Packet Forwarding Engine even when the interface is down or disabled. [PR1501321](#)
- An error message **PFE_ERROR_FAIL_OPERATION: IFD et-1/0/8: RS credits failed to return: init=192 curr=193 chip=5** is observed. [PR1502716](#)
- When you want to delete a YANG package, event-options (if configured) hierarchy has to be deactivated before issuing the **request system yang delete** command. [PR1502939](#)
- On a dual Routing Engine GRES or NSR enabled PTX10008 or PTX10016 router, a few TCP-based application sessions such as BGP or LDP might flap upon Routing Engine mastership switch. [PR1503169](#)
- Unable to bring the ports up when plugging the optic QSFP-100G-LR4-T2 (740-061409) into PTX3000 or PTX5000 routers. [PR1511492](#)
- The routes update might fail because of an HMC memory issue, and traffic impact might be seen. [PR1515092](#)
- On PTX10002-60C and PTX1000 routers, sFlow adaptive-sampling with rate limiter statement enabled, crosses the sample rate 65535. [PR1525589](#)

Interfaces and Chassis

- When multiple CFM sessions are configured on a physical interface, SNMP walk of ieee8021CFMStack table fails. [PR1517046](#)

MPLS

- 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 in a rare condition in an SR-TE scenario. [PR1493721](#)

- If the automatic bandwidth adjustment fails due to bandwidth unavailability, during the subsequent retries, it tries to bring up the LSP with the same bandwidth that was last requested. [PR1504916](#)
- SNMP trap is sent with incorrect OID jnxSpSvcSetZoneEntered. [PR1517667](#)

Network Management and Monitoring

- SNMP response packets have Don't Fragment (DF) flag set by default. [PR1514156](#)

Routing Protocols

- The **show dynamic-tunnels database** command does not reflect the current value of traffic statistics. It shows the cached value of traffic statistics, which might not be equal to the current value. [PR1445705](#)
- On PTX3000 and PTX5000 routers, the ppmmd process generates a core file after configuring the sbfd responder on the RE-DUO-2600. [PR1477525](#)
- The BGP route-target family might prevent the route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)

SEE ALSO

What's New 206
What's Changed 218
Known Limitations 224
Open Issues 227
Documentation Updates 234
Migration, Upgrade, and Downgrade Instructions 235

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for PTX Series routers.

SEE ALSO

What's New 206
What's Changed 218
Known Limitations 224
Open Issues 227

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.3 | 235](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 238](#)
- [Upgrading a Router with Redundant Routing Engines | 239](#)

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 20.3

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 20.3R3:

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-20.3R3.9.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-20.3R3.9-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 20.3 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

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

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

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 | 206](#)

[What's Changed | 218](#)

[Known Limitations | 224](#)

[Open Issues | 227](#)

[Resolved Issues | 229](#)

[Documentation Updates | 234](#)

Junos OS Release Notes for QFX Series

IN THIS SECTION

● [What's New | 240](#)

● [What's Changed | 253](#)

● [Known Limitations | 259](#)

- Open Issues | 261
- Resolved Issues | 265
- Documentation Updates | 280
- Migration, Upgrade, and Downgrade Instructions | 281

These release notes accompany Junos OS Release 20.3R3 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 20.3R3 | 241
- What's New in Release 20.3R2 | 241
- What's New in Release 20.3R1 | 241

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

NOTE: The following QFX Series platforms are supported in Release 20.3R3: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, 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 20.3R3

There are no new features or enhancements to existing features for QFX Series Switches in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features or enhancements to existing features for QFX Series Switches in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Hardware

- We've added the following features to the QFX5120-48T in Junos OS Release 20.3R1.

Table 7: Features Supported by the QFX5120-48T

Feature	Description
Firewall filters and policers	<ul style="list-style-type: none"> • Support for MPLS firewall filter on loopback interface. MPLS firewall filter can be applied to a loopback interface on a label-switching router (LSR). [See Overview of MPLS Firewall Filters on Loopback Interface.] • Support for flexible-match-mask match condition. Flexible-match-mask match condition allows you to filter by specifying the length of the match (4 bytes maximum) starting from a Layer 2 or Layer 3 packet offset. [See Firewall Filter Flexible Match Conditions.]
Timing and synchronization	<ul style="list-style-type: none"> • Precision Time Protocol (PTP) transparent clock is supported on the QFX5120-48T. [See Transparent Clock Overview.]

- **Support for QSFP-100G-DR transceivers (QFX5200, QFX5120-32C, QFX5120-48Y, QFX10002-72, and QFX10002-60C)**—Starting in Junos OS Release 20.3R1, we provide support for the QSFP-100G-DR transceivers. These transceivers interoperate with 400-Gbps breakout optics. For example, the QDD-400G-DR4 interconnects with up to four QSFP-100G-DR transceivers. The QSFP-100G-DR transceivers interconnect in single links (QSFP-100G-DR to QSFP-100G-DR or to QSFP-100G-FR) and interoperate at the shortest link length.

NOTE: The QSFP-100G-DR transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

- **Support for the QSFP-4X10GE-SR and JNP-QSFP-4X10GE-LR transceivers (QFX5120)**—Starting in Junos OS Release 20.3R1, QFX5120 switches support the QSFP-4X10GE-SR and JNP-QSFP-4X10GE-LR transceivers.

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

- **Support for QSFP-100G-FR transceivers (QFX5200 and QFX10002-72)**—Starting in Junos OS Release 20.3R1, we provide support for the QSFP-100G-FR transceivers. These transceivers interoperate with the QDD-4X100G breakout optics. For example, the QDD-4X100G-FR interconnects with up to four QSFP-100G-FR transceivers. The QSFP-100G-FR transceivers interconnect in single links (QSFP-100G-FR to QSFP-100G-FR or to QSFP-100G-DR) and interoperate at the shortest link length. The QSFP-100G-FR transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).

NOTE: The QSFP-100G-FR transceivers are not compatible with earlier-generation 100-Gbps transceivers (for example, QSFP-100G-CWDM4 and QSFP-100G-LR4).

[See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

EVPN

- **Support for creating remote VXLAN VTEP per underlay (QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, and QFX10016)**—Starting in Junos OS Release 20.3R1, you can create one VTEP logical interface per remote provider edge (PE) device, regardless of the number of routing instances. For example, if there are X number of PE devices and Y number of routing instances, this would currently result in having $(X - 1) * Y$ remote VTEPs. Starting in Junos OS Release 20.3R1, however, there are only $X - 1$ remote VTEPs. This change reduces the number of next hops and hardware tokens from the quadratic level to the linear level.

For existing platforms that support EVPN-VXLAN, configure the **shared-tunnels** statement at the **[edit forwarding-options evpn-vxlan]** hierarchy level. For changes to take effect, reboot the device.

- **Seamless EVPN-VXLAN stitching (QFX10002-36Q, QFX10002-72Q, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.3R1, we support the seamless stitching of unicast and broadcast, unknown unicast, and multicast (BUM) routes in the following use cases:
 - Interconnected EVPN-VXLAN points of delivery (PODs) in a data center.
 - Interconnected EVPN-VXLAN data centers (data center interconnect [DCI]).

NOTE: We do not currently support the assisted replication of BUM traffic in the described use cases.

In these use cases, the QFX10000 switches, either single-homed or multihomed in all-active mode, can serve as either a spine or super spine device that interconnects the PODs or data centers through an EVPN-VXLAN WAN network.

When configuring the interconnection, you can set up a single routing instance of type **virtual-switch** or **evpn** on each spine or super spine device. Or, you can use the default switching instance. In this instance, you include elements described in [interconnect](#).

After you configure the interconnection, the EVPN control plane stitches the EVPN routes from the POD or data center network and from the WAN network into a single MAC forwarding table.

- **Enhancement in the number of supported VLANs and ports (QFX5110 and QFX5120)**—Starting with Junos OS Release 20.3R1, we've increased the combined total number of VLANs and ports that can be supported on the QFX5110 and QFX5120 switches. The number of supported VLANs remains at 4093, but Junos OS no longer limits the number of ports that can be configured in conjunction with the number of configured VLANs on EVPN-VXLAN. This enhancement applies only when you use the enterprise style of configuration when configuring the interfaces.

[See [Understanding EVPN with VXLAN Data Plane Encapsulation](#).]

- **Filter-based forwarding in EVPN-VXLAN networks (QFX5110 and QFX5120)**—Starting in Junos OS Release 20.3R1, QFX5110 and QFX5120 switches support the use of firewall filters along with routing instances to specify different routes for IPv4 VXLAN-encapsulated traffic in your EVPN-VXLAN network.

To set up this feature:

- Create an input filter.
- Specify one or more of these match criteria:
 - Source or destination IP address
 - Source or destination Layer 4 port
 - Time to live (TTL)
 - IP protocol
- For the action, specify the routing instance to which to send packets. (We also support the accept, count, and discard actions.)
- Apply the filter to an IRB interface with or without a virtual gateway address or an anycast address.

For example:

```
set firewall family inet filter filter-irb term t1 from source-address 192.168.1.2/32
set firewall family inet filter filter-irb term t1 then count FBF-1-packet-count
set firewall family inet filter filter-irb term t1 then routing-instance FBF-1
set interfaces irb unit 10 family inet filter input filter-irb
```

When the Juniper Networks switch receives incoming traffic from the specified address on interface `irb.10`, it counts and then forwards the traffic to the FBF-1 routing table. According to the routing table, the packet is forwarded to the next hop that corresponds to the destination address entry in the table.

[See [Understanding Filter-Based Forwarding](#).]

- **Dynamic load balancing in an EVPN-VXLAN network (QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, and QFX5220)**—Starting with Junos OS Release 20.3R1, the listed QFX switches support dynamic load balancing in an EVPN-VXLAN network. When your EVPN-VXLAN network includes a multihomed device that can be reached through multiple virtual tunnel endpoints (VTEPs) that share a common Ethernet segment identifier (ESI), dynamic load balancing works as follows:
 - The EVPN control plane (overlay) identifies the common ESI as the next hop for a destination device with a particular MAC address.
 - Based on the parameters in a packet, the forwarding plane in the switch (hardware) dynamically chooses one of the VTEPs associated with the ESI. The VTEP then forwards the packet along the selected underlay path to the destination device.

By default, the listed QFX switches have dynamic load balancing enabled.

[See [Dynamic Load Balancing in an EVPN-VXLAN Network](#).]

- **Increased number of ARP and neighbor discovery entries and token spaces for IRB and aggregated Ethernet interfaces (QFX10002-60C)**—Starting in Junos OS Release 20.3R1, we've increased the number of token spaces to 96,000, and the number of ARP and neighbor discovery entries to 256,000. We've also enabled both 96,000 token spaces and 256,000 ARP and neighbor discovery entries by default for VXLAN Layer 3 gateway scenarios. The token spaces are also shared with the ARP and neighbor discovery entries, which helps with the default ARP scale as well as with multidimensional scale.

To disable the sharing of token spaces with the ARP and ND entries, enable the **no-arp-enhanced** statement at the **[edit system]** hierarchy level. Reboot the device for changes to take effect.

[See [Increasing ARP and Network Discovery Protocol Entries for Enhanced MC-LAG and Layer 3 VXLAN Topologies](#).]

- **Layer 2 egress filtering on EVPN-VXLAN interfaces (QFX5110 and QFX5120)**—Starting in Junos OS Release 20.3R1, QFX5110 and QFX5120 switches support the filtering of Layer 2 traffic exiting access interfaces on which EVPN-VXLAN is running.

To set up this feature:

- Create a Layer 2 egress filter.
- In the filter, specify one or more of these match criteria:
 - Source or destination MAC address
 - Ethernet type
 - VLAN ID

- Specify one or more of these actions:
 - Accept
 - Count
 - Discard
- Apply the filter to a physical interface or an aggregated Ethernet interface.

The following sample configuration creates a Layer 2 egress firewall filter named `epacl`, which you apply to interface `xe-0/0/10.0`. The first term specifies that the interface accepts and counts packets from source MAC address `00:00:5e:00:53:a1/48`. The second term specifies that the interface discards all other packets and counts them.

```
set firewall family ethernet-switching filter epacl term t1 from source-mac-address 00:00:5e:00:53:a1/48
set firewall family ethernet-switching filter epacl term t1 then accept
set firewall family ethernet-switching filter epacl term t1 then count epacl-accept
set firewall family ethernet-switching filter epacl term t2 then discard
set firewall family ethernet-switching filter epacl term t2 then count epacl-discard
set interfaces xe-0/0/10 unit 0 family ethernet-switching filter output epacl
```

[See [Overview of Firewall Filters \(QFX Series\)](#).]

High Availability (HA) and Resiliency

- **Higher scale and performance in RIFT (MX240, MX480, MX960, vMX, QFX5100, QFX5110, QFX5120-32C, QFX5120-48T, QFX5120-48T-VC, QFX5120-48Y, QFX5120-48Y-VC, QFX5120-24YM, QFX5120-48YM, QFX5130-48C, QFX5200, QFX5210, and QFX10008)**—Starting in Junos OS Release 20.3R1, we've made the following improvements to increase the scalability and performance in Routing in Fat Tree (RIFT):
 - Prefixes in RIFT
 - Peers in RIFT
 - Convergence improvement with RIFT
 - BFD sessions with RIFT

[See [RIFT Overview](#).]

IP Tunneling

- **Support for IP over IP next hop based tunneling (MX Series, PTX1000, PTX10000, and QFX10000)**—Starting in Junos OS Release 20.3R1, we support an IP-over-IP encapsulation to facilitate IP overlay construction over IP transport network. An IP network contains edge devices and core devices. To achieve higher scale and reliability among these devices, you need to logically isolate the core network from the external network that the edge devices interact with, by using an overlay encapsulation. Among the other overlay encapsulations supported, IP over IP encapsulation is the only kind where transit

devices are able to parse the inner payload and use inner packet fields for hash computation and customer edge devices are able to route traffic into and out of the tunnel without any throughput reduction. IP over IP relies on a next hop-based infrastructure to support higher scale.

On MX Series routers, routing protocol daemon(RPD) sends the encapsulation header with tunnel composite nexthop and the Packet Forwarding Engine finds the tunnel destination address and forwards the packet. On PTX Series routers and QFX10000 switches, RPD sends fully resolved next hop-based tunnel to PFE. BGP protocol is used to distribute routes and signal dynamic tunnels.

[See [Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation](#).]

- **Support for filter-based decapsulation of IPv4 and IPv6 unicast traffic encapsulated in IPv4 IP-in-IP tunnels (MX Series, PTX1000, PTX10002, and QFX10002)**—Junos OS supports decapsulating IPv4 and IPv6 unicast traffic that has been encapsulated in IPv4 IP-in-IP tunnels using firewall filters. If the outer IPv4 header address matches the firewall configuration and the packet has **ipip** set as the protocol type, then the outer IPv4 header is removed and the packet is routed based on the inner IPv4 or IPv6 address. If the packet does not have the expected **ipip** header, the packet is dropped.

Configure this feature using the following CLI statements at the **[edit firewall family inet filter filter-name term term-name]** hierarchy:

- **from protocol ipip**: Set the protocol type as IP-IP.
- **then decapsulate ipip**: Decapsulate the IP-IP packet. The inner IP destination address is routed using the inet.0 routing table by default.
- **then decapsulate ipip routing-instance routing-instance-name**: Decapsulate the IP-IP packet and route the inner destination address using the specified routing instance.

Use **show firewall** to view the configuration.

[See [filter \(Firewall Filters\)](#) and [Configuring IP Tunnel Interfaces](#).]

Juniper Extension Toolkit (JET)

- **Juniper Extension Toolkit (JET) supports BFD Service APIs for routing protocol process (rpd) programmability (MX Series, PTX Series, QFX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can use programmable rpd (prpd) BFD APIs to add, update, and delete BFD sessions and subscribe to BFD events from outside applications. These APIs enable the integration of rpd with software-defined networking (SDN) controllers and increase the flexibility of your network. The prpd BFD APIs support BFD Echo-Lite sessions in single-hop IPv4 and IPv6 modes.

The following BFD Service APIs are supported:

- Initialize
- SessionAdd
- SessionUpdate
- SessionDelete

- SessionDeleteAll
- Subscribe
- Unsubscribe

Use the **show bfd session extensive** command to view BFD sessions. BFD sessions added through prpd BFD APIs are labeled with **PRPD:<session-id>** in the client field. The **<session-id>** is 1 for the first BFD session that is added, 2 for the second, and so on.

[See [show bfd session extensive](#) and [JET APIs on Juniper EngNet.](#)]

alpha

- **Support for REST API over nondefault virtual routing and forwarding (VRF) instance (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can execute Junos OS operational commands using the REST API over a nondefault VRF instance. The nondefault VRF instance can be a user-defined instance or the management instance, mgmt_junos.

The REST API allows you to execute Junos OS operational commands over HTTP(S). If you don't specify a routing instance, REST API requests are sent over the default routing instance. Use a nondefault VRF instance to improve security and make it easier to troubleshoot.

Use the **routing-instance routing-instance** statement at the **[edit system services rest]** hierarchy level to specify a nondefault VRF instance for REST API requests.

[See [Management Interface in a Nondefault Instance](#) and [rest.](#)]

Junos Telemetry Interface

- **Support for aggregated Ethernet interface ON_CHANGE with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX3000, PTX5000, PTX10001-36MR, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.3R1, Junos telemetry interface (JTI) supports ON-CHANGE statistics for aggregated Ethernet interfaces for minimum links and member interfaces.

To export these statistics to an outside collector using remote procedure call (gRPC) services and JTI, include the following resource paths in a subscription:

- **/interfaces/interface/aggregation/state/min-links/**
- **/interfaces/interface/aggregation/state/member/**

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

Layer 3 Features

- **Support for BGP Layer 3 VPN over IP-IP Tunnel (MX Series, PTX1000, QFX10002, and QFX10008)**—Starting in Junos OS Release 20.3R1, we support BGP Layer 3 VPN over IP over IP (IP-IP) tunnels to create a new transport service. IP-IP tunnels terminate into service-layer VRF, so you do not need to use a service label. This feature allows interoperability between the new VRF and traditional VRF, so both types of overlays can coexist in your network. You can use this feature to transition from an MPLS network to an IP fabric core network and to protect your network from distributed denial-of-service (DDoS) attacks.

To use VPN over an IP-IP tunnel, configure the **tunnel-attribute** statement at the **[edit policy-options policy-statement *policy-name* term *term-name* then]** or **[edit policy-options policy-statement *policy-name* then]** hierarchy level.

To configure the receiver to program the dynamic tunnel using the tunnel attribute, use the **extended-nexthop-tunnel** statement at the **[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family (inet-vpn | inet6-vpn) unicast]** hierarchy level.

[See [BGP Layer 3 VPN over IP-IP Tunnels Overview](#), [family \(Protocols BGP\)](#), [policy-statement](#), [vrf-export](#), and [Configuring IP Tunnel Interfaces](#).]

MPLS

- **New output fields added in the show path-computation-client lsp extensive command (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.3R1, you'll see association details such as **Association type**, **ID**, and **source** in the output of the **show path-computation-client lsp** command when you use the command with the **extensive** option.

[See [show path-computation-client lsp](#).]

Network Management and Monitoring

- **Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can use the **probe** command to query the status of the probed interface. The proxy interface resides on the same node as the probed interface, or it can reside on a node to which the probed interface is directly connected.

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

To enable the **probe** command, configure the **extended-echo** statement under the **[edit system]** hierarchy.

[See [Using the Probe command](#).]

- **Remote port mirroring to an IP address (using GRE) with ToS and DSCP (QFX10002, QFX10008, and QFX10016)**—You use port mirroring to send traffic to applications that analyze traffic to monitor compliance, enforce policies, detect intrusions, and so on. Starting in Junos OS Release 20.3R1, you can

configure remote port mirroring to send sampled packets to a remote IP address. You send the packets using GRE. You can set type-of-service (ToS) and DiffServ code point (DSCP) values to provide the necessary priorities in the network for these packets. You can also apply policing to sampled packets that are leaving the interface where the GRE destination was learned. Configure the settings you need in the **[edit forwarding-options port-mirroring instance *instance-name* output]** hierarchy.

[See [instance \(Port Mirroring\)](#) and [traffic-class \(Tunnels\)](#).]

- **Enhancements to sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.3R1, devices running Junos OS with upgraded FreeBSD support the following enhancements to sessions over outbound HTTPS:
 - Connecting to multiple outbound HTTPS clients by configuring one or more clients at the **[edit system services outbound-https]** hierarchy level
 - Configuring multiple backup gRPC servers for a given outbound HTTPS client
 - Establishing a csh session
 - Establishing multiple, concurrent NETCONF and csh sessions between the device running Junos OS and an outbound HTTPS client
 - Configuring a shared secret that the outbound HTTPS client uses to authenticate the device running Junos OS
 - Authenticating the client using certificate chains in addition to self-signed certificates

[See [NETCONF and Shell Sessions over Outbound HTTPS](#).]

Routing Policy and Firewall Filters

- **Loopback firewall filter scale optimization (EX4650 and QFX5120-48Y)**—Starting with Junos OS Release 20.3R1, you can configure up to 768 loopback filter terms for IPv6, and up to 1152 terms for IPv4. To do so, you configure an ingress firewall filter, apply it to the loopback interface, and then enable the **loopback-firewall-optimization** statement at the **[edit chassis]** hierarchy level (this triggers the Packet Forwarding Engine to restart).

The switches do not support terms that include a reserved multicast destination, for example 224.0.0.x/24, and terms with a time-to-live (TTL) of 0/1. You need to configure a separate filter for these terms. So, for example, to count OSPF packets on the loopback interface, you would create a separate filter with terms for the protocol (OSPF) to count packets destined to a reserved multicast address (such as 224.0.0.6).

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

Routing Protocols

- **PTP over IRB (QFX-5110-48s and QFX-5200-32q)**—Starting in Junos OS Release 20.3R1, we support PTP boundary clock to IRB interfaces for PTP over multicast for broadcast profiles.

[See [Understanding IEEE 1588 Precision Timing Protocol \(PTP\) over IRB for Broadcast profiles](#).]

- **ECMP nexthop update rate throttling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.3R1, 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.

[See [pause-computation-during-churn](#).]

Security

- **Source MAC filtering on aggregated Ethernet interfaces (QFX5100, QFX5120-32C, and QFX5120-48Y switches)**—Starting in Junos OS Release 20.3R1, you can configure source media access control (MAC) filtering on an aggregated Ethernet interface on QFX5100, QFX5120-32C, and QFX5120-48Y switches. Ingress packets are matched on the source MAC address list you have configured under the **accept-source-mac mac-address** hierarchy level on the logical interface of the aggregated Ethernet interface.

[See [Understanding MAC Limiting on Layer 3 Routing Interfaces](#) and [accept-source-mac](#).]

Services Applications

- **Support for IPv4 and IPv6 inline active flow monitoring (QFX10002-60C)**—Starting in Junos OS Release 20.3R1, you can configure inline active flow monitoring for IPv4 and IPv6 traffic. We support both the IPFIX and the version 9 formats of the IPv4 and IPv6 templates. To configure the template properties for inline active flow monitoring, configure the options for the **flow-monitoring (version-ipfix | version9) template *template-name*** statement at the **[edit services]** hierarchy level.

[See [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#).]

Software Licensing

- **Juniper Agile Licensing (QFX5120 and QFX5200)**—Starting in Junos OS Release 20.3R1, we’re moving toward license-based software features. We now use Juniper Agile Licensing to support soft enforcement for software features on the listed devices.

Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can install and manage licenses for hardware and software features using Juniper Agile Licensing.

From this release onwards, you can now opt to use the Juniper Agile License Manager to significantly improve the ease of license management for an entire network of supported devices.

If you are upgrading to this release, you need new license keys to use the features on the listed devices. Contact [Customer Care](#) to exchange license keys for Junos OS releases earlier than Junos OS Release 20.3R1.

[Table 8 on page 252](#) describes the licensing support on the QFX5120 and QFX5200 devices.

Table 8: Licensed Features on the QFX5120 and QFX5200 Devices

QFX Switch License Model	Detailed Features
Standard license for integrated SKUs (standard hardware and software platform)	Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), QoS (Layer 2 and Layer 3), and SNMP
Advanced license for integrated and advanced SKUs	Advanced 1: BGP, FBF, GRE, IS-IS, JTI, MC-LAG, OSPF, sFlow, VRF, and VRRP
	Advanced 2: Includes Advanced 1 features + CFM, Layer 2 and Layer 3 multicast, OAM, Packet Timestamping, PTP, and Q-in-Q PTP is supported only on QFX5120-48Y and QFX5200-32C.
Premium license for integrated and premium SKUs	Includes Advanced 2 features + EVPN-MPLS, MPLS, Layer 2 circuit, Layer 3 VPN, LDP, RSVP, segment routing, and SR-TE

[See [Supported Features on QFX5120 and QFX5200 Devices](#), [Juniper Agile Licensing Guide](#), [Configuring Licenses in Junos OS](#), and [Managing Licenses](#).]

Virtual Chassis

- **Support for Virtual Chassis (QFX5120-32C)**—Starting in Junos OS Release 20.3R1, you can interconnect two QFX5120-32C switches into a Virtual Chassis managed as a single device. The Virtual Chassis:
 - Contains only QFX5120-32C switches.
 - Has two member switches in the Routing Engine role (one master and one backup).

- Supports any of the 32 network ports installed with 100-Gbps QSFP28 or 40-Gbps QSFP+ transceivers as Virtual Chassis ports (VCPs) to connect the member switches.
- Supports NSSU.

A QFX5120-32C Virtual Chassis supports the same protocols and features as the standalone switches in Junos OS Release 20.3R1, except for the following:

- EVPN-VXLAN
- Junos telemetry interface (JTI)
- Multichassis link aggregation (MC-LAG)

Configuration and operation are the same as for other QFX Series Virtual Chassis.

[See [Virtual Chassis Overview for Switches](#).]

SEE ALSO

What's Changed 253
Known Limitations 259
Open Issues 261
Resolved Issues 265
Documentation Updates 280
Migration, Upgrade, and Downgrade Instructions 281

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 254](#)
- [What's Changed in Release 20.3R2 | 255](#)
- [What's Changed in Release 20.3R1 | 257](#)

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

What's Changed in Release 20.3R3

Interfaces and Chassis

- **Blocking duplicate IP detection in the same routing instance (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. Refer to the table mentioned in the topic [inet \(interfaces\)](#). When you try to configure same IP on two logical interfaces inside same routing instance, the commit will be blocked with the error displayed as shown below: `edit user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24 edit user@host# commit commit complete edit user@host# set interfaces ge-0/0/2 unit 0 family inet address 2.2.2.2/24 edit user@host# commit edit interfaces ge-0/0/2 unit 0 family inet 'address 2.2.2.2/24' identical local address found on rt_inst default, intfs ge-0/0/2.0 and ge-0/0/1.0, family inet. error: configuration check-out failed`

[See [inet\(interfaces\)](#).]

Junos OS XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or **set refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

Layer 2 Ethernet Services

- **Link selection support for DHCP (QFX Series)**—We've introduced **link-selection** statement at the **edit forwarding-options dhcp-relay relay-option-82** hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Earlier to this release, the DHCP relay drops packets during the renewal DHCP process as the DHCP Server uses the leaf's address as a destination to acknowledge DHCP renewal message.

[See [relay-option-82](#).]

Network Management and Monitoring

- **Change in OID ifHighSpeed**—Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

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

Platform and Infrastructure

- **Juniper Agile Licensing (QFX5120-48Y, QFX5110-32Q, and QFX5110-48S)**—Starting from this release onwards, the QFX switch supports following features:
 - **Standard:** BFD, Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), QoS (Layer 2 and Layer 3), and SNMP
 - **Advanced 1:** Standard features, BGP, IS-IS, FBF, VRRP, MC-LAG, Layer 3 (static), GRE tunnel, OSPF, RIP, sFlow, and Virtual Chassis
 - **Advanced 2:** Advanced 1 features, CFM, Q-in-Q, VXLAN, PCEP, ESI-LAG, Timing, Ethernet OAM, EVPN-VXLAN, IGMP version 1, IGMP version 2, and IGMP version 3, PIM, and Multicast Listener Discovery (MLD) version 1 or version 2
 - **Premium:** Advanced 2 features, Layer 3 VPN, LDP, RSVP, Layer 2 circuit, EVPN-MPLS, Segment routing, MPLS, and MACsec

[See [Flex Software License for QFX Series Switches](#) and [Juniper Agile Licensing Guide](#).]

What's Changed in Release 20.3R2

General Routing

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

- **Only manual channelization support on QSFP-100G-SR4-T2 optics (QFX5120-48T and QFX5120-32C)**— We recommend you to use the Active Optical Cable (AOC) for auto-channelization. The QSFP-100G-SR4-T2 cables do not support auto-channelization. If you want to use the QSFP-100G-SR4-T2 optics with an external breakout cable, then you must configure the channelization manually by running **channel-speed** statement at the **edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)** hierarchy level.

[See [channel-speed](#).]

- **Support only for manual channelization on QSFP-100G-SR4-T2 optics (QFX5120-48T and QFX5120-32C)**— We recommend that you use the active optical cable (AOC) for auto-channelization. The QSFP-100G-SR4-T2 cables do not support auto-channelization. To use the QSFP-100G-SR4-T2 optics with an external breakout cable, you must configure the channelization manually by running the

channel-speed statement at the **edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)** hierarchy level.

[See [channel-speed](#).]

Junos XML API and Scripting

- The **jcs:invoke()** function supports suppressing root login and logout events in system log files for **SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log **UI_LOGIN_EVENT** and **UI_LOGOUT_EVENT** messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root **UI_LOGIN_EVENT** and **UI_LOGOUT_EVENT** messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- The **jcs:invoke()** function supports suppressing root login and logout events in system log files for **SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The **jcs:invoke()** extension function supports the **no-login-logout** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log **UI_LOGIN_EVENT** and **UI_LOGOUT_EVENT** messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root **UI_LOGIN_EVENT** and **UI_LOGOUT_EVENT** messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

MPLS

- **The show mpls lsp extensivel and show mpls lsp detail commands display next-hop gateway LSPid** — When you use the **show mpls lsp extensivel** and **show mpls lsp detail** commands, you'll see next-hop gateway LSPid in the output.

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the **verbose** statement at the **edit system export-format json** hierarchy level. The default format to export configuration data in JSON changed from **verbose** format to **ietf** format starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **edit system export-format json** hierarchy level. Although the **verbose** statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See [export-format..](#)]

What's Changed in Release 20.3R1

Class of Service (CoS)

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

General Routing

- **Priority-based flow control (PFC) support (QFX5120-32C)**—Starting with Junos OS Release 20.3R1, QFX-5120-32C switches support priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic.

High Availability (HA) and Resiliency

- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the **show rift tie** output.

Interfaces and Chassis

- **Autonegotiation status displayed correctly (QFX5120-48Y)**—In Junos OS Release 20.3R1, 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.

Junos OS XML, API, and Scripting

- **Changes to Junos XML RPC request tag names (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've updated the Junos XML request tag name for some operational RPCs to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include:
 - Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
 - Uppercase characters are converted to lowercase.

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

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.
- **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.

SEE ALSO

[What's New | 240](#)

[Known Limitations | 259](#)

[Open Issues | 261](#)

[Resolved Issues | 265](#)

[Documentation Updates | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)

Known Limitations

IN THIS SECTION

- [Class of Service \(CoS\) | 259](#)
- [Layer 2 Features | 259](#)
- [Layer 2 Ethernet Services | 259](#)
- [Platform and Infrastructure | 260](#)
- [Routing Protocols | 260](#)

Learn about known limitations in Junos OS Release 20.3R3 for QFX Series Switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

Class of Service (CoS)

- Traffic might be drop by the destination device. [PR1568333](#)

Layer 2 Features

- On the QFX5000 devices with storm control, significant difference between the configured rate and actual rate is observed. [PR1526906](#)

Layer 2 Ethernet Services

- If the configuration or image file name has special characters (for example: #, %, and @), ZTP over http or https does not work. [PR1503588](#)
- The DHCPv6 client binding does not occur after the image is upgraded and rebooted in an image + script scenario. [PR1532304](#)

Platform and Infrastructure

- The interface goes down and does not come up after configuring and deleting the Ethernet loopback configuration. [PR1353734](#)
- When you commit the **request system reboot** statement, routers gets stuck in the **Rebooting** state occasionally. [PR1385970](#)
- On the QFX10000 line of switches, the analyzer does not mirror after adding the child member to an aggregated Ethernet interface. [PR1417694](#)
- Issues with TCAM calculation on the codes 18 and 19 are observed. The same filter works without any issues on the code 17. [PR1469515](#)
- To reach tunnel out-going route destination through the BGP-Over-BGP route, recursive resolution is not supported. [PR1498085](#)
- Changing the scaled firewall profiles on the fly does not release the TCAM resources as expected. [PR1512242](#)
- Output interface index in the sFLOW packet is zero when transit traffic is observed on the IRB interface with VRRP enabled. [PR1521732](#)

Routing Protocols

- Node protection for the RSVP LSP on FTI interfaces does not work. [PR1456350](#)
- On QFX5000 , the PIP decapsulation in the forward filter does not work when the **NO from matching conditions AND when subnet masks < 32** statement is used. [PR1511893](#)
- Label based next hop load balancing does not occur at ingress node in case of the single hop LSPs. [PR1516170](#)
- On the QFX5100 and QFX5200 line of switches, traffic gets load balanced based on the final list of the next hops programmed in the hardware, for the route if the fti logical interface is used as next hop. [PR1517519](#)
- On the QFX5100 line of switches, if the you do not run the QFX-5e codes (non-TVP architecture) when an image with the Broadcom SDK upgrade (6.5.x) is installed, the CPU utilization might go up by around 5 percent. [PR1534234](#)

SEE ALSO

[What's New | 240](#)

[What's Changed | 253](#)

[Open Issues | 261](#)

[Resolved Issues | 265](#)

[Documentation Updates | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)

Open Issues

IN THIS SECTION

- [EVPN | 261](#)
- [High Availability \(HA\) and Resiliency | 262](#)
- [Infrastructure | 262](#)
- [Layer 2 Features | 262](#)
- [Platform and Infrastructure | 262](#)
- [Routing Protocols | 264](#)
- [User Interface and Configuration | 265](#)

Learn about open issues in Junos OS Release 20.3R3 for QFX Series Switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

EVPN

- The dev-longevity l2ald process generates the core file at `l2ald_next_bd_member`. [PR1570757](#)
- The MAC address of the end-host gets incorrectly programmed in the forwarding table after the ESI failovers. [PR1584595](#)

High Availability (HA) and Resiliency

- On the QFX5200-32C line of switches, the reboot time is degraded from 205 seconds in Junos OS Release 20.2R1 to 260 seconds in Junos OS Release 20.3. [PR1511607](#)

Infrastructure

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

Layer 2 Features

- On the QFX5000 Virtual Chassis, multicast traffic gets flooded even when the IGMP report times out. [PR1431893](#)
- On the QFX5110 and QFX5120 line of switches, changing lo0 IP address might sometimes either result in stale entry of IP in mpls_entry table or missing IP entry, which results in traffic drop for VXLAN traffic. [PR1472333](#)
- On the QFX5100 line of switches, traffic might be dropped in the Packet Forwarding Engine after making changes related to TPID in the dcd process. [PR1477156](#)
- In the QFX5100 line of switches, Q-in-Q with IRB do not work. [PR1481648](#)
- On the QFX5000 line of switches with storm control, significant difference between the configured rate and actual rate is observed. [PR1526906](#)

Platform and Infrastructure

- On the QFX5100-48T-6Q line of switches, the port LEDs might not work. [PR1317750](#)
- On the QFX10000 line of switches, the source MAC and TTL values do not get updated for the routed multicast packets in EVPN-VXLAN. [PR1346894](#)
- On the QFX5000 line of switches, firewall filter does not get programmed after you delete a large filter and add a new one in a single commit. [PR1583440](#)
- The EVPN VXLAN CE interface with RSTP configured might cause LACP or BFD issues. [PR1572504](#)
- On the QFX10008 line of switches, the ksyncd process generates core file at `../../../../src/junos/usr.sbin/ksyncd/ksyncd_msg.c` when rebooting the box with the EVPN-VXLAN configurations. [PR1574594](#)
- On the QFX10000 line of switches, the port might not get brought down immediately during some abnormal type of line card reboot. [PR1577315](#)
- The offer message from the server reaching the relay agent occurs. [PR1530160](#)

- The MAC addresses might not be relearned successfully after MAC address age timeouts. [PR1567723](#)
- The backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- On the QFX10000 line of switches, the Aruba wireless access point (AP) heartbeat packets get dropped. As a result, the Aruba wireless AP cannot work. [PR1352805](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- Due to the transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Those reported events had been disabled to prevent alarms and possibly unnecessary hardware replacements. [PR1384435](#)
- Either unicast RPF in the **Strict** mode or ICMP redirect does not work. [PR1417546](#)
- When the spine underlay is tagged and untagged, the inner packet comes and goes over the type-2 tunnel resulting in IPv4 discarding traffic silently on PECHIP. [PR1435864](#)
- On the QFX5200 line of switches, the ISSU might fail. [PR1438690](#)
- On the QFX10000 line of switches, removal of the EVPN-VXLAN Layer 3 gateway on the IRB interface from the spine switches might cause traffic to be silently discarded. [PR1446291](#)
- After renaming VLAN on the trunk interface, the local host MAC learning gets halt for more than 30 seconds. [PR1454274](#)
- **storm-control** does not rate-limit ARP packets. [PR1461958](#)
- On the QFX5110 line of switches, the VXLAN VNI (mcast) scaling causes traffic issue. [PR1462548](#)
- On the QFX10002-60C line of switches, the Packet Forwarding Engine installation or deletion, and link flap convergence time are reduced in Junos OS Release 19.4 compared to Junos OS Releases 19.3R1 and 19.2R1. [PR1464572](#)
- On the QFX10000 line of switches, the loopback-based filter with decapsulation GRE does not work as expected. [PR1479613](#)
- Disabled interfaces might still transmit power after the device reboots. [PR1487554](#)
- On the QFX10000 line of switches, the **show pfe filter hw filter-name** command does not retrieve the Packet Forwarding Engine program. [PR1495712](#)
- On the QFX5210 line of switches, port LED lights do not behave as expected after the upgrade. [PR1498175](#)
- Unable to scale 96000 ARP/ND with **DDoS-protection global disable-fpc**. [PR1507355](#)
- On the QFX5100 line of switches, degradation is observed during the system reboot time and FPC online time. [PR1513540](#)
- On the QFX10002-60C line of switches, degradation during system reboot time is observed. [PR1516086](#)
- SNMP trap of power failure might not be sent out. [PR1520144](#)

- The phc daemon might crash while committing the phone-home client configuration. [PR1522862](#)
- The ONIE testing fails for AS7816-64x with Junos OS Release 20.3R1.8 image. [PR1542827](#)
- On the QFX5120 line of switches, the MPLS traffic might not be forwarded after the aggregate Ethernet interface flaps. [PR1589840](#)
- On the QFX10000 line of switches, need to move WRL7 SDK to RCPL31. [PR1547565](#)
- In Junos OS Release 20.2, some features show up as a licensed feature. While using the features, alarms and commit warnings are displayed. However, there are no functional impact. [PR1558017](#)
- Unexpected multicast traffic streams is observed after enabling EVPN. [PR1570689](#)
- On the QFX1000 line of switches with P node, enabling link-protection on a P2MP LSP causes traffic drop. [PR1589488](#)
- BUM traffic are looped for a short time when the interface is newly added as the CE interface. [PR1493650](#)
- On the QFX10002-60C line of switches, NTP packets from the **set ntp date** command using outgoing interface IP instead of configured source-address. [PR1545022](#)
- On the QFX5100 line of switches that does not run the QFX-5E codes (non TVP architecture), the CPU utilization might go up by around five percent when image with Broadcom SDK upgrade (6.5.X) is installed. [PR1534234](#)
- Verification of the filter counter statistics fails when the received packets gets doubled. [PR1590009](#)
- Packet drop might occur in the ECMP next-hop flap scenario. [PR1594030](#)
- The re-installation of the Type-5 tunnels might fail in the EVPN-VXLAN scenario. [PR1595197](#)
- The mcsnoopd process might crash when you delete or add the Layer 2 forwarding configuration after ISSU. [PR1596483](#)
- On the QFX5110, QFX5120, QFX5200, and QFX5210 line of switches, the DDoS violations might be reported for the IP multicast miss traffic (IPMCAST-MISS) incorrectly. [PR1598678](#)
- During image upgrade, sometimes system goes for NMI or auto vmcore generation. [PR1601867](#)

Routing Protocols

- On the QFX5100 Virtual Chassis, the instability issues due to disabling DDoS-protection is observed. [PR1238875](#)
- On the QFX-5100 Virtual Chassis or Virtual Chassis Fan, the following error is observed in the hardware with the mini-PDT base configurations: **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:l3 nh 6594 unintsall failed.** [PR1407175](#)

- Traffic destined to first level IPv6 route, which is the tunnel destination (while the IPv6 route itself is not resolved over tunnel) is silently discarded. [PR1510053](#)
- The multi-hop BFD session might flap if you execute the RSI (Request Support Information) collection command. [PR1589765](#)

User Interface and Configuration

- The **file copy**' command uses a default staging-directory and this staging directory is located at **/var/tmp** for the root user and **/var/home/** for all non-root users. When you use **file copy**' command, in case if the storage system is full, you can use the **staging-directory** command line option to choose a different staging directory (other than the default). [PR1494489](#)

SEE ALSO

[What's New | 240](#)

[What's Changed | 253](#)

[Known Limitations | 259](#)

[Resolved Issues | 265](#)

[Documentation Updates | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 20.3R3 | 266](#)
- [Resolved Issues: Release 20.3R2 | 272](#)
- [Resolved Issues: Release 20.3R1 | 276](#)

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

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

Resolved Issues: Release 20.3R3

Class of Service (CoS)

- The dscp classifier does not work and all packets are sent to a single queue. [PR1585361](#)
- Unable to configure policer with bandwidth-limit greater than 50G. [PR1575049](#)
- The Buffer allocation for the VCP ports might not get released in the Packet Forwarding Engine after physically moving the port location. [PR1581187](#)

EVPN

- **global-mac-ip-table-aging-time** changes from a high to low value might not take effect. [PR1562925](#)
- Policy with mac-filter-list might not work if you commit any change unrelated to the policy in the EVPN scenario. [PR1567623](#)
- After the device reboots in the EVPN-VXLAN setup with the graceful restart, EVPN routes do not get advertised to the EVPN peers until the rpd process gets up for 180 seconds. [PR1586246](#)

Forwarding and Sampling

- The **archive transfer-on-commit** configuration fails when running Junos OS Release 18.2R3-S6.5. [PR1563641](#)

Interfaces and Chassis

- MAC address entry issue might be observed after the MC-LAG interface. [PR1562535](#)
- Removal of the configuration from the interface stanza might cause the dcpfe process to crash. [PR1594356](#)
- Newly added MC-LAGs do not come up after the Routing Engine switchovers. [PR1583547](#)

Layer 2 Ethernet Services

- DHCP packet drop might occur when you configure the DHCP relay on a leaf device. [PR1554992](#)
- The DHCP client becomes be offline for 120 seconds after sending the **DHCPINFORM** message in the DHCP relay scenario. [PR1575740](#)
- DHCP relay drops packets during the renewal DHCP process. [PR1576417](#)

Layer 2 Features

- On the QFX5120 line of switches, packets with VLAN ID 0 are dropped. [PR1566850](#)
- On the QFX5000 line of switches, software-forwarded VXLAN de-encapsulated packets have illegal length. [PR1574435](#)
- On the QFX5110-32Q line of switches, LACP does not come up in the **Non-Oversubscribed** mode for a set of ports. [PR1563171](#)

- MAC addresses learnt from the MC-LAG client device might keep flapping between the logical child interface and MC-AE interface after you disable one child link in the MC-AE interface. [PR1582473](#)
- In the OVSDB VXLAN scenario, inner VLAN tag 8 gets added unexpectedly into the encapsulated Ethernet header. [PR1531319](#)
- Traffic drop might occur on the aggregated Ethernet interface. [PR1585320](#)
- On the QFX5000 line of switches, DF might not forward BUM traffic. [PR1575976](#)
- Traffic forwarding for VLAN 2 might not be correct when you remove a VLAN member from the ESI interface. [PR1570446](#)
- The dcpfe process crashes in the VXLAN scenario. [PR1571170](#)
- Packets received on a port in the **LACP Detached** and **Broadcom STP Blocked** states might get forwarded. [PR1553570](#)
- LACP gets into the **Detached** state when you delete VLAN on the configured aggregate Ethernet interface on the SP style. [PR1555862](#)

Network Management and Monitoring

- Slow memory leak might occur for the snmpd process. [PR1575790](#)

Platform and Infrastructure

- On the QFX5000 line of switches, the egress ACL filter entries is only 512 in Junos OS Release 19.4R1. [PR1472206](#)
- The OSPF neighborship gets stuck in the **Start** state after configuring the EVPN-VXLAN. [PR1519244](#)
- Some inter-VLAN traffic flows do not converge after rebooting a spine (QFX10002) device in an EVPN-VXLAN non-collapsed scaled scenario. [PR1522585](#)
- On the QFX10000 line of switches, the Denial of Service (DoS) upon receipt of DVMRP packets gets received on the multi-homing ESI in VXLAN. [PR1539194](#)
- Traffic loss might be observed on interfaces in a VXLAN environment. [PR1524955](#)
- On the QFX5100 Virtual Chassis and Virtual Chassis fan, after NSSU while performing GRES, backup can generate core file and go to the database prompt. [PR1533874](#)
- The dcpfe process might crash and cause FPC to restart due to the traffic burst. [PR1534340](#)
- The BFD neighborship fails with the EVPN_VXLAN configuration after the Layer 2 learning restarts. [PR1538600](#)
- On the QFX10000 line of switches, traffic might get dropped when the **set routing-options forwarding-table no-ecmp-fast-reroute** configuration is changed to 128 ECMP entries. [PR1547457](#)
- On the QFX10002 line of switches, traffic loss occurs after loading the EVPN configuration. [1590919](#)
- The 40GbE interface might be channelized after restarting the Virtual Chassis member. [PR1548267](#)

- The interface filter with source-port 0 matches everything instead of port 0. [PR1551305](#)
- The dcpfe process might crash when you enable auto-channelization. [PR1484336](#)
- Traffic might not passed due to the addition of the VLAN tag 2 while passing through the Virtual Chassis port. [PR1555835](#)
- Traffic might be dropped when a firewall filter rule uses the then VLAN action. [PR1556198](#)
- The MAC addresses learned in a Virtual Chassis might fail in the MAC scaling environment. [PR1558128](#)
- PRBS (Pseudo Random Binary Sequence) test on the QFX5200 line of switches fails for the 100GbE interfaces with the default settings. [PR1560086](#)
- Few IPv6 ARP ND fails after loading the base configurations. [PR1560161](#)
- When configuring the static MAC and static ARP on the EVPN core aggregate interface, the underlay next-hop programming might not be updated in the Packet Forwarding Engine. [PR1561084](#)
- The tunable optics SFP+-10G-T-DWDM-ZR does not work. [PR1561181](#)
- PTP BC with G.8275.2.enh profile_2 512 clients does not come up. [PR1561348](#)
- PTP lock status gets stuck at the **Acquiring** state instead of the **Phase Aligned** state. [PR1561372](#)
- On the QFX10000 line of switches, the dcpfe process might crash during the configuration changes. [PR1561746](#)
- Traffic loss might occur in a large-scaled EVPN scenario when the next-hop type changes between discard and unicast. [PR1562425](#)
- On the QFX5000 line of switches, port mirroring might not work as expected. [PR1562607](#)
- On the QFX5120 line of switches, storm control with IRB interface might not work correctly. [PR1564020](#)
- On the QFX5100 Virtual Chassis, the following continuous message is observed:
agentd-pfe-proxy_telemetry_publisher. [PR1566528](#)
- On the QFX5100 line of switches, the following internal comment is displayed: **Placeholder for QFX platform configuration**. [PR1567037](#)
- On the QFX10002 line of switches, discrepancy in inet.1 versus Packet Forwarding Engine reports multicast routes. [PR1567353](#)
- On the QFX10000 line of switches, the firewall log incorrectly gets populated from the Packet Forwarding Engine for the IPv6 traffic. [PR1569120](#)
- PTP management message with SMTLV is sent only to the first port number to go active in the member multicast-mode l2-ifl. [PR1571283](#)
- On the QFX10008 chassis, the dcpfe process generates core file. [PR1572889](#)
- On the QFX10000 line of switches, a high rate of 802.3X pause frames are sent out of the interfaces. [PR1575280](#)
- The dcpfe process crashes while checking the virtual tunnel-nh packet status. [PR1580114](#)

- When analyzers are mapped to the channelized port, mirroring might not happen properly. [PR1580473](#)
- Issue is observed in telemetry when the **set services analytics streaming-server <> <>** configuration is present and server is not reachable. [PR1581192](#)
- When the deleted aggregated Ethernet interface is not getting deleted (mirror trunk group) in the hardware for the analyzer input aggregated Ethernet interface. [PR1589579](#)
- The interface might not come up with 1G optics. [PR1554098](#)
- In the QFX10002-72Q line of switches, SNMP walk jnxOperatingEntry displays only two PSU even if four PSU are installed. [PR1555852](#)
- Unexpected packet loss might occur if you delete the subunit of the physical interface. [PR1571286](#)
- TACACS traffic might be dropped. [PR1578579](#)
- On the QFX10000 line of switches, the firewall filter logs are incorrectly populated the protocol 8847 entries. [PR1582780](#)
- FPC might crash in a scaled-firewall configuration. [PR1586817](#)
- On the QFX10000 line of switches, the dcpfe process might crash. [PR1546572](#)
- The traffic does not get be load-balanced properly in the EVPN overlay-ecmp setup. [PR1550020](#)
- The VGA might be down when you configure the IRB interface with multi VGA. addresses [PR1555338](#)
- You need the Layer 2 static license though the license is included in the base license. [PR1557631](#)
- Upon receipt of the specific sequences of genuine packets destined to the device, the kernel crashes and restarts. [PR1557881](#)
- The VCF might become not stable. [PR1559172](#)
- The subscriber management infrastructure daemon (smid) process might become nonresponsive at 100 percent. [PR1559402](#)
- On the QFX10000-60S-6Q line of switches, the line card takes more than 15 minutes to boot up after the panic or watchdog reboots. [PR1559725](#)
- The VXLAN queue, DDos violation, and RARP packets flood might occur if you receive the RARP packets more than the supported DDos bandwidth. [PR1560243](#)
- The dcpfe process might crash after committing the EVPN-VXLAN profile configuration and ARP resolution might fail causing traffic issues. [PR1561588](#)
- On the QFX5000 line of switches, the dcpfe process might crash after deleting the VXLAN configuration. [PR1562692](#)
- The received encapsulated VXLAN packets might get dropped if you load the EVPN-VXLAN configuration. [PR1562814](#)
- Sampled memory might leak when the analyzer is in the **Down** state. [PR1564790](#)
- Traffic loss might occur in the MC-LAG scenario. [PR1565287](#)

- The DF (Designated Forwarder) might not forward traffic. [PR1567752](#)
- The 100G port with module QSFP 100G-SR4-T2 converts to two channelized interfaces without any channelized configuration. [PR1567937](#)
- On the QFX10002-60C line of switches, another port shutdowns after shutting down one port. [PR1568294](#)
- The BFD session flaps between leaf and core during the Spine reboot causing other protocols to flap. [PR1568615](#)
- The dcpfe process might crash if the TYPE-5 tunnel fails to be installed for EVPN-VXLAN. [PR1570136](#)
- Traffic loss might occur due to faulty FPC. [PR1574779](#)
- Port-mirroring might not work when the analyzer output is a trunk interface. [PR1575129](#)
- The BFD session might flap occasionally during the Spine reboot. [PR1575296](#)
- On the QFX5000 line of switches, analyzer does not work. [PR1576327](#)
- Multicast Packets with TTL=1 gets dropped when you enable the VXLAN interface when igmp-snooping or MLD-snooping. [PR1576775](#)
- Traffic might not be load-balanced properly in an EVPN overlay-ecmp setup. [PR1582017](#)
- The DHCP offer packets might be dropped on the Spine device in the VXLAN multi-homing setup. [PR1585715](#)
- Inter and intra VNI traffic might drop in the Spine device with the EVPN-VXLAN CRB configuration. [PR1586537](#)
- The VXLAN DDoS violation might occur when you disable the port mirror analyzer output interface. [PR1590150](#)
- The Routing Engine kernel might crash due to the child logical interface of the aggregated interface leading to Junos kernel failure. [PR1592456](#)
- Multiple crashes with **toe_interrupt_errors** error message might occur. [PR1593025](#)
- The existing ECMP route traffic might be dropped if you configure a static ECMP route with the same number of next-hops as the existing ECMP route. [PR1594573](#)

Routing Policy and Firewall Filters

- The rpd process might crash when the deletion of routing table occurs. [PR1565629](#)

Routing Protocols

- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. [PR1514966](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- BGP LU session might flap with the AIGP used scenario. [PR1558102](#)

- On the QFX5110 line of switches, the ARP resolution might fail if **native-vlan-id** is configured on the VXLAN interface. [PR1563569](#)
- The dcpfe process might crash when the size of the Local Bias Filter Bitmap string exceeds 256 characters. [PR1568159](#)
- The GRE egress traffic might not be forwarded between the different routing-instances. [PR1573411](#)
- The QFX5000 line of switches might drop the DHCP packets in the static VXLAN scenario. [PR1576168](#)
- The rpd process might crash after committing with the configured static group 224.0.0.0. [PR1586631](#)
- The remaining BFD sessions of the aggregated Ethernet interface flaps continuously if one of the BFD sessions is deleted. [PR1516556](#)
- The dcpfe process might crash when any interface flaps. [PR1579736](#)
- The BFD sessions over IRB interface becomes nonresponsive in the **Init** state with the incremental FRR errors. [PR1541851](#)
- The BFD session on the Layer 3 sub-interface of the ESI aggregated Ethernet interface might flap when the upstream underlay or overlay BGP flaps. [PR1544982](#)
- The fpc and dcpfe processes might crash when you insert and remove the channelized interface of QSFP. [PR1547231](#)
- On the QFX5000 line of switches, the continuous traffic destined to a device configured with MC-LAG causes nodes to lose their control connection, which can impact traffic. [PR1552877](#)
- There might be traffic loss when the GRE interface flaps. [PR1566428](#)
- On the QFX5000 line of switches, memory leak might occur. [PR1566483](#)
- Memory leak might occur in the MSDP scenario. [PR1571906](#)
- On the QFX5000 line of switches, traffic loss might occur in the EVPN-VXLAN scenario. [PR1580005](#)
- The IS-IS packet might be corrupted on the provider edge device over the Layer 2 circuit tunnel. [PR1580047](#)
- The DHCP packets might be dropped if you apply the **dyn-dhcpv4_v6_trap** dynamic filter on the interface. [PR1580352](#)
- With IGMP snooping implemented, unexpected jitter issue that might cause traffic loss occurs. [PR1583207](#)

User Interface and Configuration

- During rare circumstances, the mgd process might crash and generate a core file on the Junos devices connected with the Contrail Service Orchestration (CSO). [PR1569903](#)

Resolved Issues: Release 20.3R2

EVPN

- ARP table might not be updated after VMotion or network loop is performed. [PR1521526](#)
- All the ARP reply packets toward some address are flooded across the entire fabric. [PR1535515](#)
- EVPN-VXLAN registers MAC-move counters under **system statistics bridge** even though there is no actual MAC-move for the multi-homed clients. [PR1538117](#)
- The l2ald process might generate core file when changing the EVPN-VXLAN configuration. [PR1541904](#)
- The l2ald daemon might crash when **forwarding-options evpn-vxlan shared-tunnels** is configured. [PR1548502](#)
- The l2ald process generates core file at **l2ald_iff_rtm_delete_subintf_ifbds** during dci fusion run. [PR1550109](#)

Forwarding and Sampling

- The l2ald process might crash due to next-hop issue in the EVPN-MPLS. [PR1548124](#)
- Configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. [PR1563641](#)

Interfaces and Chassis

- The logical interface might flap after the addition or deletion of the native VLAN configuration. [PR1539991](#)

Infrastructure

- The output of the **show interfaces extensive** command might display 0 temporarily during a race condition when SNMP query for JnxCos is also issued. [PR1533314](#)

Layer 2 Features

- Flow control is enabled in the Packet Forwarding Engine irrespective of the interface configuration and the fix causes a very small amount of packet loss when a parameter related to an interface such as **interface description** on any port is changed. [PR1496766](#)
- The MAC address in the hardware table might become out of synchronization between the primary device and member in the Virtual Chassis after the MAC flaps. [PR1521324](#)
- Check traffic with VXLAN encap header fails. [PR1541316](#)
- On the QFX5120 device, packets with VLAN ID 0 are dropped. [PR1566850](#)

Platform and Infrastructure

- On the QFX5000 line of switches, the egress ACL filter entries is only 512 in Junos OS Release 19.4R1. [PR1472206](#)
- On the QFX10000 device, the chassisd process might generate core files on the backup Routing Engine after commit for 200 seconds due to the following error message: **CHASSISD_MAIN_THREAD_STALLED**. [PR1481143](#)
- The SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. [PR1484322](#)
- IRB MAC is not programmed in hardware when the MAC persistence timer expires. [PR1484440](#)
- The archival function might fail in certain conditions. [PR1507044](#)
- The output VLAN push might not work. [PR1510629](#)
- On the QFX5000 line of switches, multicast traffic loss is observed due to few multicast routes missing in the spine node. [PR1510794](#)
- DHCP traffic might not be forwarded correctly when sending the DHCP unicast packets [PR1512175](#)
- Channelized interfaces might fail to come up. [PR1512203](#)
- The output of the **show chassis forwarding-options** command displays incorrect display issue, Virtual Chassis environment, and configured num-65-127-prefix values. [PR1512712](#)
- The 100-Gigabit Ethernet AOC non-breakout port might be auto-channelized to a different speed. [PR1515487](#)
- On the QFX5100 device, the cprod process timeout triggers high CPU utilization. [PR1520956](#)
- The output interface index in the sFLOW packet is zero when the transit traffic is observed on the IRB interface with VRRP enabled. [PR1521732](#)
- Some inter-VLAN traffic flows do not converge after rebooting a spine (QFX10002) device in an EVPN-VXLAN non-collapsed scaled scenario. [PR1522585](#)
- On the QFX10002, QFX10008, and QFX10016 line of switches, the following error message is observed during specific steps while clearing and loading the scaled configuration again:
PRDS_SLU_SAL:jprds_sl_u_sal_update_lncnt(),1379:jprds_sl_u_sal_update_lncnt call failed. [PR1522852](#)
- Packet loss is observed while validating the policer after restarting the chassis control. [PR1531095](#)
- High rate of ARP or NS packets might be observed between a device that runs Junos OS and host when the device that runs Junos OS receives an ARP or NS packet on an interface in transition. [PR1534796](#)
- The following Packet Forwarding Engine error message is observed in the **BRCM-VIRTUAL,brcm_virtual_tunnel_port_create(),489: Failed NW vxlan port token(45) hw-id(7026) status(Entry not found)**. [PR1535555](#)
- Software recovery or installation using the **Bootable USB Flash Drive** option might fail. [PR1536799](#)

- The filter instance does not get removed from the Packet Forwarding Engine after deactivating VLAN and IRB. [PR1537108](#)
- On the QFX5100-48T, interfaces are not created after 10g channel-speed is applied across the 48 to 53 ports. [PR1538340](#)
- The **Management Ethernet link down** alarm is seen while verifying system alarms in a Virtual Chassis setup. [PR1538674](#)
- Unable to take RSI properly due to the authentication error. [PR1539654](#)
- Traffic loss might be seen in the OVSDb VXLAN scenario. [PR1540208](#)
- Inter VLAN traffic drop might be observed in an EVPN-VXLAN scenario. [PR1541406](#)
- On the QFX5100-Virtual Chassis, the **End Segment Not Present** message is not reported for the ping overlay function with the local host MAC. [PR1542226](#)
- On the QFX5000 device running EVPN-VXLAN, the Packet Forwarding Engine error message might be seen: **bd_platform_irb_ifl_attach_detach: platform specific irb ifl attach/detach failed (-1)**. [PR1543812](#)
- On the QFX10002-60C device, the **show pfe filter** command is unavailable. [PR1545019](#)
- On QFX10000 device, traffic might get dropped when the **set routing-options forwarding-table no-ecmp-fast-reroute** configuration is changed to 128 ECMP entries. [PR1547457](#)
- On the QFX5100 Virtual Chassis, the backup Routing Engines clear the reporting alarm for a PEM failure intermittently for a missing power source. [PR1548079](#)
- The VXLAN encapsulated packet might be sent on the network port with an incorrect inner VLAN id 4095. [PR1548218](#)
- The 40G interface might be channelized after restarting the Virtual Chassis member. [PR1548267](#)
- The Neighbor Solicitation might be dropped from the peer device. [PR1550632](#)
- On the QFX5110 and QFX5120 devices, the DHCPv6 traffic received over vtep might not be forwarded. [PR1551710](#)
- On the QFX5000 devices, the ARP resolution might fail. [PR1552671](#)
- Traffic might be dropped when a firewall filter rule uses the **then VLAN** action. [PR1556198](#)
- On the QFX5120-48YM device, the **Multiple License Warning Messages** are observed. [PR1556816](#)
- On the QFX5120 device, amber LEDs are displayed for the fan modules after upgrading to Junos OS Release 20.2R1. [PR1558407](#)
- Few IPv6 ARP ND fails after loading the base configurations. [PR1560161](#)
- PTP BC with **G.8275.2.enh profile_2** 512 clients do not come up. [PR1561348](#)
- PTP lock status gets stuck at the **Acquiring** state instead of the **Phase Aligned** state. [PR1561372](#)
- The interface filter with source-port 0 matches everything instead of port 0. [PR1551305](#)

- Traffic might not be passed due to the addition of the VLAN tag 2 while passing through the Virtual Chassis port. [PR1555835](#)
- Packet drops might be seen while verifying the LFM operation during the graceful switchover. [PR1515280](#)
- The l2cpd process might crash if the ERP is deleted after the switchover. [PR1517458](#)
- Traffic loss might be observed on interfaces in a VXLAN environment. [PR1524955](#)
- On the QFX10000 devices, channelizing the 40G port to 10G port might bring down another interface. [PR1527814](#)
- On the QFX10002 devices, the firewall log incorrectly gets populated from the Packet Forwarding Engine. [PR1533814](#)
- The ARP request might be dropped in leaf in the EVPN-VXLAN scenario. [PR1539278](#)
- The Broadcom chip FPC might crash during the system booting. [PR1545455](#)
- The **action-shutdown** statement of the storm control does not work for the ARP broadcast packets. [PR1552815](#)
- Traffic storm might be caused by analyzer due to link flap. [PR1557274](#)
- On the QFX5000 device, the firewall filter might fail to work. [PR1558320](#)
- On the QFX10000 device, the dcpfe process might crash during configuration changes. [PR1561746](#)
- Traffic loss might occur in a large-scaled EVPN scenario when the next hop Type changes between discard and unicast. [PR1562425](#)
- On the QFX5000 devices, port mirroring might not work as expected. [PR1562607](#)
- The RPD memory might leak on the backup Routing Engine due to the flapping of the link. [PR1539601](#)

Routing Policy and Firewall Filters

- The policy configuration might be mismatched between the rpd and mgd process when **policy-options prefix-list** is deactivated and is involved in the configuration sequence. [PR1523891](#)

Routing Protocols

- On the QFX 5100-48T-6Q Virtual Chassis or Virtual Chassis fan, the following error message is observed while copying the image to the Virtual Chassis fan member and trying to downgrade the image: **rcp for member 14, failed**. [PR1486632](#)
- The OSPF neighborhood gets stuck in the **Start** state when EVPN-VXLAN is configured. [PR1519244](#)
- On the QFX5110-32Q device, the following syslog error message is observed after loading the NC T5 EVPN VXLAN configuration: **BCM-L2,pfe_bcm_l2_sp_bridge_port_tpid_set() Config TPID New/Old (8100:8100) Other-Tpid's ba49, 4aa0, 80f**. [PR1558189](#)
- The dcpfe process might crash when the size of the Local Bias Filter Bitmap string exceeds 256 characters. [PR1568159](#)

- The dcpfe process might crash while updating VRF for multicast routes during IRB uninit. [PR1546745](#)
- The OSPFv3 session might keep flapping and OSPFv3 hellos might be dropped in the host-path. [PR1547032](#)
- The BGP LU session flap might be seen with the AIGP used scenario. [PR1558102](#)
- On the QFX5110 device, the ARP resolution might fail if **native-vlan-id** is configured on the VXLAN interface. [PR1563569](#)

User Interface and Configuration

- The configuration under groups stanza is not inherited properly. [PR1529989](#)

Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. [PR1490552](#)
- On the QFX5120 and QFX5210 devices, unexpected storm control events might occur. [PR1519893](#)

Resolved Issues: Release 20.3R1

Class of Service (CoS)

- On QFX5120 switches, the priority-based flow control (PFC) feature is not supported on 2-member Virtual Chassis currently because of the hardware limitation. [PR1431895](#)
- Traffic might be forwarded to an incorrect queue when fixed classifier is used. [PR1510365](#)

EVPN

- On QFX10002-60C EVPN/VXLAN multicast, the **show** command issued for the VTEP interface does not show the mesh-group ID. [PR1498052](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)
- Unable to create a new VTEP interface. [PR1520078](#)

Infrastructure

- The OID ifOutDiscards reports zero and sometimes shows valid value. [PR1522561](#)

Interfaces and Chassis

- Traffic over MC-LAG drops because the next-hop points ICL link instead of MC-LAG. [PR1486919](#)
- MC-LAG consistency check fails if multiple IRB units are configured with the same VRRP group. [PR1488681](#)
- Error message is not generated while verifying the GRE limitation. [PR1495543](#)

Layer 2 Features

- MAC learning might not work correctly on QFX5120 switches. [PR1441186](#)
- On QFX5120 switches Q-in-Q, the third VLAN tag is not pushed onto the stack and SWAP is being done instead. [PR1469149](#)
- On QFX5200 switches, MAC learning rate is degraded by 88 percent. [PR1494072](#)
- Traffic imbalance might be observed on QFX5000 switches if **ash-params** is not configured. [PR1514793](#)
- MAC address in hardware table might become out of sync between master and member in Virtual Chassis after MAC flap. [PR1521324](#)

Layer 2 Ethernet Services

- Issues with DHCPv6 relay processing confirm and reply packets. [PR1496220](#)
- The MC-LAG might become down after disabling and then enabling the force-up. [PR1500758](#)
- The aggregated Ethernet interface might not come up after switch is rebooted. [PR1505523](#)

MPLS

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

Platform and Infrastructure

- Port LEDs do not work on the QFX5100 switch in a QFX5110-QFX5100 mixed mode Virtual Chassis. [PR1317750](#)
- A VM core is seen on QFX Series Virtual Chassis. [PR1421250](#)
- SFP-LX10 stays down until autonegotiation is disabled. [PR1423201](#)
- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB. [PR1442587](#)
- In the EVPN-VXLAN scenario, changing the VLAN name associated with the access ports might prevent the MAC addresses from being learned. [PR1454095](#)
- On the QFX5100 switch, the interface output counter is double counted for self-generated traffic. [PR1462748](#)
- On the QFX5100 switch, traffic loss might be seen with framing errors or runts if MACsec is configured. [PR1469663](#)
- On the QFX5000 switch, the DSCP marking might not work as expected if the fixed classifiers are applied to interfaces. [PR1472771](#)
- The sFlow could not work correctly if the received traffic goes out of more than one interface. [PR1475082](#)
- The dcpfe process might generate core file with the non-oversubscribed mode after SDK upgrade. [PR1485854](#)
- The 10 GbE VCP ports do not become active in a QFX5100 Virtual Chassis scenario. [PR1486002](#)

- On QFX5100 switches, If more than one UDF filter or term is configured, then only the first filter or term will be programmed in the hardware because of SDK 6.5.16 upgrade. [PR1487679](#)
- The queue statistics are not as expected after configuring the IFD and logical-interface shaping with the transmit rate and scheduler map [PR1488935](#)
- High CPU load due to receipt of specific multicast packets on Layer 2 interface. [PR1491905](#)
- Traffic loss could be observed in mixed Virtual Chassis setup of QFX5100 and EX4300 switches. [PR1493258](#)
- Traceroute monitor with MTR version v.69 shows a false 10 percent loss. [PR1493824](#)
- On the QFX5120 switch, traffic loss might be seen in a MC-LAG scenario. [PR1494507](#)
- On the QFX5120 switch, SNMP polling for CPU utilization and CPU state of backup Routing Engine do not show in a two-member Virtual Chassis. [PR1495384](#)
- Kernel routing table queue become nonresponsive after J-Flow sampling of a malformed packet. [PR1495788](#)
- ARP does not get refreshed after timeout on QFX10002-60C. [PR1497209](#)
- Extra carrier transitions are seen on the peer when negative triggers are performed on QFX5100 and QFX5110 switches. [PR1497380](#)
- Virtual Chassis is not stable with 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces. [PR1497563](#)
- 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 an SFP of the aggregated Ethernet member interface is unplugged or plugged. [PR1497993](#)
- The **request-pfe-execute** command takes longer than 5 seconds to get a reply in on the QFX5100 platform. [PR1498092](#)
- Firewall filter might not get applied on QFX5100 and QFX5110 switches. [PR1499647](#)
- BFD sessions flap after deactivating or activating the aggregated Ethernet interface or executing GRES. [PR1500798](#)
- On QFX5000 switches, ERPS might not work correctly. [PR1500825](#)
- The interface becomes physically down after changing to FEC none mode. [PR1502959](#)
- LLDP packets are not acquired when **native-vlan-id** and tagged VLAN-ID are the same on a port. [PR1504354](#)
- The l2cpd might crash if the ERP configuration is added or removed, and l2cpd is restarted. [PR1505710](#)
- The archival function might fail in certain conditions. [PR1507044](#)
- On QFX5100 switches, the fxpc process might crash while installing image through ZTP. [PR1508611](#)

- Traffic might be affected on QFX10002, QFX10008, and QFX10016 platforms because of PECHIP wedge caused by deactivating CoS ETS configuration. [PR1509220](#)
- ARP replies might be flooded through the EVPN-VXLAN network as unknown unicast ARP reply. [PR1510329](#)
- The QFX10000-36Q line card used on QFX10008 and QFX10016 switches might fail to detect any QSFP. [PR1511155](#)
- In VXLAN configuration, the firewall filters might not be loaded into the TCAM with the message **DFWE ERROR DFW: Cannot program filter ..** because of the TCAM overflow after upgrading to Junos OS Release 18.1R3-S1, 18.2R1 and later. [PR1514710](#)
- The routes update might fail upon HMC memory issue and affects the traffic. [PR1515092](#)
- The MAC learning might not work properly after multiple MTU changes on the access port in a VXLAN scenario. [PR1516653](#)
- The dcpfe process might crash because of memory leak. [PR1517030](#)
- The VGD core file might be generated when the OVSDDB server restarts. [PR1518807](#)
- Traffic forwarding might be affected when adding or removing or modifying the VLAN and VNI configurations such as VLAN-ID, VNI-ID and ingress-replication statement. [PR1519019](#)
- On QFX10002, QFX10008, and QFX10016 switches, **PRDS_SLU_SAL:jprds_slu_sal_update_lrcnt(),1379: jprds_slu_sal_update_lrcnt call failed** syslog error messages might be seen when clearing and loading the scaled configuration again. [PR1522852](#)
- On QFX10002-60C switches, sFlow adaptive-sampling with the rate limiter statement enabled crosses sample rate 65535. [PR1525589](#)
- Packet loss is seen while validating the policer after restarting chassis control. [PR1531095](#)

Routing Protocols

- The FPC process goes to the **NotPrsnt** state after upgrading the QFX5100 Virtual Chassis/Virtual Chassis Fabric. [PR1485612](#)
- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- The rpd process generates core file at **rt_nh_resolve_add_gen** in `../../../../src/junos/usr/sbin/rpd/lib/rt/rt_resolve_ind.c`: with the EVPN-DHCP configurations. [PR1494005](#)
- On the QFX5000 platform, high CPU load because of receipt of specific Layer 2 frames in an EVPN-VXLAN deployment and when deployed in a Virtual Chassis configuration. [PR1495890](#)
- Firewall filter might not work in certain conditions in a Virtual Chassis. [PR1497133](#)
- Traffic drop might be observed after modifying the FBF firewall filter. [PR1499918](#)

- With the **egress-to-ingress** configuration statement, you cannot configure 2000 scale and the scale is reduced to 1000. [PR1514570](#)
- Enabling IPv6 flow-based Packet Forwarding Engine hashing gives commit error. [PR1519018](#)
- Firewall **sample** configuration gives the warning as unsupported on QFX10002-36Q switches and does not work. [PR1521763](#)
- On QFX5000, the fxpc process might crash if VXLAN interface flaps. [PR1528490](#)

User Interface and Configuration

- The version information under the configuration is changed starting in Junos OS Release 19.1. [PR1457602](#)

SEE ALSO

What's New 240
What's Changed 253
Known Limitations 259
Open Issues 261
Documentation Updates 280
Migration, Upgrade, and Downgrade Instructions 281

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for the QFX Series Switches.

SEE ALSO

What's New 240
What's Changed 253
Known Limitations 259
Open Issues 261
Resolved Issues 265
Migration, Upgrade, and Downgrade Instructions 281

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 281
- Installing the Software on QFX10002-60C Switches | 284
- Installing the Software on QFX10002 Switches | 284
- 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 | 285
- Installing the Software on QFX10008 and QFX10016 Switches | 287
- Performing a Unified ISSU | 291
- Preparing the Switch for Software Installation | 292
- Upgrading the Software Using Unified ISSU | 292
- Upgrade and Downgrade Support Policy for Junos OS Releases | 294

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 **20.3** in the Release pull-down list to the right of the Software tab on the Download Software page.

4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.3 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Download the software to a local host.

8. Copy the software to the device or to your internal software distribution site.

9. Install the new jinstall package on the device.

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

Customers in the United States and Canada use the following command:


```
user@host> request system software add
source/jinstall-host-qfx-5-x86-64-20.3-R1.n-secure-signed.tgz reboot
```

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

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

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 20.3 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.

NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add <pathname><source>** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-20.3R3.9.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-x86-64-20.3R3.9.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-x86-64-20.3R3.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-x86-64-20.3R3.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-x86-64-20.3R3.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-x86-64-20.3R3.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 292](#)
- [Upgrading the Software Using Unified ISSU on page 292](#)

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-x86-64-20.3R3.n-secure-signed.tgz*.

NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
```

```

ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

Upgrade and Downgrade Support Policy for Junos OS Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- • Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and 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 | 240](#)

[What's Changed | 253](#)

[Known Limitations | 259](#)

[Open Issues | 261](#)

[Resolved Issues | 265](#)

[Documentation Updates | 280](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 296](#)
- [What's Changed | 306](#)
- [Known Limitations | 311](#)
- [Open Issues | 313](#)
- [Resolved Issues | 315](#)
- [Documentation Updates | 324](#)
- [Migration, Upgrade, and Downgrade Instructions | 324](#)

These release notes accompany Junos OS Release 20.3R3 for the SRX Series Services Gateways. 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 20.3R3 | 297](#)
- [What's New in Release 20.3R2 | 297](#)
- [What's New in Release 20.3R1 | 297](#)

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

What's New in Release 20.3R3

There are no new features for SRX in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features for SRX in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Application Security

- **Listing of micro-applications and non-configurable applications (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.3R1, we've introduced the following operational commands to display applications details:

- **show services application-identification application micro-applications** to display the list of micro-applications.
- **show services application-identification application non-configurable** to display the list of non-configurable applications.

[See [show services application-identification application micro-applications](#) and [show services application-identification application non-configurable](#).]

- **Application signature package rollback (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.3R1, you can roll back the current version of the application signature package to the previous version by using one of the following methods:

- **Automatic**—The system automatically rolls back to the previous version of the application signature package when the signature package installation fails on your security device.
- **Manual**—You can roll back the application signature package to its previous version on your security device using the **request services application-identification rollback** command.

[See [Predefined Application Signatures for Application Identification](#).]

Authentication and Access Control

- **Enhanced user identity information loading rate (SRX Series)**— Starting in Junos OS Release 20.3R1, for SRX300 Series devices with eUSB (SRX300, SRX320, SRX340, and SRX345), the authentication entry database moves from disk memory to internal memory. This enhancement reduces disk usage and increases the read-write speed of loading authentication entries.

For SRX1500, SRX380, SRX300, SRX320, SRX340, SRX345, SRX4100, SRX4200, SRX4600, SRX550HM, SRX5400, SRX5600, SRX5800 devices and vSRX 3.0 instances, the user firewall database operations on disk are enhanced; this results in reduced disk usage and increases disk lifetime.

[See [Active Directory Authentication Tables](#).]

Chassis Clustering

- **Wi-Fi Mini-Physical Interface Module (Mini-PIM) (SRX320, SRX340, SRX345, SRX380, and SRX550M)**—Starting in Junos OS Release 20.3R1, we provide support for the Wi-Fi Mini-PIM in High Availability (HA) cluster configuration.

[See [Wi-Fi Mini-Physical Interface Module Overview](#).]

- **Support for single PSU operation without alarms (SRX4100 and SRX4200)**—Starting in Junos OS Release 20.3R1, a new argument **pem-absence** is available at the **[edit chassis alarm]** hierarchy level. You can use **[set chassis alarm pem-absence ignore]** to ignore the power supply unit (PSU) alarm. By default, the PSU alarm is raised when any PSU is missing or not energized.

[See [Understanding Chassis Alarms](#), [show chassis alarms](#), and [pem-absence](#).]

Flow-Based and Packet-Based Processing

- **SPU Forwarding in PowerMode IPsec (SRX5400, SRX5600, and SRX5800 Devices)**—Starting in Junos OS Release 20.3R1, you can implement PowerMode IPsec (PMI) SPU forwarding on both encryption and decryption data paths. The PMI SPU supports the following features:
 - Encrypt the clear-text packets in the PMI data path on a different SPU.
 - Forward the decrypted IPsec packets to a clear-text session in the PMI data path.
 - Fat-tunnel mode and NAT-T

[See [Fragmentation Packets with PowerMode IPsec](#), [Route-Based and Policy-Based VPNs with NAT-T](#)]

Installation and Upgrade

- **Support for enhanced file-signing with veriexec (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550)**—Starting in Junos OS Release 20.3R1, we've enhanced the secure BIOS functionality to support Verified Exec, also known as veriexec, to validate the Junos OS software image. Veriexec is a file-signing and verification scheme that protects the Junos OS from unauthorized software and activity that might compromise the integrity of your device.

[See [Veriexec overview](#).]

Interfaces and Chassis

- **Support for Ethernet OAM LFM (SRX4100, SRX4200, and SRX4600)**—The IEEE 802.3ah standard defines OAM Link Fault Management (LFM). Ethernet LFM functions at the transport layer of OAM. You use Ethernet LFM to monitor link operations for physical or emulated point-to-point Ethernet links that connect peer OAM entities.

Starting in Junos OS Release 20.3R1, we support the following OAM LFM features:

- Discovery
- Link monitoring
- Fault signaling and detection
- Action profile

[See [Configuring Link Fault Management](#).]

Intrusion Detection and Prevention (IDP)

- **IDP support for pass-through GRE and IP-IP tunnel traffic in the TAP mode (SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 20.3R1, the Terminal Access Point (TAP) mode for IDP support is available for pass-through GRE and IP over IP (IP-IP) tunnel traffic. The TAP mode for IDP allows you to passively monitor traffic flows inside the IP-IP tunnel.

[See [TAP Mode for IDP](#).]

Junos Telemetry Interface

- **Packet Forwarding Engine and Routing Engine sensor support on JTI (SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)**—Junos OS Release 20.3R1 provides streaming support for revenue interface statistics through Packet Forwarding Engine (PFE) sensors and pseudo interface statistics through Routing Engine sensors. Sensors are supported through Junos telemetry interface (JTI) and remote procedure calls (gRPC) or gRPC Network Management Interface (gNMI) services. gNMI service is also enabled for other supported Routing Engine sensors.

Using JTI and gRPC or gNMI services, you can stream telemetry statistics to an outside collector.

These interface sensors are supported:

- Physical interfaces (IFD) (resource path **/interfaces/interface/**).
- Logical interfaces (IFL) (resource path **/interfaces/interface/subinterfaces/**).

These Routing Engine sensors are supported using gNMI services (previously, only gRPC services were supported):

- System events (resource path **/junos/events**).
- BGP peer information (resource path **/network-instances/network-instance/protocols/protocol/bgp/**).
- Memory utilization for routing protocol task (resource path **/junos/task-memory-information/**).
- Operational state of Routing Engines, power supply modules, Switch Fabric Boards, Control Boards, Switch Interface Boards, Modular Interface Cards, and Physical Interface Cards (resource path **/components/**).
- Link Layer Discovery Protocol (LLDP) (resource path **/lldp/**).
- Address Resolution Protocol (ARP) statistics for IPv4 routes (resource path **/arp-information/**).
- Network Discovery Protocol (NDP) table state information for IPv6 routes (resource path **/nd6-information/**).
- NDP router-advertisement statistics (resource path **/ipv6-ra/**).
- IS-IS routing protocol statistics (resource path **/network-instances/network-instance/protocols/protocol/isis/levels/level/** and **network-instances/network-instance/protocols/protocol/isis/interfaces/interface/levels/level/**).

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

Junos OS XML API and Scripting

- **Support for REST API over nondefault virtual routing and forwarding (VRF) instance (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can execute Junos OS operational commands using the REST API over a nondefault VRF instance. The nondefault VRF instance can be a user-defined instance or the management instance, `mgmt_junos`.

The REST API allows you to execute Junos OS operational commands over HTTP(S). If you don't specify a routing instance, REST API requests are sent over the default routing instance. Use a nondefault VRF instance to improve security and make it easier to troubleshoot.

Use the **routing-instance** *routing-instance* statement at the `[edit system services rest]` hierarchy level to specify a nondefault VRF instance for REST API requests.

[See [Management Interface in a Nondefault Instance](#) and [rest](#).]

J-Web

- **Remote access VPN (SRX Series)**—Starting in Junos OS Release 20.3R1, J-Web VPN supports remote access to allow users, who work at home or travel, to connect to the corporate office and its resources. Using J-Web, you can configure Juniper Secure Connect or NCP Exclusive Client remote access VPN. You can access these menus at VPN > IPsec VPN > Create VPN > Remote Access.

[See [Create a Remote Access VPN—Juniper Secure Connect](#) and [Create a Remote Access VPN—NCP Exclusive Client](#).]

- **AppQoS (SRX Series)**—Starting in Junos OS Release 20.3R1, J-Web supports application quality of service (AppQoS). Using AppQoS, you can prioritize and meter application traffic to provide better service for business-critical or high-priority application traffic. You can access this menu at Network > Connectivity > AppQoS.

[See [About the Application QoS Page](#).]

- **Enhanced Security Policies page (SRX Series)**—Starting in Junos OS Release 20.3R1, we've enhanced the Security Policies page at Security Policies & Objects > Security Policies for an improved user experience. You can edit the fields on the Security Policies page inline to create or edit a policy rule.

[See [About the Rules Page](#).]

- **Change in Configuration tab architecture (SRX Series)**—Starting in Junos OS Release 20.3R1, we've removed the existing Configuration tab. The menus under the Configuration tab are classified into the following new tabs for an enhanced user experience:
 - Device
 - Network
 - Security Rules & Objects

- Security Services
- VPN

The new tabs include the corresponding configuration menu and sub-menu options.

[See [Configure Basic Settings](#).]

- **Improved Access Profile page (SRX Series)**—Starting in Junos OS Release 20.3R1, we've enhanced the Access Profile page for an improved user experience. The Access Profile page now supports the newly added Local and RADIUS authentication services.

[See [About the Access Profile Page](#).]

Layer 2 Features

- **Support for different MAC addresses on IRB interfaces (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM)**—Starting in Junos OS Release 20.3R1, you can assign a different MAC address to an IRB interface.

To assign a MAC address to an IRB interface, enable the **mac** statement at the **[edit interfaces irb unit unit-number]** hierarchy level.

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

Logical Systems and Tenant Systems

- **Support for root system's stream configuration for user logical systems and tenant systems (SRX Series, and vSRX)**—Starting in Junos OS Release 20.3R1, you can configure the **root-streaming** option at the **[edit logical-systems logical-systems-name security log]** and **[edit tenants tenants-name security log]** hierarchy levels in the stream mode for user logical system and tenant system. The **root-streaming** option allows the user logical systems and tenant systems to generate logs using the root system's stream configuration.

[See [root-streaming](#).]

Network Management and Monitoring

- **Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.3R1, you can use the **probe** command to query the status of the probed interface. The proxy interface resides on the same node as the probed interface, or it can reside on a node to which the probed interface is directly connected.

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

To enable the **probe** command, configure the **extended-echo** statement under the **[edit system]** hierarchy.

[See [Using the Probe command](#).]

- **SNMP support to export statistics of user firewall (SRX Series and vSRX)**—Starting in Junos OS Release 20.3R1, the following four new OIDs of MIB jnxUserFirewalls provide statistics of user firewall counters to SNMP:

- jnxUserFwDomainAuthTable
- jnxUserFwADDomCtrlTable
- jnxUserFwLDAPTable
- jnxUserFwProbeTable

The OID jnxUserFwDomainAuthTable provides statistics from multiple sources such as Active Directory (AD), Clearpass, and JIMS. The other three OIDs provide the statistics of AD only.

[See [Enterprise-Specific SNMP MIBs Supported by Junos OS](#).]

- **Enhancements to sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.3R1, devices running Junos OS with upgraded FreeBSD support the following enhancements to sessions over outbound HTTPS:
 - Connecting to multiple outbound HTTPS clients by configuring one or more clients at the **[edit system services outbound-https]** hierarchy level
 - Configuring multiple backup gRPC servers for a given outbound HTTPS client
 - Establishing a csh session
 - Establishing multiple, concurrent NETCONF and csh sessions between the device running Junos OS and an outbound HTTPS client
 - Configuring a shared secret that the outbound HTTPS client uses to authenticate the device running Junos OS
 - Authenticating the client using certificate chains in addition to self-signed certificates

[See [NETCONF and Shell Sessions over Outbound HTTPS](#).]

- **Real-time performance monitoring (RPM) with IP monitoring withdraw option (SRX380, SRX300, SRX320, SRX340, SRX345, and SRX550HM)**—When an RPM probe is successful, IP monitoring adds one or many primary routes to the routing table. Starting in Junos OS Release 20.3R1, we've introduced the option **withdraw** to remove the primary routes when RPM fails to probe the destination. When the primary routes are withdrawn, the traffic can choose other routes in the routing table. If no other routes exists, then the traffic is dropped.

In Junos OS releases before Release 20.3R1, when an RPM probes fails, IP monitoring adds a backup route. If the probe is later successful, the backup route is deleted.

To enable the **withdraw** option, use the **set services ip-monitoring policy *policy-name* then preferred-route withdraw** command.

[See [ip-monitoring \(Services\)](#) and [show services ip-monitoring status](#).]

Routing and Forwarding Options

- **Distributed mode support for BFD (SRX5000 line of devices with SPC3 card)**—Starting in Junos OS Release 20.3R1, we've introduced distributed mode for BFD failure detection. This mode provides faster BFD failure detection of 300 (3 x 100) ms. You can enable distributed mode when you configure the BFD failure detection timer to a value less than 500 ms.

For optimization and performance enhancement, you must configure the BFD failure detection timer value in multiples of 50 ms.

[See [detection-time \(BFD Liveness Detection\)](#).]

Security

- **Support for TLS profiles in Dynamic Address Feed Servers(NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.3R1, you can secure the communication channel between an SRX Series device and a feed server using a Transport Layer Security (TLS) profile. When you configure the **tls-profile** statement at the **[edit security dynamic-address feed-server]** hierarchy level, the SRX Series device and the feed server verify the server certificate and the client certificate in order to download dynamic address feed data on the device.

A valid CA certificate must be present on the SRX Series device. The device needs a client certificate configured in the SSL initiation profile to connect to the feed server.

[See [Encrypt Traffic Using SSL Proxy and TLS](#) and [tls-profile](#).]

Unified Threat Management (UTM)

- **UTM service inspection for pass-through IP-IP and GRE tunnel in TAP mode (SRX Series and vSRX)**—Starting in Junos OS Release 20.3R1, unified threat management (UTM) can inspect IP over IP (IP-IP) and GRE inner tunnel traffic in Terminal Access Point (TAP) mode by de-encapsulating the outer and inner IP headers up to two levels. You can configure up to eight TAP interfaces on SRX Series devices.

[See [SRX TAP Mode Support Overview](#).]

VPNs

- **IKEv2 configuration payload improvements with ikev2 process (SRX5000 line of devices with SPC3 and vSRX running ikev2)**—Starting in Junos OS Release 20.3R1, we've improved the IKEv2 configuration payload to support the following features:
 - IPv4 and IPv6 local address pool (you can also assign a fixed IP address to a peer).
 - Additional IKEv2 configuration attributes **INTERNAL_IP6_ADDRESS** and **INTERNAL_IP6_DNS**. See [Understanding Internet Key Exchange Version 2](#).
 - Allow the administrator to configure the RADIUS server with a framed pool associated with a peer or user.
 - Additional option, **none** introduced for **authentication-order**. See [authentication-order \(Access Profile\)](#).

- RADIUS accounting start and stop messages to indicate IKEv2 peer session up and down events.
- Introduction of IPv6 support allows dual stack tunnels using configuration payload.

[See [show security ike active-peer](#).]

- **Tunnel distribution profile and redistribution (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.3R1, we've optimized tunnel redistribution. After tunnel redistribution, the data path might not be optimal. You can use VPN session affinity to optimize the data path after tunnel redistribution. Note that the data path that is being optimized experiences a higher packet delay until it is fully optimized.

SRX Series devices don't support VPN session affinity by default. To enable this feature, use the **set security flow load-distribution session-affinity ipsec** command.

[See [session-affinity](#).]

- **Extended Sequence Number using IKEv2 (SRX5400, SRX5600, and SRX5800 devices)**—Starting from Junos OS Release 20.3R1, we provide support for Extended Sequence Number (ESN) in Mixed mode of SPC3 and SPC2 service cards.

[See [Understanding Extended Sequence Number \(ESN\)](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 307](#)
- [What's Changed in Release 20.3R2 | 307](#)
- [What's Changed in Release 20.3R1 | 309](#)

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

What's Changed in Release 20.3R3

Junos OS XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or **set refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

What's Changed in Release 20.3R2

Flow-Based and Packet-Based Processing

- On SRX Series devices in earlier releases, when the session table was full there was no alarm set to indicate this. Starting from this release, when the percent of flow session table utilization is 95% on FPC and PIC, an alarm message `? Flow session table is almost full on FPC <number> PIC <number>?` is set. Similarly, when the percent of DCP session table utilization is 95% on FPC and PIC, an alarm message `? DCP session table is almost full on FPC <number> PIC <number>?` is set.
- **Self-generated IKE packets chooses outgoing interface matching source IP Address (SRX Series)** — A self-generated Internet Key Exchange (IKE) packet always select the ECMP outgoing interface that matches source IP address. Note that filter-based forwarding for self-generated traffic with rerouting is not supported.

General Routing

- **Support for fully qualified domain name (FQDN) for log server (SRX Series)**—Starting in Junos OS Release, you can configure TTL value for a DNS server cache with hostname or IP address.

[See [Configuring the TTL Value for DNS Server Caching](#).]

MPLS

- **The `show mpls lsp extensive` and `show mpls lsp detail` commands display next-hop gateway LSPid** — When you use the `show mpls lsp extensive` and `show mpls lsp detail` commands, you'll see next-hop gateway LSPid in the output.

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppressing root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the **no-login-logout** parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

- **The `jcs:invoke()` function supports suppressing root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the **no-login-logout** parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified RPC. If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are logged in system log files.

[See [invoke\(\) Function \(SLAX and XSLT\)](#).]

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the **verbose** statement at the **edit system export-format json** hierarchy level. The default format to export configuration data in JSON changed from **verbose** format to **ietf** format starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the **edit system export-format json** hierarchy level. Although the **verbose** statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See [export-format](#).]

VPNs

- **The junos-ike package installed by default (SRX5400, SRX5600, and SRX5800)**— For the SRX5400, SRX5600, and SRX5800 devices with RE3 installed, the junos-ike package is installed by default. As a result, the iked and ikemd processes run on the Routing Engine by default instead of the IPsec key management daemon (kmd). In earlier Junos OS releases, the junos-ike package is an optional package for SRX5400, SRX5600, and SRX5800 devices with RE3, and IPsec Key Management Daemon (KMD) runs by default.

[See [Enabling IPsec VPN Feature Set on SRX5K-SPC3 Services Processing Card](#).]

- **IKE index displayed in show security ipsec security-associations detail output (SRX5400, SRX5600, and SRX5800)**—When you execute the **show security ipsec security-associations detail** command, a new output field, **IKE SA Index**, corresponding to every IPsec Security Association (SA) within a tunnel is displayed under each IPsec SA information.

[See [show security ipsec security-associations](#).]

What's Changed in Release 20.3R1

Authentication and Access Control

- **SSH protocol version 1 option deprecated from CLI (SRX Series)**—Starting in Junos OS Release 20.3R1, 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 20.3R1, continue to support the **v1** option to remotely manage systems and applications.

[See [protocol-version](#).]

Junos OS XML API and Scripting

- **Changes to Junos XML RPC request tag names (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've updated the Junos XML request tag name for some operational RPCs

to ensure consistency across the Junos XML API. Devices running Junos OS still accept the old request tag names, but we recommend that you use the new names going forward. The changes include:

- Most, but not all, request tag names that start with **show** replace **show** with **get** in the name.
- Uppercase characters are converted to lowercase.

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

J-Web

- **Change in the J-Web browser tab title (SRX Series)**—The J-Web browser tab title displays the device model and hostname. These details are also displayed when you hover over the J-Web browser tab.

For example, when you access J-Web for an SRX320 device with a host name `srx320-xyz`, the J-Web browser tab displays the title as *J-Web (srx320 - srx320-xyz)*.

If the hostname isn't configured, the J-Web browser tab title displays the host URL or IP address; for example, *J-Web (srx320 - <device IP address>)*.

Network Address Translation (NAT)

- **Port block allocation support (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 20.3R1, you can configure the port block allocation size from 1 through 64512. To save system memory, the recommended port block allocation size is 64. If you configure the port block allocation size to be lesser than 64, the system displays the warning message, **warning: To save system memory, the block size is recommended to be no less than 64.**

In releases earlier than Junos OS Release 20.3R1, you can configure port block allocation size from 1 through 64512 on SRX5400, SRX5600, and SRX5800 only.

[See [Configure Port Block Allocation Size.](#)]

System Logs

- **Option change-log is changed to default (SRX Series)**— Starting in Junos OS Release 20.3R1, the **change-log** is a default option at `[edit system syslog file name]` hierarchy for SRX Series devices. As the default option, **change-log** records all the configuration changes. In Junos OS releases earlier than 20.3R1, you need to configure change-log.

[See [file \(System Logging\).](#)]

Known Limitations

IN THIS SECTION

- [General Routing | 312](#)
- [J-Web | 312](#)
- [VPNs | 312](#)

Learn about known limitations in this release 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.

General Routing

- IPC cannot afford pushing 50,000 NAT rules with complicated matching criteria from Routing Engine to Packet Forwarding Engine. It's hard to simply provide how much rules with complicated matching criteria can be supported, because memory size consumption depends on different rule content. [PR1481845](#)
- Due to internal message failures between the Routing Engine and Packet Forwarding Engine, some packets get missed in the PCAP files while using the JDPI unknown packet capture feature. [PR1491919](#)
- Committing a large number of custom applications with a single member, a single context, and a varying pattern might result in significant time taken for completion of commit. Commit status can be checked using the show services application-identification commit-status command. [PR1493127](#)
- On SRX4600 device, chassis cluster does not support 1G SFPs for fabric links. For the HA fabric links one can use either the dedicated 10G HA fabric ports or the 40G revenue ports, both of which do not support 1G SFPs. In case 1G SFPs are accidentally used nevertheless, ISSU Junos OS upgrade is known to fail. Note that the dedicated HA control links can be configured to support 1G SFPs. [PR1496630](#)
- When a SRX Series device is configured with connection-mode as always, the secure client on the Android platform will always establish the connection even though the client is manually disconnected. [PR1537815](#)

J-Web

- For a spoke device in a hub-and-spoke topology, the UI will show VPN topology as Site to Site. [PR1495973](#)

VPNs

- When multiple traffic selectors are configured on a particular VPN, theiked process checks for a maximum of 1 DPD probe that is sent to the peer for the configured DPD interval. The DPD probe is sent to the peer if traffic flows over even one of the tunnels for the given VPN object. [PR1366585](#)
- On the SRX5000 line of devices with an SPC3 card, sometimes IKE SA is not seen on the device when the st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)
- In SPC2 and SPC3 mixed-mode HA deployments, tunnel per second (TPS) is getting affected while dead peer detection (DCD) is being served on existing tunnels. This limitation is due to a large chunk of CPU

being occupied by infrastructure (gencfg) used by IKED to synchronize its DPD state to the backup nodes. [PR1473482](#)

- IKE DH group24 and IPsec PFS group24 are not supported from Juniper Secure client, though these are supported on the SRX Series device. [PR1506966](#)

Open Issues

IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 313](#)
- [General Routing | 313](#)
- [Interfaces and Chassis | 314](#)
- [Intrusion Detection and Prevention \(IDP\) | 314](#)
- [J-Web | 314](#)
- [User Interface and Configuration | 315](#)
- [VPNs | 315](#)

Learn about open issues in this release 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.

Flow-Based and Packet-Based Processing

- Use an antireplay window size of 512 for IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). Hence, there are no out-of-order packets with 512 antireplay window size. [PR1470637](#)

General Routing

- On an SRX340 device with J-Flow version 9 configured, the flowd process might generate core files frequently when the device is busy. [PR1463689](#)
- Installing a custom YANG package on the SRX Series devices might fail with error messages because the relevant file system is available as Read-Only. [PR1484870](#)

- On SRX1500, SRX4100, and SRX4200 devices, the Routing Engine CPU busy may cause the chassis cluster FPC PIC to go down. There is a continuous TCP keepalive which gets exchanged between the Routing Engine, Junos OS VM, and Linux. If that keepalive expires without any ACK, then the connection between the Routing Engine and flowd process is declared down and FPC goes down. The TCP keepalive possibly can expire if the Routing Engine CPU is highly utilized due to which no other actions are served. [PR1507114](#)
- Junos OS DHCP server pool change br-lan will wait for the lease timeout before getting a new IP in the new DHCP pool. [PR1516241](#)
- In macOS platforms, when the client connects successfully, the client is not getting minimized to the tray icon and it stays connected and you need to manually minimize it. [PR1525889](#)
- IPsec rekey fails when the SRX Series device is configured with kilobyte-based lifetime in remote access solution. [PR1527384](#)
- When IKEv1 and IKEv2 gateways that are bound to remote access profiles have the same local address, the connection fails with the Juniper Secure Connect client. [PR1539323](#)
- The LLDP protocol can be configured on SRX4000 and SRX5000 lines of devices, which it is not actually supported on those platforms. [PR1540797](#)
- Kernel might stop, with VM core files generated, and the system might reboot continuously after five child interfaces are added to the reth interface on one node. This might cause service impact. [PR1551297](#)
- For SRX4200 device to enable GRE queuing need to enable tunnel-queuing. [PR1569471](#)

Interfaces and Chassis

- On SRX1500 device, the IRB ping through the IRB interface is not working due to Packet Forwarding Engine (PFE) failure. [PR1582989](#)

Intrusion Detection and Prevention (IDP)

- Starting from Junos OS Release 21.1, either greater-than or less-than are allowed for age-of-attack filter of dynamic attack group configuration. The age-of-attack field in signatures will be changed to CVE dates from activation dates. Anomalies and generic attacks will be part of all groups created. [PR1397599](#)

J-Web

- Configuration of global settings options of IPsec VPN such as TCP encap profile, IPsec power mode and IKE package installation are not supported from J-Web. [PR1496439](#)
- J-Web GUI does not allow you to save a rule if the cumulative shared objects are more than 2500 before the policy grid is saved. When there are several shared objects, there will be a noticeable delay in opening sources and destinations of a rule, and performing rule action. [PR1540047](#)

User Interface and Configuration

- Configure an interface with a logical unit under a logical system. Enter the logical system CLI and try to delete the same configuration. It won't allow as expected. Exit logical system CLI, enter normal CLI and delete the configuration. It results in mustd core sometimes. [PR1467818](#)

VPNs

- In the output of the show security ipsec inactive-tunnels command, Tunnel Down Reason is not displayed as this functionality is not supported in Junos OS Release 18.2R2 and later. [PR1383329](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, a new behavior has been introduced that differs from the behavior on the older SPC2 card. The SRX Series device with AutoVPN configuration can now accept multiple IPsec tunnels from a peer device (with the same source IP address and port number) using different IKE-IDs. [PR1407356](#)
- 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 and SPC2 mixed mode, with a very large amount of IKE peers (60,000) with dead peer detection (DPD) enabled, IPsec tunnels might flap in some cases when IKE and IPsec rekeys are happening at the same time. [PR1473523](#)
- Some TCP connections going through IPsec tunnels are getting struck after RG1 failover. [PR1477184](#)

Resolved Issues

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

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

Resolved Issues: 20.3R3

Forwarding and Sampling

- The configuration archive transfer-on-commit fails when running Junos OS Release 18.2R3-S6.5. [PR1563641](#)

General Routing

- The JNH memory leak could be observed on MPCs or MICs. [PR1542882](#)
- The kmd process might crash when the interface flaps. [PR1544800](#)

- FPCs might not boot-up on MX960 or EX9214 switches in a certain condition. [PR1545838](#)
- SRX1500 reports fans running at over speed. [PR1546132](#)
- On SRX4100 and SRX4200 devices, if PEM0 is removed, the output of jnxOperatingDescr.2 command might be incomplete. [PR1547053](#)
- On vSRX2.0, vSRX3.0, SRX1500, SRX4100, SRX4200, SRX4600 running chassis cluster in Junos OS Release 18.3 or later, multiple messages of "LCC: ch_cluster_lcc_set_context:564: failed to lock chassis_vmx mutex 11" are generated in the chassisd log file. These messages may recur after every few seconds and they do not have any impact on system operation. [PR1547953](#)
- PKI CMPv2 client certificate enrolment does not work on SRX when using root-CA. [PR1549954](#)
- On SRX1500, SRX-SFP-1GE-T(Part#740-013111) for a copper cable might be corrupted after reboot. [PR1552820](#)
- Application identity unknown packet capture utility does not function on SRX Series devices when enhanced-services mode is enabled. [PR1558812](#)
- The show security log report top idp group-by threat-severity order-by count top-number 5 where-attack command display changes. [PR1560027](#)
- The pkid process runs at 100 percent when the device is unable to connect to a particular URL. [PR1560374](#)
- The idpd process might stop when committing IDP configuration under logical systems and tenant systems during RGs failover. [PR1561298](#)
- The flowd process might pause and generates a core dump if JFlow version 9 is configured. [PR1567871](#)
- Wi-Fi mPIM on SRX Series devices is reaching out to NTP and DNS servers. [PR1569680](#)
- MACsec not using network-control queue. [PR1571977](#)
- Traffic going through the VRRP interface might be dropped when VRRP enabled IRB interface goes down. [PR1572920](#)
- The 1G interfaces might not come up after device reboot. [PR1585698](#)
- On SRX-Series devices, the protocol-version command which controls TLS-versions (1.1, 1.2, 1.3, etc) within SSL-Proxy has been unhidden. [PR1587149](#)

Interfaces and Chassis

- Backup Routing Engine or backup node may get stuck in bad status with improper backup-router configuration. [PR1530935](#)
- The configuration check out failed with error message: identical local address found on rt_inst [default], intfs. [PR1581877](#)

Intrusion Detection and Prevention (IDP)

- The IDP policy process might become unresponsive and fail to compile the IDP policy after an IDP automatic update. [PR1577684](#)

Network Management and Monitoring

- The mib2d process crashes and generates a core file on backup Routing Engine. [PR1557384](#)
- SSH connection might become unresponsive and logs show kern.maxfiles limit exceeded by uid messages. [PR1567634](#)

Routing Policy and Firewall Filters

- The junos-defaults construct within a unified-policies application match criteria now restricts the ports and protocols of a flow on a per-dynamic-application basis. [PR1551984](#)
- SecIntel connection name resolution errors due to SecIntel memory leaks. [PR1566128](#)
- Traffic loss might be seen when a big number of applications or addresses is referenced by one policy. [PR1576038](#)

Unified Threat Management (UTM)

- UTM license expiry event loss may cause the device to not quit the advanced service mode and maximum-sessions is decreased by half. [PR1563874](#)

VPNs

- IPsec traffic might get dropped after RG0 failover. [PR1522931](#)
- Traffic going through policy-based IPsec tunnel might be dropped after RG0 failover. [PR1550232](#)
- Theiked process might stop with Multinode High Availability setup. [PR1559121](#)
- A session might be closed when the session is created during the IPsec rekey. [PR1564444](#)
- When there are multiple IPsec SA, backup SA start IPsec rekey. [PR1565132](#)
- SPI mismatch caused by simultaneous rekeys under kmd stress. [PR1571105](#)

Resolved Issues: 20.3R2

Application Layer Gateways (ALGs)

- The SCCP ALG does not work on SRX Series devices running with Junos OS Release 17.3R1 and later. [PR1535356](#)

Chassis Clustering

- Disabled node on chassis cluster sent out ARP request packets. [PR1548173](#)

Flow-Based and Packet-Based Processing

- The rst-invalidate-session command does not work if configured together with the no-sequence-check command. [PR1541954](#)

General Routing

- A condition within TCP proxy could result in downloads becoming permanently stuck or not completing. TCP proxy is used by multiple services, including Juniper ATP Cloud in block mode, ICAP, SSL proxy, antivirus, content filtering, and antispam. [PR1502977](#)
- In a dual CPE scenario, if the rule match is completed before application identification is done, AppQoE moves the session to the other node. [PR1514973](#)
- FQDN-based security log stream does not dynamically update the IP address. [PR1520071](#)
- The TCP packet might be dropped if syn-proxy protection is enabled. [PR1521325](#)
- On SRX Series devices with chassis cluster, high CPU usage might be seen due to the llmd process. [PR1521794](#)
- Certificate validation might fail when OCSP is used and the OCSP server is a dual-stack device. [PR1525924](#)
- On the SRX1500 device, the traffic rate shown in the CLI command is not accurate. [PR1527511](#)
- The MAC table is null in Layer 2 mode after one pass-through session is created successfully. [PR1528286](#)
- On SRX4100 and SRX4200 devices, four out of eight fans might not work. [PR1534706](#)
- The firewall filter SA and DA tags are not in the log messages as expected in port details. [PR1539338](#)
- Packet drop might be seen when a packet with destination port 0 is received on the SRX380 device. [PR1540414](#)
- The nsd process might crash when DNS-based allowlisting is configured under SSL proxy. [PR1542942](#)
- Need syslog to indicate signature download completion. [PR1545580](#)
- The flowd process might generate core files when the user changes the flow mode configuration to packet mode. [PR1546653](#)
- On SRX4100 and SRX4200 devices, if PEM0 is removed, the output of jnxOperatingDescr.2 command might be incomplete. [PR1547053](#)
- Advanced anti-malware file/email statistics not gets increment with latest PB version. [PR1547094](#)

- On vSRX2.0, vSRX3.0, SRX1500, SRX4100, SRX4200, SRX4600 running chassis cluster in Junos OS Release 18.3 or later releases, multiple messages of "LCC: ch_cluster_lcc_set_context:564: failed to lock chassis_vmx mutex 11" are generated in the chassisd log file. These messages may recur after every few seconds and they do not have any impact on system operation. [PR1547953](#)
- LCMD log "gw_cb_presence:136: PEM(slot = 0): error detecting presence (fruid = 15, drv_id = 30, status = -11)" generates every second on the SRX4100 and SRX4200 devices. [PR1550249](#)
- On SRX1500, SRX-SFP-1GE-T(Part#740-013111) for a copper cable might be corrupted after reboot. [PR1552820](#)
- An ipfd core file might be generated when using adaptive threat profiling. [PR1554556](#)
- On SRX550M device, the dumpdisklabel command fails with message "ERROR: Unknown platform srx550m". [PR1557311](#)
- The outbound-ssh routing-instance command shown as unsupported. [PR1558808](#)
- Application identity unknown packet capture utility does not function on SRX Series devices when enhanced-services mode is enabled. [PR1558812](#)
- The PKID process runs at 100 percent when the device is unable to connect to a particular URL. [PR1560374](#)

Interfaces and Chassis

- When SRX Series devices receive proxy ARP requests on VRRP interfaces, SRX Series devices send ARP replies with the underlying interface MAC address. [PR1526851](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srxpfe process might generate core files during the idpd process commit on SRX Series devices. [PR1521682](#)
- The srxpfe process might stop when the adaptive threat profiling is configured within an IDP rule base and security log is enabled. [PR1532737](#)
- Adaptive threat profiling incorrectly classifies hosts when Server-to-Client (S2C) IDP signatures are used. [PR1533116](#)
- IDP Policy load might fail post image upgrade for Junos OS 15.1X49 releases. [PR1546542](#)

J-Web

- Sometimes, when you edit the local gateway in the remote access VPN workflow under VPN>IPsec VPN, J-Web might not display one or more drop-down values. [PR1521788](#)
- In the SRX5000 line of devices, J-Web can take up to 60 seconds to 90 seconds to load 60000 security policies. [PR1521841](#)
- The "+" button is not shown at J-Web interface menu. [PR1550755](#)

Platform and Infrastructure

- Syslog reporting "PFE_FLOWD_SELFPING_PACKET_LOSS: Traffic impact: Selfping packets loss/err: 300 within 600 second" error messages in node 0 and node 1 control panel. [PR1522130](#)
- Commit not failing as expected after removing the reth interface. [PR1538273](#)

Routing Policy and Firewall Filters

- The show security dynamic-address feed-name command could not list secprofiling feed. [PR1537714](#)
- The flowd or srpxfe process might crash when an SRX Series or NFX Series device running Junos OS Release 18.2R1 or later supports the unified policy feature. [PR1544554](#)
- Traffic might be dropped unexpectedly when the URL category match condition is used on a security policy. [PR1546120](#)
- Global policies working with multi-zones cause high CPU utilization. [PR1549366](#)
- NSD process stops when the secprofiling feed name is 64 bytes. [PR1549676](#)
- On SRX5000 line of devices, the secondary node might get stuck in performing ColdSync after a reboot, upgrade or if ISSU is performed [PR1558382](#)
- The traffic might be dropped due to inserting one global policy above others on SRX Series devices. [PR1558827](#)

Subscriber Access Management

- Incorrect counter type (counter instead of gauge) specified for some values in MIB jnxUserAAAMib. [PR1533900](#)

Unified Threat Management (UTM)

- Stream buffer memory leak might happen when UTM is configured under unified policies. [PR1557278](#)

VPNs

- On SRX5000 line of devices with SPC3 card, when the encryption algorithm is not configured in IPsec proposal, the output of show security ipsec security-associations command might display empty space instead of keyword null for encryption algorithm. [PR1507270](#)
- The traffic might be dropped when IPsec VPN with NAT-T enabled. [PR1522017](#)
- IPsec traffic may get dropped after RGO failover. [PR1522931](#)
- On all SRX series devices using IPsec with NAT Traversal, MTU size for the external interface might be changed after IPsec SA is re-established. [PR1530684](#)
- After IPsec tunnel using policy-based VPN is overwritten by another VPN client, traffic using this IPsec tunnel will be dropped. [PR1546537](#)
- Traffic goes through policy-based IPsec tunnel might be dropped after RGO failover. [PR1550232](#)

- Theiked process might stop with Multinode High Availability setup. [PR1559121](#)
- A session might be closed when the session is created during the IPsec rekey. [PR1564444](#)

Resolved Issues: 20.3R1

Application Security

- AppQoE support for dynamic-application. [PR1503400](#)

Chassis Clustering

- If a cluster ID of 16 or multiple of 16 is used, the chassis cluster might not come up. [PR1487951](#)
- The ISSU fails with timeout due to cold synchronization failure. [PR1502872](#)

Flow-Based and Packet-Based Processing

- The show security group-vpn server statistics |display XML is not in expected format. [PR1349959](#)
- Constant memory leak might lead to FPC memory exhaustion. [PR1381527](#)
- ECMP load balancing does not happen when RG1 node 0 is secondary. [PR1475853](#)
- On Web proxy, memory leak is observed in association hash table and DNS hash table. [PR1480760](#)
- CLI autocomplete is now available for both secintel and advanced anti-malware products. [PR1487419](#)
- Risk of service interruption is probable on SRX Series devices with a dual-stacked CA server. [PR1489249](#)
- GRE or IPsec tunnel might not come up when the set security flow no-local-favor-ecmp command is configured. [PR1489276](#)
- Not able to clear the warm sessions on the peer SRX Series devices. [PR1493174](#)
- SRX Series devices now keep a local copy of configuration changes within /var/log/configuration-log. [PR1493842](#)
- Phone client stop seen while configuring SRX345 device ZTP with CSO. [PR1496650](#)
- Outbound SSH connection flap or memory leak issue might be observed while pushing the configuration to ephemeral DB with a high rate. [PR1497575](#)
- Unexpected flow logging traffic beyond the packet filter. [PR1497939](#)
- Traffic interruption happens due to MAC address duplication between two devices running Junos OS. [PR1497956](#)
- Don't use uppercase characters for source-identity when using the show security match-policies command. [PR1499090](#)
- J-Flow version 9 does not display correct outgoing interface for APBR traffic. [PR1502432](#)

- A condition within TCP proxy could result in downloads becoming permanently stuck or not completing. TCP proxy is used by multiple services, including Juniper ATP Cloud in block mode, ICAP, SSL proxy, anti-virus, content filtering, and anti-spam. [PR1502977](#)
- The cfmd core observed when LTM is triggered for the session configured on ethernet-switching interface without bridge domain configuration. [PR1503696](#)
- Layer 2 ping is not working with remote mep. [PR1504986](#)
- SOF asymmetric scenario is not working with the phase 1 solution. [PR1507865](#)
- VRRP does not work on the redundant Ethernet interface with a VLAN ID greater than 1023. [PR1515046](#)
- PCAP file generated using packet capture was improper on the SRX5000 line of devices. [PR1515691](#)
- A logic issue was corrected in SSL proxy that could lead to an srxpfe or flowd core file under load. [PR1516903](#)
- The PPPoE session does not come up after return to zero on SRX Series devices. [PR1518709](#)
- TAP mode behavior has been improved and the configuration has been greatly simplified. [PR1521066](#)
- Adaptive Threat Profiling would stop submitting new IP addresses to a feed after a limit of 10,000 has been reached. [PR1524284](#)
- Commit confirmed rollback is not working. [PR1527848](#)

Infrastructure

- The installation fails when upgrading from legacy Junos OS to specific BSDx-based Junos OS. [PR1505864](#)

Interfaces and Chassis

- All interfaces remain in the down status after the SRX300 line of devices power up or reboot. [PR1488348](#)
- Continuous drops are seen in control traffic when high amount of data queues in one SPC2 PIC. [PR1490216](#)
- PPO IPv6 route does not work. [PR1495839](#)
- Fabric interface might be monitored down after chassis cluster reboot. [PR1503075](#)

Intrusion Detection and Prevention (IDP)

- When intelligent inspection status changes, syslog is not getting generated on SRX300 and SRX500 line of devices. [PR1448365](#)
- Configuring anomaly occurs in CLI. [PR1490437](#)
- The IDP attack detection might not work in a specific situation. [PR1497340](#)
- IDP's custom-attack time-binding interval command was mistakenly hidden within the CLI. [PR1506765](#)
- Adaptive Threat Profiling incorrectly classifies hosts when Server-to-Client (S2C) IDP signatures are used. [PR1533116](#)

J-Web

- You cannot configure Redundant PSU and Power Budget Statistics on the SRX380 device which is in HA mode through J-Web. [PR1493713](#)
- The J-Web users might not be able to configure PPPoE using the PPPoE wizard. [PR1502657](#)
- J-Web chassis status widget is incorrectly reporting temperature alarms. [PR1507156](#)
- The parameters show another LSYS at J-Web in a multiple LSYS scenario. [PR1518675](#)

MPLS

- BGP session flaps between two directly connected BGP peers because of the wrong TCP-MSS in use. [PR1493431](#)

Network Address Translation (NAT)

- Not all NAT sessions are synchronized from Node 1 to Node 2. [PR1473788](#)

Platform and Infrastructure

- The SRX1500 and the SRX4000 line of devices might boot up with rescue configuration after a power outage. [PR1490181](#)
- Packets get dropped when the next hop is IRB over It interface. [PR1494594](#)

Routing Policy and Firewall Filters

- On SRX Series devices, in a very rare condition, security policies don't synchronize between the Routing Engine and Packet Forwarding Engine. This issue might cause traffic loss. [PR1453852](#)
- Session-close security-logging is now enabled by default for pre-id-default-policy. [PR1491698](#)
- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. [PR1492436](#)
- The srpxfe or flowd process might stop due to memory corruption within JDPI. [PR1500938](#)
- Traffic might fail to hit policies if match dynamic-application and match source-end-user-profile options are configured under the same security policy name. [PR1505002](#)
- Junos OS upgrade may encounter failure in certain conditions when enabling ATP. [PR1519222](#)

Routing Protocols

- The BGP route target family might prevent the route reflector 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 an SRX4200 device, 35 percent of drop is seen in all TPS cases. [PR1481625](#)

- On SRX Series devices with SPC3, when overlapping traffic-selectors are configured, multiple IPsec SAs get negotiated with the peer device. [PR1482446](#)
- Some options under IKE and IPsec policy and proposal help text description should change to NOT RECOMMENDED. [PR1487515](#)
- Use different XML tags for local and remote IKE IDs to avoid confusion. [PR1493368](#)
- Issue with XML rpc show security ipsec tunnel-distribution summary output. [PR1494274](#)
- The SRX5000 line of devices with SPC3 was not supporting simultaneous IKE negotiation. [PR1497297](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R3 documentation for the SRX Series.

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

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 19.2 to the next three releases – 19.3, 19.4 and 20.1 or downgrade to the previous three releases – 19.1, 18.4 and 18.3.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 19.2 is an EEOL release. Hence, you

can upgrade from 19.2 to the next two EEOL releases – 19.3 and 19.4 or downgrade to the previous two EEOL releases – 19.1 and 18.4.4.

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade and Downgrade to subsequent 3 releases	Upgrade and Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see <https://www.juniper.net/support/eol/junos.html>.

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

Junos OS Release Notes for vMX

IN THIS SECTION

- [What's New | 326](#)
- [What's Changed | 327](#)
- [Known Limitations | 328](#)
- [Open Issues | 329](#)
- [Resolved Issues | 329](#)
- [Licensing | 330](#)
- [Upgrade Instructions | 330](#)

These release notes accompany Junos OS Release 20.3R3 for vMX. 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 20.3R3 | 326](#)
- [What's New in Release 20.3R2 | 326](#)
- [What's New in Release 20.3R1 | 326](#)

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

What's New in Release 20.3R3

There are no new features for vMX in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features for vMX in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Installation and Upgrade

- **RHEL version 7.7 support (vMX)**—Starting with Junos OS Release 20.3R1, we provide support for Red Hat Enterprise Linux (RHEL) version 7.7. You can use RHEL v7.7 to install vMX on a kernel-based virtual machine (KVM) by running the installation scripts provided by Juniper Networks.

[See [Minimum Hardware and Software Requirements](#).]

Juniper Extension Toolkit (JET)

- **Juniper Extension Toolkit (JET) supports BFD Service APIs for routing protocol process (rpd) programmability (MX Series, PTX Series, QFX Series, and vMX)**—Starting in Junos OS Release 20.3R1, you can use programmable rpd (prpd) BFD APIs to add, update, and delete BFD sessions and subscribe to BFD events from outside applications. These APIs enable the integration of rpd with software-defined networking (SDN) controllers and increase the flexibility of your network. The prpd BFD APIs support BFD Echo-Lite sessions in single-hop IPv4 and IPv6 modes.

The following BFD Service APIs are supported:

- Initialize
- SessionAdd

- SessionUpdate
- SessionDelete
- SessionDeleteAll
- Subscribe
- Unsubscribe

Use the **show bfd session extensive** command to view BFD sessions. BFD sessions added through prpd BFD APIs are labeled with **PRPD:<session-id>** in the client field. The **<session-id>** is 1 for the first BFD session that is added, 2 for the second, and so on.

[See [show bfd session extensive](#) and [JET APIs on Juniper EngNet.](#)]

System Management

- **Higher scale and performance in RIFT (QFX5100, QFX5110, QFX10000, MX960, and vMX)**— Starting in Junos OS Release 20.3R1, we've made the following improvements to increase the scalability and performance in Routing in Fat Tree (RIFT):
 - Prefixes in RIFT
 - Peers in RIFT
 - Convergence improvement with RIFT
 - BFD sessions with RIFT

[See [RIFT Overview.](#)]

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 328](#)
- [What's Changed in Release 20.3R2 | 328](#)
- [What's Changed in Release 20.3R1 | 328](#)

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

What's Changed in Release 20.3R3

Junos OS XML API and Scripting

- Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or set **refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

What's Changed in Release 20.3R2

There are no changes in behavior or syntax for vMX in Junos OS Release 20.3R2.

What's Changed in Release 20.3R1

There are no changes in behavior or syntax for vMX in Junos OS Release 20.3R1.

Known Limitations

Learn about known limitations in this release for vMX.

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

There are no known behavior or limitation for vMX in Junos OS Release 20.3R3.

Open Issues

IN THIS SECTION

- [Platform and Infrastructure](#) | 329

Learn about open issues in this release for vMX.

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

Platform and Infrastructure

- Issue is seen only in VMX setups with the blockpointer in the ktree infra is getting corrupted leading to core file generation. There is no function impact such as fpc restart or system down and the issues won't be observed in hardware setups. [PR1525594](#)

Resolved Issues

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

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

Resolved Issues: 20.3R3

There are no resolved issues for vMX in Junos OS Release 20.3R3.

Resolved Issues: 20.3R2

Routing Protocols

- When admin-color based policy evaluation happens with the policy lfa configuration, the backup next hop chosen (among the different backup next hops possible) might not be correct. [PR1558581](#)

Resolved Issues: 20.3R1

There are no resolved issues for vMX in Junos OS Release 20.3R1.

Licensing

Starting in Junos OS Release 19.2R1, Juniper Agile Licensing introduces a new capability that significantly improves the ease of license management network wide. The Juniper Agile License Manager is a software application that runs on your network and provides an on-premise repository of licenses that are dynamically consumed by Juniper Networks devices and applications as required. Integration with Juniper's Entitlement Management System and Portal provides an intuitive extension of the existing user experience that enables you to manage all your licenses.

- The Agile License Manager is a new option that provides more efficient management of licenses, but you can continue to use individual license keys for each device if required.
- To use vMX or vBNG feature licenses in Junos OS Release 19.2R1 version, you need new license keys. Previous license keys will continue to be supported for previous Junos OS releases, but for the Junos OS 19.2R1 release and later you need to carry out a one-time migration of existing licenses. Contact [Customer Care](#) to exchange previous licenses. Note that you can choose to use individual license keys for each device, or to deploy Agile License Manager for more efficient management of licenses.
- For more information about Agile Licensing keys and capabilities, see [Juniper Agile Licensing portal FAQ](#).

See [Juniper Agile Licensing Guide](#) for more details on how to obtain, install, and use the License Manager.

Upgrade Instructions

You cannot upgrade Junos OS for the vMX router from earlier releases using the **request system software add** command.

You must deploy a new vMX instance using the downloaded software package.

Remember to prepare for upgrades with new license keys and/or deploying Agile License Manager.

Junos OS Release Notes for vRR

IN THIS SECTION

- [What's New | 331](#)
- [What's Changed | 333](#)
- [Known Limitations | 334](#)
- [Open Issues | 334](#)
- [Resolved Issues | 335](#)

These release notes accompany Junos OS Release 20.3R3 for vRR. 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 20.3R3 | 331](#)
- [What's New in Release 20.3R2 | 332](#)
- [What's New in Release 20.3R1 | 332](#)

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

To learn about common BGP or routing Junos features supported on vRR for Junos OS 20.3R3, see [What's New](#) for MX Series routers.

What's New in Release 20.3R3

There are no new features for vRR in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features for vRR in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Routing Protocols

- **Support for implicit filter for default EBGp route propagation behavior without policies (ACX Series, JRR200, MX204, vRR and PTX5000)**—Starting in Junos OS Release 20.3R1, we've introduced a new configuration hierarchy, **defaults ebgp no-policy** at the existing **[edit protocols bgp]** hierarchy level. The configuration option separates the default policy for **receive** and **advertise**, into separate clauses (**accept**, **reject**, or **reject-always**) to allow the route propagation behavior of EBGp speakers to vary independently from its default behavior.

In earlier releases, the default behavior of BGP was to receive and advertise all routes. With the introduction of this feature, the default behavior still remains to “accept” all routes for both **receive** and **advertise**, but you also have an option to reject routes by default.

With the **reject** configuration, you can reject routes of type inet unicast and inet6 unicast in instance types master, vrf, virtual-router, and non-forwarding. With the **reject-always** configuration, you can reject all routes from being received or getting advertised, irrespective of address family or instance type. By using this feature, you can control traffic in leaf autonomous systems (AS) and thereby, prevent them from having to accidentally function as transit autonomous systems.

NOTE: The introduction of this implicit filter does not affect the existing deployments that rely on the default behavior.

[See [Implicit Filter for Default EBGp route propagation behavior without policies](#) and [defaults](#)]

- **BGP sharding for IPv4 and Ipv6 L3VPN, BGP-LU (MX Series, PTX-Series and vRR)**—Starting in Release 20.3R1, Junos OS supports BGP sharding and update IO features for these IPv4 and Ipv6 address families:
 - inet-vpn unicast
 - inet-vpn multicast (vrf.inet.2)
 - inet6-vpn unicast
 - inet6-vpn multicast (vrf.inet.2)
 - inet labeled-unicast
 - inet6 labeled-unicast

To enable BGP sharding, configure **rib-sharding** at the **[edit system processes routing bgp]** hierarchy level. Sharding is dependent on the update I/O thread feature. To enable update I/O, configure **update-threading** at the **[edit system processes routing bgp]** hierarchy level.

BGP Sharding is supported only on 64-bit routing protocol process (rpd) where the Routing Engine has at least 4 CPU cores and 16 GB of memory. To enable your device to always use 64-bit mode, use **set force-64-bit** at **[edit system processes routing]** hierarchy level. If you configure rib-sharding on a routing engine, RPD creates sharding threads. By default, the number of sharding threads created is the same as the number of CPU cores on the routing engine. Optionally, you can specify the number-of-shards you want to create. To set the number of sharding threads, use **set number-of-shards <number-of-shards>** at **[edit system processes routing bgp rib-sharding]** hierarchy level. To set the number of update threads, use **set number-of-threads <number-of-threads>** at the **[edit system processes routing bgp update-threading]** hierarchy level. To enable your device to always use 64-bit mode, use **set force-64-bit** at **[edit system processes routing]** hierarchy level.

[See [rib-sharding](#) and [update-threading](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 333](#)
- [What's Changed in Release 20.3R2 | 333](#)
- [What's Changed in Release 20.3R1 | 333](#)

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

To learn more about common BGP or routing changes in behavior or syntax in Junos OS 20.3R3, see [What's Changed](#) for MX Series routers.

What's Changed in Release 20.3R3

There are no changes in behavior or syntax for vRR in Junos OS Release 20.3R3.

What's Changed in Release 20.3R2

There are no changes in behavior or syntax for vRR in Junos OS Release 20.3R2.

What's Changed in Release 20.3R1

There are no changes in behavior or syntax for vRR in Junos OS Release 20.3R1.

Known Limitations

IN THIS SECTION

- [Routing Protocols](#) | 334

Learn about known limitations in this release for vRR.

To learn more about common BGP or routing known limitation in Junos OS 20.3R3, see [Known Limitations](#) 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.

Routing Protocols

- Commit check fails when rib-sharding is configured with these statements:
 - `routing-instances <name> routing-options multipath`
 - `routing-instances <name> routing-options policy-multipath`
 - `routing-instances <name> protocols mvpn.`

Open Issues

Learn about open issues in this release for vRR.

To learn more about common BGP or routing open issues in Junos OS 20.3R3, see [Open Issues](#) 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.

There are no known issues for vRR in Junos OS Release 20.3R3.

Resolved Issues

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

To learn more about common BGP or routing resolved issues in Junos OS 20.3R2, see [Resolved Issues](#) 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: 20.3R3

There are no resolved issues for vRR in Junos OS Release 20.3R3.

Resolved Issues: 20.3R2

General Routing

- 6PE prefixes might not be removed from RIB upon reception of withdrawal from a BGP neighbor when RIB sharding is enabled. [PR1556271](#)
- BGP flap and rpd crash might be observed. [PR1545837](#)

Resolved Issues: 20.3R1

There are no resolved issues for vRR in Junos OS Release 20.3R1.

Junos OS Release Notes for vSRX

IN THIS SECTION

- [What's New | 336](#)
- [What's Changed | 340](#)
- [Known Limitations | 341](#)
- [Open Issues | 342](#)
- [Resolved Issues | 343](#)
- [Migration, Upgrade, and Downgrade Instructions | 346](#)

These release notes accompany Junos OS Release 20.3R3 for vSRX. 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 20.3R3 | 336](#)
- [What's New in Release 20.3R2 | 336](#)
- [What's New in Release 20.3R1 | 337](#)

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

What's New in Release 20.3R3

There are no new features for vSRX in Junos OS Release 20.3R3.

What's New in Release 20.3R2

There are no new features for vSRX in Junos OS Release 20.3R2.

What's New in Release 20.3R1

Interfaces and Chassis

- **TAP mode support for (vSRX 3.0)**—Starting in Junos OS Release 20.3R1, TAP mode is supported for IDP, UTM, and UserFW on vSRX 3.0 to generate security log information and to display the information on threats detected, application usage, and user details according to the incoming traffic.

Both client to server and server to client traffic is directed to vSRX port using switch mirror or fiber tap. In this mode, vSRX 3.0 receives packet only from the configured TAP interface. All sending packets to TAP interface are dropped silently before leaving the vSRX instance. Except the configured TAP interface, other interface can be configured as standard interface and can be used as management interface or connected to outside server.

Use the **set security forwarding-options mode tap interface <interface-name>** command to configure TAP mode on an interface.

To disable this TAP mode, delete the TAP mode for the related interface and the related zone and policy configuration of that interface.

[See [TAP Mode Support Overview](#), [TAP Mode for IDP](#), [TAP Mode for Security Zones and Policies](#), and [forwarding-options \(Security\)](#).]

Juniper ATP Cloud

- **Support for integration of AWS GuardDuty with vSRX Firewalls and Juniper ATP Cloud (vSRX)**—Starting with Junos OS Release 20.3R1, we support threat feeds from Amazon Web Services (AWS) GuardDuty. The threats are sent as a security feed to the vSRX firewalls in the AWS environment. The vSRX firewalls can access the feeds either by directly downloading it from the AWS S3 bucket or, if the vSRX firewall is enrolled with Juniper ATP Cloud, the feed is pushed to the firewall device along with the security intelligence (SecIntel) feeds.

[See [Integrate AWS GuardDuty with vSRX Firewalls](#).]

Junos Telemetry Interface

- **Packet Forwarding Engine and Routing Engine sensor support on JTI (SRX1500, SRX4100, SRX4200, SRX4600, and vSRX)**—Junos OS Release 20.3R1 provides streaming support for revenue interface statistics through Packet Forwarding Engine (PFE) sensors and pseudo interface statistics through Routing Engine sensors. Sensors are supported through Junos telemetry interface (JTI) and remote procedure calls (gRPC) or gRPC Network Management Interface (gNMI) services. gNMI service is also enabled for other supported Routing Engine sensors.

Using JTI and gRPC or gNMI services, you can stream telemetry statistics to an outside collector.

These interface sensors are supported:

- Physical interfaces (IFD) (resource path `/interfaces/interface/`).
- Logical interfaces (IFL) (resource path `/interfaces/interface/subinterfaces/`).

These Routing Engine sensors are supported using gNMI services (previously, only gRPC services were supported):

- System events (resource path **/junos/events**).
- BGP peer information (resource path **/network-instances/network-instance/protocols/protocol/bgp/**).
- Memory utilization for routing protocol task (resource path **/junos/task-memory-information/**).
- Operational state of Routing Engines, power supply modules, Switch Fabric Boards, Control Boards, Switch Interface Boards, Modular Interface Cards, and Physical Interface Cards (resource path **/components/**).
- Link Layer Discovery Protocol (LLDP) (resource path **/lldp/**).
- Address Resolution Protocol (ARP) statistics for IPv4 routes (resource path **/arp-information/**).
- Network Discovery Protocol (NDP) table state information for IPv6 routes (resource path **/nd6-information/**).
- NDP router-advertisement statistics (resource path **/ipv6-ra/**).
- IS-IS routing protocol statistics (resource path **/network-instances/network-instance/protocols/protocol/isis/levels/level/** and **network-instances/network-instance/protocols/protocol/isis/interfaces/interface/levels/level/**).

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

Management

- **Enhanced Service Mode Support (vSRX 3.0)**—Starting in Junos OS Release 20.3R1, vSRX 3.0 supports Enhanced Service Mode (ESM). When this mode is enabled, vSRX 3.0 can support maximum of 128K sessions for Layer 7 services with increased service memory and the number of L4 sessions will be reduced to 50%.

By default, ESM is disabled and the vSRX 3.0 is in basic firewall mode. You can enable ESM using the **set security forwarding-process enhanced-services-mode** command. After enabling this mode, you need to reboot the instance.

When you enable this configuration, you will receive a warning message **warning: You have changed enhanced services mode. You must reboot the system for your change to take effect. If you have deployed a cluster, be sure to reboot all nodes.**

[See [forwarding-process](#) and [show security flow status](#).]

Performance and Scaling

- **Scaling vSRX 3.0 using Microsoft Azure Load Balancer and Virtual Machine Scale Sets (vSRX 3.0)**—Starting in Junos OS Release 20.3R1, vSRX 3.0 can automatically scale out or scale in for internal and outbound traffic using Azure Load Balancer (LB) and Microsoft Azure Virtual Machine Scale Sets (VMSS).

vSRX 3.0 instances are inline firewalls and any throughput or connection scaling limitations on these firewalls limit the performance and scaling of the entire virtual network. In such cases autoscaling of infrastructure for traffic inside the virtual network and for the outbound traffic is required. You can use the suggested deployments with Azure Load Balancer and Virtual Machine Scale Sets to achieve vSRX 3.0 scaling and better performance for your business needs.

[See [vSRX 3.0 Scaling for Internal and Outbound Traffic Using Azure Load Balancer and Virtual Machine Scale Sets](#).]

VPNs

- **Increase in IPsec VPN tunnels (vSRX)**—Starting in Junos OS Release 20.3R1, vSRX instances support up to 10,000 IPsec VPN tunnels. Previously, vSRX instances with 17 vCPUs supported 512 IPsec VPN tunnels.

To support the increased number of IPsec VPN tunnels, a minimum of 19 vCPUs are required. Out of the 19 vCPUs, 3 vCPUs must be dedicated to RE.

You must run the **request system software add optional://junos-ike.tgz** command the first time you wish to enable increased IPsec tunnel capacity. For subsequent software upgrades of the instance, the junos-ike package is upgraded automatically from the new Junos OS releases installed in the instance. If chassis cluster is enabled then run this command on both the nodes.

You can configure the number of vCPUs allocated to Junos Routing Engine using the **set security forwarding-options resource-manager cpu re <value>**. You must reboot the system to activate the new

vCPU allocation for RE and Flow RT threads. Run the **show security forward-options resource-manager status** command to verify the vCPU allocation between routing engine and the flow RT threads.

[See [Junos OS Features Supported on vSRX, forwarding-options \(Security\)](#), and [show security forward-options resource-manager](#).]

- **Increased Tunnel Scaling (vSRX)**—Starting in Junos OS Release 20.3R1, vSRX is supported by a new architecture similar to SRX5000 line of devices with SPC3 which increases the tunnel scale.

IPsec VPN features that are supported on SRX5000 line of devices with SPC3 (SRX5K-SPC3) are also supported on vSRX instances.

By default, when the vSRX boots up, the legacy architecture is executed. To enable the new architecture its mandatory to load and install this new junos-ike package. This is an optional package that is included in the Junos OS release. As an administrator, you must execute the **request system software add optional://junos-ike.tgz** command to load the junos-ike package.

[See [IPsec VPN Features and Configurations Not Supported on SRX5K-SPC3 and vSRX Instances](#).]

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R3 | 341](#)
- [What's Changed in Release 20.3R2 | 341](#)
- [What's Changed in Release 20.3R1 | 341](#)

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

What's Changed in Release 20.3R3

Junos OS XML API and Scripting

- Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the **request system scripts refresh-from** operational mode command, include the **cert-file** option and specify the certificate path. Before you refresh a script using the **set refresh** or set **refresh-from** configuration mode command, first configure the **cert-file** statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See [request system scripts refresh-from](#) and [cert-file](#).]

What's Changed in Release 20.3R2

There are no changes in behavior or syntax for vSRX in Junos OS Release 20.3R2.

What's Changed in Release 20.3R1

There are no changes in behavior or syntax for vSRX in Junos OS Release 20.3R1.

Known Limitations

IN THIS SECTION

- [Intrusion Detection and Prevention \(IDP\) | 342](#)
- [J-Web | 342](#)

Learn about known limitations in this release for vSRX.

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

Intrusion Detection and Prevention (IDP)

- Disable IDP before upgrading vSRX from a Junos OS Release 15.1X49 to Junos OS Release 17.4 or higher releases. Due to a change in IDP database format after Junos OS Release 15.1X49, there is no IDP database initially after the upgrade and the IDP configuration may fail to load, potentially leading to the entire Junos OS configuration not to load at the first bootup after the upgrade. After the upgrade, first download and install the IDP security package before re-enabling IDP again. [PR1455125](#)

J-Web

- For a spoke device in a hub-and-spoke topology, the UI will show VPN topology as Site to Site. [PR1495973](#)

Open Issues

IN THIS SECTION

- [J-Web | 342](#)
- [User Access and Authentication | 343](#)

Learn about open issues in this release for vSRX.

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

J-Web

- Configuration of global setting options of IPSec VPN such as TCP-Encap profile, IPSec Power Mode, and IKE package installation is not supported from the UI. [PR1496439](#)
- J-Web does not allow you to save a rule if the cumulative shared objects are more than 2,500 before the policy grid is saved. When there are several shared objects, there will be a noticeable delay in opening sources and destinations of a rule, and performing rule action. [PR1540047](#)

User Access and Authentication

- On vSRX 3.0 on Azure, with Microsoft Azure Hardware Security Module (HSM) enabled, keypair generation fails if the user re-uses the certificate ID for creating a new keypair, even if the previous keypair has been deleted. [PR1490558](#)

Resolved Issues

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

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

Resolved Issues: 20.3R3

Intrusion Detection and Prevention (IDP)

- Application identification related signatures might not get triggered. [PR1588450](#)

Platform and Infrastructure

- The queue-counters-srx-reserved-buffer-bytes count is 625000 Bytes, expected buffer is 2500000. [PR1538286](#)
- IKE configure mode payload is not pushing secondary DNS and secondary WINS attributes to Xauth module with IKEv1. Hence, the client is not getting assigned with secondary DNS and secondary WINS with IKEv1. [PR1558831](#)
- Delay in vSRX CLI prompt might be observed. [PR1559741](#)
- The pkid process runs at 100 percent when the device is unable to connect to a particular URL. [PR1560374](#)
- If committing source-address <addr> routing-instance and then delete source-address <addr> in private edit mode, commit fails with warning message. [PR1582529](#)

Routing Protocols

- Traffic might be lost during mirror data transmit from the primary ppmmd or bfdd. [PR1570228](#)

VPNs

- A session might be closed when the session is created during the IPsec rekey. [PR1564444](#)

Resolved Issues: 20.3R2

Application Security

- During rare circumstances, if the AppID unknown packet capture functionality is enabled, the srpxpfe process might crash and generate a core file. [PR1538991](#)

Chassis Clustering

- The control link might be broken when there is excessive traffic load on the control link in vSRX cluster deployment. [PR1524243](#)

CLI

- On Microsoft Azure deployments, SSH public key authentication is not supported for vSRX 3.0 CLI and portal deployment. [PR1402028](#)
- Commit is not successful when the configuration is committed without active probe settings options (all options under active probe settings are optional). [PR1533420](#)
- The master-password configuration is rejected if master-encryption-password (MEK) is not set. [PR1537251](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srpxpfe process might generate core files during the idpd process commit on SRX Series devices. [PR1521682](#)
- When adaptive threat profiling is configured within an IDP rule base and logging is enabled, on the vSRX instances the Packet Forwarding Engine process might stop and generate a core file. [PR1532737](#)

Platform and Infrastructure

- Configuration integrity mismatch error in vSRX3.0 running on Azure with key-vault integrated. [PR1551419](#)
- The pkid process runs at 100 percent when the device is unable to connect to a particular URL. [PR1560374](#)

VPNs

- The flowd process might stop in an IPsec VPN scenario. [PR1517262](#)

Resolved Issues: 20.3R1

Application Security

- Application Quality of Experience (AppQoE) system log shows best-path previous-interface value as “N/A” when deactivating DBG or the link. [PR1487056](#)
- When destination-path-group is deleted in the configuration and added again, the fc-id, dscp, fc name, and loss priority fields are reset. [PR1489948](#)
- The flow performance might be reduced in the Security Intelligence scenario. [PR1491682](#)

Intrusion Detection and Prevention (IDP)

- The IDP attack detection may not work in a specific situation. [PR1497340](#)

J-Web

- While creating a firewall policy rule, the list of available dynamic applications is empty in HA on the Select Dynamic Application page. [PR1490346](#)
- Infinite loading circle may be encountered via J-Web. [PR1493601](#)

Platform and Infrastructure

- The clock drift issue might cause control link failure of a vSRX cluster running on KVM hypervisor. [PR1496937](#)
- The vSRX may restart unexpectedly. [PR1479156](#)
- In vSRX3.0 on Azure with keyvault enabled, change in MEK results in deletion of certificates. [PR1513456](#)
- With CSO SD-WAN configuration loaded, flowd process generates core files while deleting the GRE IPsec configuration. [PR1513461](#)
- Changes to the configuration command for assigning more vCPUs to the Routing Engine. [PR1505724](#)
- On vSRX the interfaces might remain shut as the FPC faces issues while coming online after an upgrade attempt on the device. [PR1499092](#)
- When SSL proxy is enabled and if the vSRX runs out of memory, then the SSL proxy module might stop. [PR1505013](#)

Routing Policy and Firewall Filters

- Traffic might fail to hit policies if match dynamic-application and match source-end-user-profile options are configured under the same security policy name. [PR1505002](#)
- Junos OS upgrade may encounter failure in certain conditions when enabling ATP. [PR1519222](#)

Unified Threat Management (UTM)

- The source and destination IP or port fields were reversed for Content-Filtering and Anti-Virus logs. [PR1499327](#)

VPNs

- On vSRX3.0 instances, when ECMP routes are configured to load balance over multiple IPsec VPNs connected to a single multipoint tunnel interface, the traffic may not flow. [PR1438311](#)
- The flowd process might stop in IPsec VPN scenario. [PR1517262](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software Packages | 347](#)
- [Validating the OVA Image | 352](#)

This section contains information about how to upgrade Junos OS for vSRX using the CLI. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

You also can upgrade to Junos OS Release 20.3R3 for vSRX using J-Web (see [J-Web](#)) or the Junos Space Network Management Platform (see [Junos Space](#)).

Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Releases 17.4, 18.1, 18.2, 18.3, 18.4, 19.1, or 19.2 is supported.

The following limitations apply:

- Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Release 19.3 and higher is not supported. For upgrade between other combinations of Junos OS Releases in vSRX and vSRX 3.0, the general Junos OS upgrade policy applies.
- The file system mounted on /var usage must be below 14% of capacity.

Check this using the following command:

```
show system storage | match "/var$" /dev/vtbd1s1f
```

2.7G	82M	2.4G	3%	/var
------	-----	------	----	------

Using the **request system storage cleanup** command might help reach that percentage.

- The Junos OS upgrade image must be placed in the directory `/var/host-mnt/var/tmp/`. Use the **request system software add /var/host-mnt/var/tmp/<upgrade_image>**
- We recommend that you deploy a new vSRX virtual machine (VM) instead of performing a Junos OS upgrade. That also gives you the option to move from vSRX to the newer and more recommended vSRX 3.0.
- Ensure to back up valuable items such as configurations, license-keys, certificates, and other files that you would like to keep.

NOTE: For ESXi deployments, the firmware upgrade from Junos OS Release 15.1X49-Dxx to Junos OS releases 17.x, 18.x, or 19.x is not recommended if there are more than three network adapters on the 15.1X49-Dxx vSRX instance. If there are more than three network adapters and you want to upgrade, then we recommend that you either delete all the additional network adapters and add the network adapters after the upgrade or deploy a new vSRX instance on the targeted OS version.

Upgrading Software Packages

To upgrade the software using the CLI:

1. Download the **Junos OS Release 20.3R3 for vSRX .tgz** file from the [Juniper Networks website](#). Note the size of the software image.
2. Verify that you have enough free disk space on the vSRX instance to upload the new software image.

```
root@vsrx> show system storage
```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/vtbd0s1a	694M	433M	206M	68%	/
devfs	1.0K	1.0K	0B	100%	/dev
/dev/md0	1.3G	1.3G	0B	100%	/junos
/cf	694M	433M	206M	68%	/junos/cf
devfs	1.0K	1.0K	0B	100%	/junos/dev/
procfs	4.0K	4.0K	0B	100%	/proc
/dev/vtbd1s1e	302M	22K	278M	0%	/config
/dev/vtbd1s1f	2.7G	69M	2.4G	3%	/var
/dev/vtbd3s2	91M	782K	91M	1%	/var/host
/dev/md1	302M	1.9M	276M	1%	/mfs
/var/jail	2.7G	69M	2.4G	3%	/jail/var

/var/jails/rest-api	2.7G	69M	2.4G	3%	/web-api/var
/var/log	2.7G	69M	2.4G	3%	/jail/var/log
devfs	1.0K	1.0K	0B	100%	/jail/dev
192.168.1.1:/var/tmp/corefiles		4.5G	125M	4.1G	3%
/var/crash/corefiles					
192.168.1.1:/var/volatile	1.9G	4.0K	1.9G	0%	
/var/log/host					
192.168.1.1:/var/log	4.5G	125M	4.1G	3%	
/var/log/hostlogs					
192.168.1.1:/var/traffic-log	4.5G	125M	4.1G	3%	
/var/traffic-log					
192.168.1.1:/var/local	4.5G	125M	4.1G	3%	/var/db/host
192.168.1.1:/var/db/aamwd	4.5G	125M	4.1G	3%	
/var/db/aamwd					
192.168.1.1:/var/db/secinteld	4.5G	125M	4.1G	3%	
/var/db/secinteld					

3. Optionally, free up more disk space if needed to upload the image.

```

root@vsrx> request system storage cleanup
List of files to delete:
Size Date   Name
11B Jan 25 14:15 /var/jail/tmp/alarmd.ts
259.7K Jan 25 14:11 /var/log/hostlogs/vjunos0.log.1.gz
494B Jan 25 14:15 /var/log/interactive-commands.0.gz
20.4K Jan 25 14:15 /var/log/messages.0.gz
27B Jan 25 14:15 /var/log/wtmp.0.gz
27B Jan 25 14:14 /var/log/wtmp.1.gz
3027B Jan 25 14:13 /var/tmp/BSD.var.dist
0B Jan 25 14:14 /var/tmp/LOCK_FILE
666B Jan 25 14:14 /var/tmp/appidd_trace_debug
0B Jan 25 14:14 /var/tmp/eedebug_bin_file
34B Jan 25 14:14 /var/tmp/gksdchk.log
46B Jan 25 14:14 /var/tmp/kmdchk.log
57B Jan 25 14:14 /var/tmp/krt_rpf_filter.txt
42B Jan 25 14:13 /var/tmp/pfe_debug_commands
0B Jan 25 14:14 /var/tmp/pkg_cleanup.log.err
30B Jan 25 14:14 /var/tmp/policy_status
0B Jan 25 14:14 /var/tmp/rtsdb/if-rtsdb
Delete these files ? [yes,no] (no) yes

```

```
<
output omitted>
```

NOTE: If this command does not free up enough disk space, see [\[SRX\] Common and safe files to remove in order to increase available system storage](#) for details on safe files you can manually remove from vSRX to free up disk space.

4. Use FTP, SCP, or a similar utility to upload the Junos OS Release 20.3R3 for vSRX .tgz file to `/var/crash/corefiles/` on the local file system of your vSRX VM. For example:

```
root@vsrx> file copy ftp://username:prompt@ftp.hostname.net/pathname/
junos-vsrx-x86-64-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE.tgz /var/crash/corefiles/
```

5. From operational mode, install the software upgrade package.

```
root@vsrx> request system software add
/var/crash/corefiles/junos-vsrx-x86-64-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE.tgz
no-copy no-validate reboot
Verified junos-vsrx-x86-64-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING:      This package will load JUNOS 20.3 software.
WARNING:      It will save JUNOS configuration files, and SSH keys
WARNING:      (if configured), but erase all other files and information
WARNING:      stored on this machine. It will attempt to preserve dumps
WARNING:      and log files, but this can not be guaranteed. This is the
WARNING:      pre-installation stage and all the software is loaded when
WARNING:      you reboot the system.
Saving the config files ...
Pushing Junos image package to the host...
Installing
/var/tmp/install-media-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31
junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31
junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-linux.tgz
Setting up Junos host applications for installation ...
=====
```

```

Host OS upgrade is FORCED
Current Host OS version: 3.0.4
New Host OS version: 3.0.4
Min host OS version required for applications: 0.2.4
=====
Installing Host OS ...
upgrade_platform: -----
upgrade_platform: Parameters passed:
upgrade_platform: silent=0
upgrade_platform:
package=/var/tmp/junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-linux.tgz
upgrade_platform: clean install=0
upgrade_platform: clean upgrade=0
upgrade_platform: Need reboot after staging=0
upgrade_platform: -----
upgrade_platform:
upgrade_platform: Checking input
/var/tmp/junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-linux.tgz ...
upgrade_platform: Input package
/var/tmp/junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-linux.tgz is
valid.
upgrade_platform: Backing up boot assets..
cp: omitting directory '.'
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK
initrd.cpio.gz: OK
upgrade_platform: Checksum verified and OK...
/boot
upgrade_platform: Backup completed
upgrade_platform: Staging the upgrade package -
/var/tmp/junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-linux.tgz..
./
./bzImage-intel-x86-64.bin
./initramfs.cpio.gz
./upgrade_platform
./HOST_COMPAT_VERSION
./version.txt
./initrd.cpio.gz
./linux.checksum
./host-version
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK

```

```

upgrade_platform: Checksum verified and OK...
upgrade_platform: Staging of
/var/tmp/junos-srx-mr-vsrx-20.3-2021-1-25.0_RELEASE_20.3_THROTTLE-linux.tgz
completed
upgrade_platform: System need *REBOOT* to complete the upgrade
upgrade_platform: Run upgrade_platform with option -r | --rollback to rollback
the upgrade
Host OS upgrade staged. Reboot the system to complete installation!
WARNING:      A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:      'request system reboot' command when software installation is
WARNING:      complete. To abort the installation, do not reboot your system,
WARNING:      instead use the 'request system software rollback'
WARNING:      command as soon as this operation completes.
NOTICE: 'pending' set will be activated at next reboot...
Rebooting. Please wait ...
shutdown: [pid 13050]
Shutdown NOW!
*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY
Shutdown NOW!
System shutdown time has arrived\x07\x07

```

If no errors occur, Junos OS reboots automatically to complete the upgrade process. You have successfully upgraded to Junos OS Release 20.3R3 for vSRX.

NOTE: Starting in Junos OS Release 17.4R1, upon completion of the vSRX image upgrade, the original image is removed by default as part of the upgrade process.

6. Log in and use the **show version** command to verify the upgrade.

```

--- JUNOS 20.3-2021-1-25.0_RELEASE_20.3_THROTTLE Kernel 64-bit
JNPR-11.0-20171012.170745_fbsd-
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@:~ # cli
root> show version
Model: vsrx
Junos: 20.3-2021-1-25.0_RELEASE_20.3_THROTTLE
JUNOS OS Kernel 64-bit [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs [20171012.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20171012.170745_fbsd-builder_stable_11]

```

```

JUNOS OS time zone information [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20171012.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20171012.170745_fbsd-builder_stable_11]
JUNOS py extensions [20171017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20171017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20171012.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20171012.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities
[20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package
[20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx Data Plane Crypto Support
[20171017.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20171017.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20171017.110007_ssd-builder_release_174_throttle]

```

Validating the OVA Image

If you have downloaded a vSRX .ova image and need to validate it, see [Validating the vSRX .ova File for VMware](#).

Note that only .ova (VMware platform) vSRX images can be validated. The .qcow2 vSRX images for use with KVM cannot be validated the same way. File checksums for all software images are, however, available on the download page.

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the [High Availability User Guide](#).

For additional information about using ISSU on security devices, see the [Chassis Cluster User Guide for SRX Series Devices](#).

For information about ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\)](#) Web application.

Licensing

Starting in 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that we've developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see [Juniper Flex Program](#).

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. <https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home

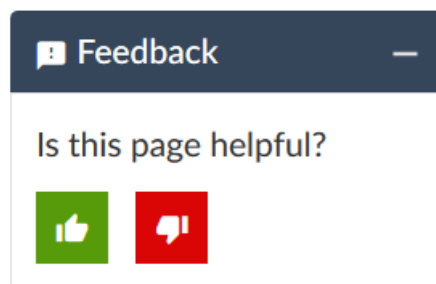
NOTE: To obtain information about the components that are supported on the devices, and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. apps.juniper.net/compliance/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

Revision History

25 November 2022—Revision 8, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

29 July 2022—Revision 7, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

24 March 2022—Revision 6, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

27 October 2021—Revision 5, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

16 September 2021—Revision 4, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

26 August 2021—Revision 3, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

15 July 2021—Revision 2, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

17 June 2021—Revision 1, Junos OS Release 20.3R3— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

22 April 2021—Revision 4, Junos OS Release 20.3R2— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

8 April 2021—Revision 3, Junos OS Release 20.3R2— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

11 March 2021—Revision 2, Junos OS Release 20.3R2— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

26 February 2021—Revision 1, Junos OS Release 20.3R2— ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

7 January 2021—Revision 7, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

2 December 2020—Revision 6, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

12 November 2020—Revision 5, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

4 November 2020—Revision 4, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

2 November 2020—Revision 3, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

8 October 2020—Revision 2, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

29 September 2020—Revision 1, Junos OS Release 20.3R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.