

# Release Notes

Published  
2022-06-30

## Junos<sup>®</sup> OS 20.3R1 Release Notes

### SUPPORTED ON

- ACX Series, cRPD, 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

### SOFTWARE HIGHLIGHTS

- Probe command to query the status of the probed interfaces (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)
- cSRX orchestration using Kubernetes (cSRX)
- Phone-home client (EX4300 Virtual Chassis)
- Enhancements to sessions over outbound HTTPS (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)
- New Juniper Secure Connect application for SRX Series and vSRX Next-Generation Firewalls (Juniper Secure Connect)
- SRv6 network programming in IS-IS (MX Series with MPC7E, MPC8E and MPC9E line cards)
- LDP Tunneling over Segment Routing Traffic Engineering (MX Series, PTX Series, and ACX5448 Routers)
- IP over IP next hop based tunneling (MX Series, PTX1000, PTX10000, and QFX10000)
- TCP authentication option (TCP-AO) for BGP and LDP connections (MX Series and PTX Series)
- Juniper Agile Licensing (QFX5120 and QFX5200)



- Seamless EVPN-VXLAN stitching (QFX10002-36Q, QFX10002-72Q, QFX10008, and QFX10016)
- Enhanced file-signing with Veriexec (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550)
- Wi-Fi Mini-PIM in High Availability (HA) cluster configuration (SRX Series)
- IKEv2 configuration payload improvements on new IKED platforms (SRX5000 line of devices with SPC3 and vSRX)
- Scaling vSRX using Azure Load Balancer and Virtual Machine Scale Sets (vSRX and vSRX 3.0)



## IN FOCUS GUIDE

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

## Day One+

- Use this [new setup guide](#) to get your Junos OS up and running in three quick steps.



# Release Notes: Junos<sup>®</sup> OS Release 20.3R1 for the 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

30 June 2022

<b>Contents</b>	<b>Introduction   16</b>
	<b>Junos OS Release Notes for ACX Series   16</b>
	<b>What's New   17</b>
	Hardware   18
	EVPNS   18
	Multicast   20
	Network Management and Monitoring   20
	Routing Protocols   20
	Segment Routing   21
	<b>What's Changed   22</b>
	General Routing   22
	Junos OS, XML, API, and Scripting   23
	Routing Protocols   23
	<b>Known Limitations   24</b>
	General Routing   24



**Open Issues | 25**

- Forwarding and Sampling | 26**
- General Routing | 26**
- Interfaces and Chassis | 27**
- Platform and Infrastructure | 27**
- Virtual Chassis | 27**

**Resolved Issues | 28**

- General Routing | 28**
- Interfaces and Chassis | 31**
- MPLS | 31**
- Routing Protocols | 31**
- VPNs | 31**

**Documentation Updates | 31****Migration, Upgrade, and Downgrade Instructions | 32**

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

**Junos OS Release Notes for cRPD | 33****What's New | 34**

- Supported Features on cRPD | 35**

**Known Limitations | 35**

- MPLS | 36**

**Open Issues | 36**

- Routing Protocols | 36**

**Resolved Issues | 36**

- Routing Protocols | 36**

**Junos OS Release Notes for cSRX | 37****What's New | 37**

- Installation and Upgrade | 38**

**What's Changed | 38**

- Download Juniper Signature Pack on cSRX | 39**

**Known Limitations | 40**

- Installation and Upgrade | 41**

**Open Issues | 41**

- Interfaces and Chassis | 42**



## Junos OS Release Notes for EX Series | 42

### What's New | 43

- Hardware | 43
- Class of Service (CoS) | 49
- EVPN | 49
- Junos OS XML, API, and Scripting | 49
- Junos Telemetry Interface | 49
- MPLS | 51
- Network Management and Monitoring | 52
- open-config | 53
- Routing Policy and Firewall Filters | 53
- Software Installation and Upgrade | 53

### What's Changed | 54

- Class of Service (CoS) | 55
- Junos OS, XML, API, and Scripting | 55
- Platform and Infrastructure | 55
- Routing Protocols | 56
- Subscriber Management and Services | 56

### Known Limitations | 56

- EVPN | 57
- General Routing | 57
- Infrastructure | 58

### Open Issues | 58

- EVPN | 59
- General Routing | 59
- Infrastructure | 59
- Layer 2 Features | 60
- Network Management and Monitoring | 60
- Platform and Infrastructure | 60

### Resolved Issues | 60

- Authentication and Access Control | 61
- EVPN | 61
- General Routing | 61
- Infrastructure | 62



Interfaces and Chassis	63
Layer 2 Ethernet Services	63
Layer 2 Features	63
MPLS	63
Platform and Infrastructure	63
Routing Protocols	63
User Interface and Configuration	64
Virtual Chassis	64
Documentation Updates	64
Migration, Upgrade, and Downgrade Instructions	65
Upgrade and Downgrade Support Policy for Junos OS Releases	65
Junos OS Release Notes for JRR Series	66
What's New	66
Routing Protocols	67
What's Changed	68
General Routing	68
Known Limitations	68
Open Issues	69
Resolved Issues	69
Documentation Updates	70
Migration, Upgrade, and Downgrade Instructions	70
Upgrade and Downgrade Support Policy for Junos OS Releases	70
Junos OS Release Notes for Juniper Secure Connect	71
What's New	72
Juniper Secure Connect	72
Open Issues	73
Platform and Infrastructure	73
VPNs	74
Junos OS Release Notes for Junos Fusion for Enterprise	74
What's New	74
What's Changed	75
Known Limitations	75
Open Issues	76
Resolved Issues	76



Documentation Updates | 77

Migration, Upgrade, and Downgrade Instructions | 77

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 78

Upgrading an Aggregation Device with Redundant Routing Engines | 79

Preparing the Switch for Satellite Device Conversion | 80

Converting a Satellite Device to a Standalone Switch | 81

Upgrade and Downgrade Support Policy for Junos OS Releases | 81

Downgrading Junos OS | 82

Junos OS Release Notes for Junos Fusion Provider Edge | 83

What's New | 83

What's Changed | 84

General Routing | 84

Known Limitations | 84

Open Issues | 85

Resolved Issues | 85

Junos Fusion Provider Edge | 86

Documentation Updates | 86

Migration, Upgrade, and Downgrade Instructions | 87

Basic Procedure for Upgrading an Aggregation Device | 87

Upgrading an Aggregation Device with Redundant Routing Engines | 90

Preparing the Switch for Satellite Device Conversion | 90

Converting a Satellite Device to a Standalone Device | 92

Upgrading an Aggregation Device | 94

Upgrade and Downgrade Support Policy for Junos OS Releases | 94

Downgrading from Junos OS Release 20.1 | 95

Junos OS Release Notes for MX Series | 95

What's New | 96

Hardware | 97

Authentication, Authorization, and Accounting | 100

Class of Service (CoS) | 100

EVPN | 100

High Availability (HA) and Resiliency | 101

Interfaces and Chassis | 101

IP Tunneling | 102



Juniper Extension Toolkit	103
Junos OS XML, API, and Scripting	104
Junos Telemetry Interface	104
Layer 2 Features	111
Layer 2 VPN	111
Layer 3 Features	111
MPLS	111
Multicast	112
Network Management and Monitoring	112
Next Gen Services	113
Port Security	114
Routing Protocols	115
Segment Routing	117
Services Applications	117
Software Defined Networking (SDN)	118
System Management	119
What's Changed	120
EVPN	121
General Routing	121
High Availability (HA) and Resiliency	122
Infrastructure	122
Interfaces and Chassis	122
Junos OS, XML, API, and Scripting	122
Routing Protocols	123
Services Applications	123
Subscriber Management and Services	123
Known Limitations	124
EVPN	124
General Routing	124
Interfaces and Chassis	125
MPLS	125
Network Management and Monitoring	125
Platform and Infrastructure	126
Routing Protocols	127



Subscriber Management and Services | 127

Open Issues | 127

Class of Service (CoS) | 128

EVPN | 128

Forwarding and Sampling | 128

General Routing | 129

Infrastructure | 132

Interfaces and Chassis | 132

Intrusion Detection and Prevention (IDP) | 133

Layer 2 Ethernet Services | 133

MPLS | 133

Network Management and Monitoring | 133

Platform and Infrastructure | 133

Routing Protocols | 134

Subscriber Access Management | 134

User Interface and Configuration | 134

Resolved Issues | 135

Application Layer Gateways (ALGs) | 136

Class of Service (CoS) | 136

EVPN | 136

Forwarding and Sampling | 137

General Routing | 137

Infrastructure | 143

Interfaces and Chassis | 143

Intrusion Detection and Prevention (IDP) | 144

J-Web | 144

Juniper Extension Toolkit (JET) | 144

Junos Fusion Provider Edge | 144

Layer 2 Ethernet Services | 144

MPLS | 145

Network Management and Monitoring | 146

Platform and Infrastructure | 146

Routing Protocols | 147

Services Applications | 148



Subscriber Access Management	149
Subscriber Management and Services	149
User Interface and Configuration	149
VPNs	149
Documentation Updates	150
Migration, Upgrade, and Downgrade Instructions	150
Basic Procedure for Upgrading to Release 20.3R1	151
Procedure to Upgrade to FreeBSD 11.x-Based Junos OS	151
Procedure to Upgrade to FreeBSD 6.x-Based Junos OS	154
Upgrade and Downgrade Support Policy for Junos OS Releases	156
Upgrading a Router with Redundant Routing Engines	156
Downgrading from Release 20.3R1	157
Junos OS Release Notes for NFX Series	157
What's New	158
Application Security	158
Wireless WAN	159
What's Changed	159
What's Changed in Release 20.3R1	160
Known Limitations	160
Interfaces	160
Open Issues	161
High Availability	161
Interfaces	161
Platform and Infrastructure	162
Virtual Network Functions (VNFs)	162
Resolved Issues	163
Application Security	163
High Availability	163
Interfaces	163
Platform and Infrastructure	164
Documentation Updates	164
Migration, Upgrade, and Downgrade Instructions	165
Upgrade and Downgrade Support Policy for Junos OS Releases	165
Basic Procedure for Upgrading to Release 20.3	165



## Junos OS Release Notes for PTX Series | 167

### What's New | 168

- Hardware | 168
- Authentication, Authorization, and Accounting | 169
- IP Tunneling | 169
- Juniper Extension Toolkit (JET) | 170
- Junos OS XML, API, and Scripting | 171
- Junos Telemetry Interface | 171
- Layer 3 Features | 175
- MPLS | 175
- Network Management and Monitoring | 175
- Port Security | 176
- Routing Protocols | 177
- Segment Routing | 178
- Services Applications | 178
- Software Defined Networking (SDN) | 179

### What's Changed | 180

- Class of Service (CoS) | 181
- General Routing | 181
- High Availability (HA) and Resiliency | 182
- Junos OS XML, API, and Scripting | 182
- Juniper Extension Toolkit (JET) | 182
- MPLS | 182
- Routing Protocols | 183
- System Management | 183

### Known Limitations | 183

- General Routing | 184
- MPLS | 185
- Routing Protocols | 185

### Open Issues | 186

- General Routing | 186
- MPLS | 187
- Routing Protocols | 187



**Resolved Issues | 187****General Routing | 188****Interfaces and Chassis | 189****MPLS | 189****Network Management and Monitoring | 189****Routing Protocols | 189****Documentation Updates | 190****Migration, Upgrade, and Downgrade Instructions | 190****Basic Procedure for Upgrading to Release 20.3 | 191****Upgrade and Downgrade Support Policy for Junos OS Releases | 193****Upgrading a Router with Redundant Routing Engines | 194****Junos OS Release Notes for QFX Series | 195****What's New | 195****Hardware | 197****EVPN | 198****High Availability (HA) and Resiliency | 201****IP Tunneling | 201****Juniper Extension Toolkit (JET) | 202****Junos OS XML, API, and Scripting | 203****Junos Telemetry Interface | 203****Layer 3 Features | 204****MPLS | 204****Network Management and Monitoring | 204****Routing Policy and Firewall Filters | 206****Routing Protocols | 206****Security | 207****Services Applications | 207****Software Defined Networking (SDN) | 207****Software Licensing | 207****Virtual Chassis | 208****What's Changed | 209****Class of Service (CoS) | 210****General Routing | 210****High Availability (HA) and Resiliency | 211**



Interfaces and Chassis	211
Junos OS XML, API, and Scripting	211
Routing Protocols	211
Known Limitations	212
Layer 2 Ethernet Services	212
Platform and Infrastructure	213
Routing Protocols	213
Open Issues	214
High Availability (HA) and Resiliency	214
Layer 2 Features	214
Platform and Infrastructure	215
Routing Protocols	217
Virtual Chassis	218
Resolved Issues	218
Class of Service (CoS)	219
EVPN	219
Infrastructure	219
Interfaces and Chassis	219
Layer 2 Features	220
Layer 2 Ethernet Services	220
MPLS	220
Platform and Infrastructure	220
Routing Protocols	222
User Interface and Configuration	223
Documentation Updates	223
Migration, Upgrade, and Downgrade Instructions	224
Upgrading Software on QFX Series Switches	224
Installing the Software on QFX10002-60C Switches	227
Installing the Software on QFX10002 Switches	227
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	228
Installing the Software on QFX10008 and QFX10016 Switches	230
Performing a Unified ISSU	234
Preparing the Switch for Software Installation	235



Upgrading the Software Using Unified ISSU | 235

Upgrade and Downgrade Support Policy for Junos OS Releases | 237

Junos OS Release Notes for SRX Series | 238

What's New | 239

Application Security | 240

Authentication and Access Control | 240

Chassis Clustering | 240

Flow-Based and Packet-Based Processing | 241

Installation and Upgrade | 241

Interfaces and Chassis | 241

Intrusion Detection and Prevention (IDP) | 242

Junos Telemetry Interface | 242

Junos OS XML API and Scripting | 243

J-Web | 243

Layer 2 Features | 245

Logical Systems and Tenant Systems | 245

Network Management and Monitoring | 245

Routing and Forwarding Options | 247

Security | 247

Unified Threat Management (UTM) | 247

VPNs | 247

What's Changed | 249

Authentication and Access Control | 249

Junos OS XML API and Scripting | 249

J-Web | 250

Network Address Translation (NAT) | 250

System Logs | 250

Known Limitations | 251

Flow-Based and Packet-Based Processing | 251

J-Web | 251

VPNs | 252

Open Issues | 252

J-Web | 253

VPNs | 253



**Resolved Issues | 254****Application Security | 254****Chassis Clustering | 254****Flow-Based and Packet-Based Processing | 254****Infrastructure | 255****Interfaces and Chassis | 255****Intrusion Detection and Prevention (IDP) | 256****J-Web | 256****MPLS | 256****Network Address Translation (NAT) | 256****Platform and Infrastructure | 256****Routing Policy and Firewall Filters | 257****Routing Protocols | 257****VPNs | 257****Documentation Updates | 258****Migration, Upgrade, and Downgrade Instructions | 258****Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | 258****Junos OS Release Notes for vMX | 259****What's New | 260****Installation and Upgrade | 260****Juniper Extension Toolkit (JET) | 260****System Management | 261****Open Issues | 261****General Routing | 262****Resolved Issues | 262****Platform and Infrastructure | 262****Licensing | 262****Upgrade Instructions | 263****Junos OS Release Notes for vRR | 263****What's New | 264****Routing Protocols | 264****What's Changed | 265**



Known Limitations	265
Routing Protocols	266
Open Issues	266
Resolved Issues	266
Junos OS Release Notes for vSRX	267
What's New	267
Interfaces and Chassis	268
Juniper ATP Cloud	268
Junos Telemetry Interface	268
Management	270
Performance and Scaling	270
VPNs	270
Known Limitations	271
Intrusion Detection and Prevention (IDP)	272
J-Web	272
Open Issues	272
J-Web	273
User Access and Authentication	273
Resolved Issues	273
Application Security	273
Intrusion Detection and Prevention (IDP)	273
J-Web	273
Platform and Infrastructure	274
Routing Policy and Firewall Filters	274
Unified Threat Management (UTM)	274
VPNs	274
Migration, Upgrade, and Downgrade Instructions	275
Upgrading Software Packages	276
Validating the OVA Image	281
Upgrading Using ISSU	281
Licensing	282
Compliance Advisor	282
Finding More Information	282
Documentation Feedback	283



Requesting Technical Support | 284

Self-Help Online Tools and Resources | 284

Creating a Service Request with JTAC | 285

Revision History | 285



# Introduction

Junos OS runs on the following Juniper Networks<sup>®</sup> hardware: ACX Series, cRPD, 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 20.3R1 for the ACX Series, Containerized Routing Protocol Process (cRPD), 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](mailto:techpubs-comments@juniper.net).
- **Important Information:**
  - [Upgrading Using ISSU on page 281](#)
  - [Licensing on page 282](#)
  - [Compliance Advisor on page 282](#)
  - [Finding More Information on page 282](#)
  - [Documentation Feedback on page 283](#)
  - [Requesting Technical Support on page 284](#)

## Junos OS Release Notes for ACX Series

### IN THIS SECTION

- [What's New | 17](#)
- [What's Changed | 22](#)
- [Known Limitations | 24](#)
- [Open Issues | 25](#)
- [Resolved Issues | 28](#)



- Documentation Updates | 31
- Migration, Upgrade, and Downgrade Instructions | 32

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- Hardware | 18
- EVPNS | 18
- Multicast | 20
- Network Management and Monitoring | 20
- Routing Protocols | 20
- Segment Routing | 21

This section describes the new features or enhancements to existing features in Junos OS Release 20.3R1 for the ACX Series.



## 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"> <li>• Support for hierarchical CoS. Support for up to three levels of hierarchical scheduling (physical interfaces, logical interfaces, and queues). [See <a href="#">Hierarchical Class of Service</a>.]</li> </ul>
Network management and monitoring	<ul style="list-style-type: none"> <li>• 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 <a href="#">Accounting Options Configurations</a>.]</li> </ul>
Timing and synchronization	<ul style="list-style-type: none"> <li>• Support for logs, alarms, counters, and SNMP traps for Precision Time Protocol (PTP)/ Synchronous Ethernet. [See <a href="#">Enterprise-Specific SNMP Traps Supported by Junos OS</a> and <a href="#">show chassis alarms</a>.]</li> <li>• Support for the g.8275.1 profile on the ACX710. [See <a href="#">Assisted Partial Timing Support</a>.]</li> <li>• Support for PTP G.8275.1. Use PTP profile-type g.8275.1 to enable the G.8275.1 profile. [See <a href="#">Profile Type</a>.]</li> <li>• Support for limited images on the ACX710. [See <a href="#">Software Installation and Upgrade Overview</a>.]</li> <li>• Support for RIP version 1, RIP version 2, and RIP next generation (RIPng) on PTX10008 routers. [See <a href="#">RIP and RIPng Overview</a>.]</li> </ul>

## 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 | 22](#)



Known Limitations	24
Open Issues	25
Resolved Issues	28
Documentation Updates	31
Migration, Upgrade, and Downgrade Instructions	32

## What's Changed

### IN THIS SECTION

- General Routing | 22
- Junos OS, XML, API, and Scripting | 23
- Routing Protocols | 23

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 20.3R1 for the ACX Series routers.

### General Routing

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



## 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.
- **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 | 17](#)

[Known Limitations | 24](#)

[Open Issues | 25](#)

[Resolved Issues | 28](#)

[Documentation Updates | 31](#)

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



## Known Limitations

### IN THIS SECTION

- [General Routing | 24](#)

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](#)
- On the ACX710 router, traffic loss is beyond the tolerance limit of 200 ms during convergence. [PR1499965](#)
- Not able to scale BFD to 1024 sessions with IPv4 and IPv6. [PR1502170](#)
- Satellites do not track intermittently with GPS-only constellation. [PR1505325](#)
- Unexpected delay counter values are seen in the output of the **show ptp statistics detail** command when the upstream master 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 master 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](#)

SEE ALSO

<a href="#">What's New   17</a>
<a href="#">What's Changed   22</a>
<a href="#">Open Issues   25</a>
<a href="#">Resolved Issues   28</a>
<a href="#">Documentation Updates   31</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   32</a>

## Open Issues

IN THIS SECTION

- [Forwarding and Sampling | 26](#)
- [General Routing | 26](#)
- [Interfaces and Chassis | 27](#)
- [Platform and Infrastructure | 27](#)
- [Virtual Chassis | 27](#)

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.

## Forwarding and Sampling

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

## General Routing

- A **jnxIfOtnOperState** trap notification is sent for all ot- interfaces. [PR1406758](#)
- The DHCP clients are not able to scale to 96,000. [PR1432849](#)
- Protocols get forwarded when you use the nonexisting SSM map source address in IGMPv3 instead of pruning. [PR1435648](#)
- Drop profile max threshold might not reach its full limit when the packet size is other than 1000 bytes. [PR1448418](#)
- The vpls-oam sessions are detected with error (RDI sent by some MEP) after changing VLANs. [PR1478346](#)
- On the ACX710 routers, traffic loss is observed after changing the ALT port cost for RSTP. [PR1482566](#)
- Mirroring does not work in Junos OS Release 19.4R2. [PR1491789](#)
- In some scenarios of port disabling and restart of clock recovery process followed by port enabling might not lock until it is reconfigured. [PR1505405](#)
- EXP marking is not as expected for the Layer 2 CKT, Layer 3 VPN, and Layer 2 VPN traffics. [PR1509627](#)
- Transit traffic MPLS EXP bits are wrongly marked as per queue-num and not as per rewrite-rule configuration. [PR1509635](#)
- The APTS mode of operation is not supported in G.8275.1. [PR1525918](#)
- The IPv6 BFD sessions flap when configured below 100 ms flaps. [PR1456237](#)
- On the ACX5448 router, the EXP rewrite for the Layer 3 VPN sends all traffic with incorrect EXP. [PR1500928](#)
- FEC field is not displayed when the interface is down. [PR1530755](#)
- The syslog error messages related to **ACX\_DFW\_CFG\_FAILED** are observed. [PR1490940](#)
- The **show class-of-service routing-instance** command does not show the configured classifier. [PR1531413](#)
- The clksyncd process generates core file. [PR1537107](#)
- Management Ethernet link down alarm is observed while verifying the system alarms in the Virtual Chassis setup. [PR1538674](#)
- In the ACX5048 router, **queue-counters-trans-bytes-rate** are more than expected while configuring the IFD and logical interface shaping with the transmit rate and scheduler-map. [PR1538934](#)



- The following unexpected error message is observed while deleting the remote stream 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: clksync_mimic_delete_clock_entry`. [PR1539953](#)
- On the ACX710 router, cTE should be tuned closer to 2Way cTE. [PR1527347](#)

## Interfaces and Chassis

- When reboot is issued to the standby MC-LAG node, one-time traffic hit of active path traffic is seen and later when this node comes up, the MC-LAG active and standby roles get changed to the other device. [PR1505841](#)

## Platform and Infrastructure

- The CFM remote MEP does not come up after configuration or if it remains in the **Start** state. [PR1460555](#)

## Virtual Chassis

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

## SEE ALSO

[What's New | 17](#)

[What's Changed | 22](#)

[Known Limitations | 24](#)

[Resolved Issues | 28](#)

[Documentation Updates | 31](#)

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



## Resolved Issues

### IN THIS SECTION

- [General Routing | 28](#)
- [Interfaces and Chassis | 31](#)
- [MPLS | 31](#)
- [Routing Protocols | 31](#)
- [VPNs | 31](#)

This section lists the issues fixed in Junos OS Release 20.3R1 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

- Policer discarded count is shown incorrectly to the enq count of the interface queue, but the traffic behavior is as expected. [PR1414887](#)
- The **gigether-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 master 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 DEI 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 LACP 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 | 17](#)

[What's Changed | 22](#)

[Known Limitations | 24](#)

[Open Issues | 25](#)

[Documentation Updates | 31](#)

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

## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for ACX Series routers.



## SEE ALSO

[What's New | 17](#)[What's Changed | 22](#)[Open Issues | 25](#)[Known Limitations | 24](#)[Resolved Issues | 28](#)[Migration, Upgrade, and Downgrade Instructions | 32](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 32](#)

This section contains the upgrade and downgrade support policy for Junos OS for 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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.

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



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

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

SEE ALSO

<a href="#">What's New   17</a>
<a href="#">What's Changed   22</a>
<a href="#">Known Limitations   24</a>
<a href="#">Open Issues   25</a>
<a href="#">Resolved Issues   28</a>
<a href="#">Documentation Updates   31</a>

# Junos OS Release Notes for cRPD

IN THIS SECTION

- [What's New | 34](#)
- [Known Limitations | 35](#)
- [Open Issues | 36](#)
- [Resolved Issues | 36](#)

These release notes accompany Junos OS Release 20.3R1 for the containerized routing protocol process (cRPD) container. 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](https://www.juniper.net/documentation/product/en_US/junos-os).



## What's New

### IN THIS SECTION

- [Supported Features on cRPD | 35](#)

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



Supported Features on cRPD

- In Junos OS Release 20.3R1, you can configure the following supported routing features on cRPD:

Supported Feature	Description
BGP flowspec	Support for the BGP flow specification method to prevent denial-of-service attacks on the cRPD environment.  [See <a href="#">Understanding BGP Flow Routes for Traffic Filtering.</a> ]
EVPN-VPWS	Support for EVPN-VPWS to provide VPWS with EVPN signaling mechanisms.  [See <a href="#">Overview of VPWS with EVPN Signaling Mechanisms.</a> ]
EVPN TYPE 5 with MPLS	Support for EVPN Type 5 for EVPN/MPLS.  [See <a href="#">EVPN Type-5 Route with MPLS encapsulation for EVPN-MPLS.</a> ]
Segment routing	Segment routing support for OSPF and IS-IS protocols to provide basic functionality with Source Packet Routing in Networking (SPRING).  [See <a href="#">Understanding Source Packet Routing in Networking.</a> ]
Layer 2 VPN	Support for Layer 2 circuit to provide Layer 2 VPN and VPWS with LDP signaling.  [See <a href="#">Configuring Ethernet over MPLS (L2 Circuit).</a> ]
MPLS	Support for MPLS to provide LDP signaling protocol configuration with the control plane functionality.  [See <a href="#">Understanding the LDP Signaling Protocol.</a> ]

Known Limitations

IN THIS SECTION

- [MPLS | 36](#)

Learn about known limitations in this release for cRPD.



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

## MPLS

- Linux kernel not installing IPv6 route on ingress when IPv6 tunneling is enabled. [PR1504466](#)

## Open Issues

### IN THIS SECTION

- [Routing Protocols | 36](#)

Learn about open issues in this release for cRPD.

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

### Routing Protocols

- Linux version Linux crpd 4.9.0-9-2-amd64 #1 SMP Debian 4.9.168-1+deb9u3 (2015-12-19) x86\_64 x86\_64 GNU/Linux has a bug with adding IPv6 routes over a v4-mapped-v6 gateway. [PR1510811](#)

## Resolved Issues

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

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

### Routing Protocols

- Assigning a logical interface to a VRF instance may not always work. [PR1476669](#)
- BGP TCP MD5 authentication support is now available. [PR1514393](#)



# Junos OS Release Notes for cSRX

## IN THIS SECTION

- [What's New | 37](#)
- [What's Changed | 38](#)
- [Known Limitations | 40](#)
- [Open Issues | 41](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- [Installation and Upgrade | 38](#)

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



## 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

- [Download Juniper Signature Pack on cSRX](#) | 39

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



## 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

### IN THIS SECTION

- [Installation and Upgrade | 41](#)

Learn about known limitations in this release for cSRX.



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

## Installation and Upgrade

- Deploying cSRX with the environment variable CSRX\_SIZE to modify the size is not currently supported. Only the default size, large is supported. [PR1527636](#)

## Open Issues

### IN THIS SECTION

- [Interfaces and Chassis](#) | 42

This section lists the known issues in hardware and software for cSRX in Junos OS Release 20.3R1.

Learn about open issues in this release for cSRX.

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



## Interfaces and Chassis

- In case of support for multiple networks under K8s deployment, the cluster network interface is created with different ifindex orders by different Container Network Interface (CNI) (Flannel and Weave), this leads cSRX to bind different interfaces as the management interface.

As a workaround, set the environment variable `CSRX_MGMT_PORT_REORDER` to yes to bind the first interface as the management port, and set it to no to bind the last interface as the management port.

[PR1509702](#)

# Junos OS Release Notes for EX Series

## IN THIS SECTION

- [What's New | 43](#)
- [What's Changed | 54](#)
- [Known Limitations | 56](#)
- [Open Issues | 58](#)
- [Resolved Issues | 60](#)
- [Documentation Updates | 64](#)
- [Migration, Upgrade, and Downgrade Instructions | 65](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).



## What's New

### IN THIS SECTION

- [Hardware](#) | [43](#)
- [Class of Service \(CoS\)](#) | [49](#)
- [EVPN](#) | [49](#)
- [Junos OS XML, API, and Scripting](#) | [49](#)
- [Junos Telemetry Interface](#) | [49](#)
- [MPLS](#) | [51](#)
- [Network Management and Monitoring](#) | [52](#)
- [open-config](#) | [53](#)
- [Routing Policy and Firewall Filters](#) | [53](#)
- [Software Installation and Upgrade](#) | [53](#)

Learn about new features introduced in this release for EX Series Switches.

**NOTE:** The following EX Series switches are supported in Release 20.3R1: EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

## 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 44](#) 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"> <li>Support for COS features (classifiers, rewrites, port queuing, and L3) on EX9204, EX9208, and EX9214 switches. [See <a href="#">Understanding Junos OS CoS Components for EX Series Switches</a>.]</li> </ul>
EVPN	<ul style="list-style-type: none"> <li>Support for EVPN-MPLS singlehoming. This feature supports single-homed devices on an EVPN-MPLS network. [See <a href="#">Introduction to EVPN Multihoming</a>.]</li> <li>Support for NDP and Proxy ARP. Junos OS supports proxy Address Resolution Protocol (ARP) and Network Discovery Protocol (NDP).</li> </ul>



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

Feature	Description
Firewalls and policers	<ul style="list-style-type: none"> <li>• Support for CCC and Layer 3 firewall forwarding. [See <a href="#">CCC Overview</a>.]</li> <li>• Support for advanced Layer 2 features:             <ul style="list-style-type: none"> <li>• Firewall filters for Layer 2 and MAC filters. [See <a href="#">Layer 2 Forwarding Tables</a>.]</li> <li>• Layer 2 firewall forwarding support. A firewall filter specifies required traffic and directs it to the mirror. [See <a href="#">Understanding Port Mirroring and Analyzers</a>.]</li> </ul> </li> <li>• 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"> <li>• Input and output filter chains</li> <li>• Actions, including policy-map filters, do-not-fragment, and prefix</li> <li>• Layer 2 policers</li> <li>• Policer overhead adjustment</li> <li>• Hierarchical policers</li> <li>• Shared bandwidth</li> <li>• Percentages</li> <li>• Logical interfaces</li> </ul> </li> </ul> <p>[See <a href="#">Traffic Policer Types</a>.]</p>
Junos telemetry interface	<ul style="list-style-type: none"> <li>• 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"> <li>• <code>/junos/system/cmerror/configuration/</code></li> <li>• <code>/junos/system/cmerror/counters/</code></li> <li>• <code>/junos/system/linecard/environment/</code></li> <li>• <code>/junos/system/linecard/optics/</code></li> <li>• <code>/junos/system/linecard/optics/optics-diag[if-name =]</code></li> <li>• <code>/junos/system/linecard/optics/optics-diag/if-name</code></li> <li>• <code>/junos/system/linecard/optics/optics-diag/snmp-if-index</code></li> <li>• <code>/junos/system/linecard/optics/lane[lane_number=]/</code></li> </ul> </li> </ul> <p>[See <a href="#">Guidelines for gRPC Sensors (Junos Telemetry Interface)</a>.]</p>



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

Feature	Description
Layer 2 features	<ul style="list-style-type: none"> <li>• Support for Layer 3 VPN, Layer 2 VPN and Layer 2 circuits. [See <a href="#">Introduction to Configuring Layer 3 VPNs</a>, <a href="#">Layer 2 VPNs Configuration Overview</a> and <a href="#">Layer 2 Circuits Configuration Overview</a>.]</li> <li>• 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 <a href="#">Understanding Layer 2 Bridge Domains</a> and <a href="#">Learning and Forwarding</a>.]</li> <li>• Support for IRB, VLAN handling, and Q-in-Q tunneling. [See <a href="#">Integrated Routing and Bridging</a>, <a href="#">Understanding Bridging and VLANs on Switches</a> and <a href="#">Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation</a>.]</li> <li>• Support for VPLS. [See <a href="#">Introduction to Configuring VPLS</a>.]</li> </ul>
Layer 3 features	<ul style="list-style-type: none"> <li>• Support for Layer 3 forwarding. Junos OS supports the following Layer 3 features on the EX9200-15C: <ul style="list-style-type: none"> <li>• BGP (Multipath/v4-v6 labelled unicast)</li> <li>• Bidirectional Forwarding Detection (excluding micro BFD and BFD sessions with authentication)</li> <li>• IPv4 (forwarding and options)</li> <li>• IPv6 (forwarding and route accounting)</li> <li>• Load balancing (ECMP and FRR)</li> <li>• L2VPN, CCC, and L2 Circuit</li> <li>• 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.</li> <li>• OSPF (node-link-protection and node-link-degradation)</li> <li>• 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)</li> <li>• Routing Instance Logical System VRF</li> <li>• Tunnel (Generic Routing Encapsulation (GRE), Logical Tunnel (LT), and Virtual Tunnel (VT))</li> </ul> </li> </ul>



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

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



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

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>• Port mirroring support for families <b>inet</b>, <b>inet6</b>, and <b>ethernet-switching</b>, configured at the <b>[edit forwarding-options port-mirroring]</b> hierarchy level. [See <a href="#">Understanding Port Mirroring and Analyzers</a>.]</li> <li>• 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"> <li>• Discovery and link monitoring</li> <li>• Distributed LFM</li> <li>• Remote fault detection and remote loopback</li> </ul> [See <a href="#">OAM Link Fault Management</a>.]</li> <li>• Support for Junos OS management and software features on the EX9200-15C: <ul style="list-style-type: none"> <li>• 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</li> </ul> </li> <li>• Support for hyper mode and non hyper mode features. [See <a href="#">Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches</a>.]</li> </ul>
Port security	<ul style="list-style-type: none"> <li>• Provides support for MACsec on ports at these speeds 10G, 25G, 40G, and 100G. [See <a href="#">Understanding Media Access Control Security (MACsec)</a>.]</li> </ul>
Services applications	<ul style="list-style-type: none"> <li>• 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 <a href="#">Configuring Flow Aggregation on MX, M, vMX and T Series Routers, EX9200 Switches, and NFX250 to Use Version 9 Flow Templates</a>.]</li> </ul>
System management	<ul style="list-style-type: none"> <li>• Support for the Display Common Language Equipment Identifier (CLEI) barcode and model number for orderable field-replaceable units (FRUs). [See <a href="#">show chassis hardware</a>.]</li> </ul>

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) 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 \(Junos Telemetry Interface\)](#).]

## 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

<a href="#">What's Changed</a>	<a href="#">54</a>
<a href="#">Known Limitations</a>	<a href="#">56</a>
<a href="#">Open Issues</a>	<a href="#">58</a>
<a href="#">Resolved Issues</a>	<a href="#">60</a>
<a href="#">Documentation Updates</a>	<a href="#">64</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">65</a>

## What's Changed

IN THIS SECTION

- [Class of Service \(CoS\)](#) | [55](#)
- [Junos OS, XML, API, and Scripting](#) | [55](#)
- [Platform and Infrastructure](#) | [55](#)
- [Routing Protocols](#) | [56](#)
- [Subscriber Management and Services](#) | [56](#)

Learn about what changed in this release for EX Series Switches in Junos OS 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](#).]

## Platform and Infrastructure

- **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**.



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.

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

<a href="#">What's New   43</a>
<a href="#">Known Limitations   56</a>
<a href="#">Open Issues   58</a>
<a href="#">Resolved Issues   60</a>
<a href="#">Documentation Updates   64</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   65</a>

Known Limitations

IN THIS SECTION

- [EVPN | 57](#)
- [General Routing | 57](#)
- [Infrastructure | 58](#)



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

## EVPN

- Partial traffic loss is seen with a single link between the leaf devices. [PR1480847](#)

## General Routing

- Channels show false up even when peer end is down and different speed configured; LED also shows green. [PR1530061](#)
- The LED behavior of the following DAC/AOC/LX4 optics-vendor part number is different from the standard LED behavior of the EX9208 platform. DAC cables 1. JNP-100G-DAC-5M 740-061002 1P1C43A4212RJ QSFP28-100G-CU5M JUNIPER-AMPHENOL 2. JNP-100G-DAC-5M 740-061002 1RC434030JN QSFP28-100G-CU5M JUNIPER-LEONI 3. JNP-100G-DAC-1M 740-061000 1P1C40A4192G3 QSFP28-100G-CU1M JUNIPER-AMPHENOL 4. JNP-100G-DAC-1M 740-061000 1RC40513116 QSFP28-100G-CU1M JUNIPER-LEONI 5. JNP-100G-DAC-3M 740-061001 1RC4251201L QSFP28-100G-CU3M JUNIPER-LEONI 6. JNP-100G-DAC-3M 740-061001 1PC423161C5 QSFP28-100G-CU3M JUNIPER-AMPHENOL LED behaviour: ----- 1. when admin down in DUT: DUT(EX9208): amber, Non-DUT(EX9208): amber 2. when admin down in NON-DUT: DUT(EX9208): amber, Non-DUT(EX9208): amber Reason for behaviour change: For the above DAC cables (1-6), the above mentioned LED behaviour is expected since there is no LOS capability for these cables on peer side. AOC optics ----- 1. JNP-100G-AOC-5M 740-065632 1FCS251700A QSFP28-100G-AOC-5M JUNIPER-INNO 2. JNP-100G-AOC-15M 740-068217 1FCS5515004 QSFP28-100G-AOC-15M JUNIPER-INNO 3. JNP-100G-AOC-10M 740-061411 1FCS4517015 QSFP28-100G-AOC-10M JUNIPER-INNO 3. JNP-100G-AOC-30M 740-064980 1FCS7450001 QSFP28-100G-AOC-30M JUNIPER-INNO 4. JNP-100G-AOC-3M 740-065631 1FCS152004L QSFP28-100G-AOC-3M JUNIPER-INNO 5. JNP-100G-AOC-7M 740-065633 1FCS3520025 QSFP28-100G-AOC-7M JUNIPER-INNO LED behaviour: ----- 1. when admin down in DUT: DUT(EX9208): off, Non-DUT(EX9208): off 2. when admin down in Non-DUT: DUT(EX9208): off, Non-DUT(EX9208): off Reason for behaviour change: For the above the AOC optics part number (1-5), the above mentioned LED behaviour is expected because there is a false LOS seen on DUT(EX9208). This false LOS is due to low power mode when interface is disabled. In low power mode, the correctness/Working of alarm flags is dependent upon the vendor implementation 40g LX4 ----- 1. JNP-QSFP-40G-LX4 740-056705 1FCP251000Z QSFP+40GE-LX4 JUNIPER-INNO LED behaviour: ----- 1. when admin down in DUT: DUT(EX9208): off, Non-DUT(EX9208): off 2. when admin down in Non-DUT: DUT(EX9208): off, Non-DUT(EX9208): off Reason for behaviour change: For the above the LX4 optics vendor part number, the above mentioned LED behaviour is expected because there is a false LOS seen on DUT(EX9208). This false LOS is due to low power mode when the interface is disabled. In low power mode, the correctness/Working of alarm flags is dependent on the vendor implementation. [PR1532930](#)



Infrastructure

- On EX4300-MP, 9000 IPv6 MC routes can be installed. If you try to add more IPv6 MC routes, error messages will be seen. [PR1493671](#)

SEE ALSO

<a href="#">What's New</a>	<a href="#">  43</a>
<a href="#">What's Changed</a>	<a href="#">  54</a>
<a href="#">Open Issues</a>	<a href="#">  58</a>
<a href="#">Resolved Issues</a>	<a href="#">  60</a>
<a href="#">Documentation Updates</a>	<a href="#">  64</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  65</a>

Open Issues

IN THIS SECTION

- [EVPN](#) | [59](#)
- [General Routing](#) | [59](#)
- [Infrastructure](#) | [59](#)
- [Layer 2 Features](#) | [60](#)
- [Network Management and Monitoring](#) | [60](#)
- [Platform and Infrastructure](#) | [60](#)

Learn about open issues in Junos OS Release 20.3R1 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.



## EVPN

- After a reboot during recovery process the ESI LAGs come up before the BGP sessions and routes/ARP entries are not synchronized. [PR1487112](#)

## General Routing

- The output of the **show interface ge-x/x/x** command displays **Duplex: Half-duplex** when link-mode is set to automatic or is not set. This is a display issue and has no service impact. [PR1364659](#)
- On EX4300-48MP, EX4650, and EX4300, platforms, unicast RPF check in strict mode might not work properly. [PR1417546](#)
- When VLAN members are specified as a string, the 'IF\_MSG\_IFL\_VADDR' TLV is not generated with the VLAN information, and the MX Series MPCs is not updated with the native VLAN-ID and native VLAN enable flag, the packets are still treated as untagged. When packets reach the trunk egress interface they are getting dropped as the trunk interface does not allow untagged traffic to pass through. The issue is specific to platforms with ZT line cards. [PR1506403](#)
- A 35-second delay is added in reboot time from Junos OS Release 20.2R1 compared to Junos OS Release 19.4R2. [PR1514364](#)
- In the Layer 2 circuit termination scenario with input-vlan-map or output-vlan-map and family ccc, the output-vlan-map push operation might not work. It has a traffic impact. [PR1510629](#)

## Infrastructure

- 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](#)
- "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](#)



Layer 2 Features

- GARPs are being sent whenever there is a MAC (fdb) operation (add or delete). This is now updated to send GARP when the interface is up and the Layer 3 interface attached to the VLAN. [PR1192520](#)

Network Management and Monitoring

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

Platform and Infrastructure

- On EX9208 switches, 33 percent degradation in MAC learning rate is seen in Junos OS Release 19.3R1 when compared to Junos OS Release 18.4R1. [PR1450729](#)
- After GRES, interfaces might flap due to which DHCP bindings might be lost. [PR1515234](#)

SEE ALSO

<a href="#">What's New</a>	<a href="#">  43</a>
<a href="#">What's Changed</a>	<a href="#">  54</a>
<a href="#">Known Limitations</a>	<a href="#">  56</a>
<a href="#">Resolved Issues</a>	<a href="#">  60</a>
<a href="#">Documentation Updates</a>	<a href="#">  64</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  65</a>

Resolved Issues

IN THIS SECTION

- [Authentication and Access Control](#) | 61
- [EVPN](#) | 61
- [General Routing](#) | 61
- [Infrastructure](#) | 62
- [Interfaces and Chassis](#) | 63
- [Layer 2 Ethernet Services](#) | 63



●	Layer 2 Features   63
●	MPLS   63
●	Platform and Infrastructure   63
●	Routing Protocols   63
●	User Interface and Configuration   64
●	Virtual Chassis   64

This section lists the issues fixed in Junos OS Release 20.3R1 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.

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

SEE ALSO

<a href="#">What's New   43</a>
<a href="#">What's Changed   54</a>
<a href="#">Known Limitations   56</a>
<a href="#">Open Issues   58</a>
<a href="#">Documentation Updates   64</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   65</a>

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for EX Series switches.

SEE ALSO

<a href="#">What's New   43</a>
<a href="#">What's Changed   54</a>
<a href="#">Known Limitations   56</a>
<a href="#">Open Issues   58</a>
<a href="#">Resolved Issues   60</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   65</a>



# Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 65](#)

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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.

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

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

SEE ALSO

<a href="#">What's New   43</a>
<a href="#">What's Changed   54</a>
<a href="#">Known Limitations   56</a>
<a href="#">Open Issues   58</a>
<a href="#">Resolved Issues   60</a>



# Junos OS Release Notes for JRR Series

## IN THIS SECTION

- What's New | 66
- What's Changed | 68
- Known Limitations | 68
- Open Issues | 69
- Resolved Issues | 69
- Documentation Updates | 70
- Migration, Upgrade, and Downgrade Instructions | 70

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- Routing Protocols | 67

Learn about new features introduced in Junos OS Release 20.3R1 for JRR Series Route Reflectors.



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

<a href="#">What's Changed</a>	<a href="#">  68</a>
<a href="#">Known Limitations</a>	<a href="#">  68</a>
<a href="#">Open Issues</a>	<a href="#">  69</a>
<a href="#">Resolved Issues</a>	<a href="#">  69</a>
<a href="#">Documentation Updates</a>	<a href="#">  70</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  70</a>



# What's Changed

IN THIS SECTION

- [General Routing | 68](#)

Learn about what changed in this release for JRR200 Route Reflectors.

## General Routing

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

SEE ALSO

<a href="#">What's New   66</a>
<a href="#">Known Limitations   68</a>
<a href="#">Open Issues   69</a>
<a href="#">Resolved Issues   69</a>
<a href="#">Documentation Updates   70</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   70</a>

# Known Limitations

There are no known limitations JRR Series in Junos OS Release 20.3R1 for JRR Series Route Reflectors.

SEE ALSO

<a href="#">What's New   66</a>
<a href="#">What's Changed   68</a>
<a href="#">Open Issues   69</a>
<a href="#">Resolved Issues   69</a>



<a href="#">Documentation Updates</a>	<a href="#">70</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">70</a>

## Open Issues

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

### SEE ALSO

<a href="#">What's New</a>	<a href="#">66</a>
<a href="#">What's Changed</a>	<a href="#">68</a>
<a href="#">Known Limitations</a>	<a href="#">68</a>
<a href="#">Resolved Issues</a>	<a href="#">69</a>
<a href="#">Documentation Updates</a>	<a href="#">70</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">70</a>

## Resolved Issues

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

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

### SEE ALSO

<a href="#">What's New</a>	<a href="#">66</a>
<a href="#">What's Changed</a>	<a href="#">68</a>
<a href="#">Known Limitations</a>	<a href="#">68</a>
<a href="#">Open Issues</a>	<a href="#">69</a>
<a href="#">Documentation Updates</a>	<a href="#">70</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">70</a>



## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for JRR200 Route Reflectors.

### SEE ALSO

[What's New | 66](#)[What's Changed | 68](#)[Known Limitations | 68](#)[Open Issues | 69](#)[Resolved Issues | 69](#)[Migration, Upgrade, and Downgrade Instructions | 70](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 70](#)

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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2,



19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.

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

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

#### SEE ALSO

<a href="#">What's New   66</a>
<a href="#">What's Changed   68</a>
<a href="#">Known Limitations   68</a>
<a href="#">Open Issues   69</a>
<a href="#">Resolved Issues   69</a>
<a href="#">Documentation Updates   70</a>

## Junos OS Release Notes for Juniper Secure Connect

#### IN THIS SECTION

- [What's New | 72](#)
- [Open Issues | 73](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).



## What's New

### IN THIS SECTION

- [Juniper Secure Connect | 72](#)

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

### 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 logon	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.



Table 3: Feature Support for Juniper Secure Connect (*continued*)

Feature	Description
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](#).]

## Open Issues

### IN THIS SECTION

- [Platform and Infrastructure](#) | 73
- [VPNs](#) | 74

This section lists the known issues in hardware and software in Junos OS Release 20.3R1 for Juniper Secure Connect.

Learn about open issues in this release 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.

### Platform and Infrastructure

- 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](#)
- 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](#)



## VPNs

- IKE DH group24 and IPsec PFS group24 are not supported from Juniper Secure client, though these are supported on the SRX Series device. [PR1506966](#)
- 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](#)

# Junos OS Release Notes for Junos Fusion for Enterprise

## IN THIS SECTION

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

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

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



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

SEE ALSO

<a href="#">What's Changed</a>	<a href="#">75</a>
<a href="#">Known Limitations</a>	<a href="#">75</a>
<a href="#">Open Issues</a>	<a href="#">76</a>
<a href="#">Resolved Issues</a>	<a href="#">76</a>
<a href="#">Documentation Updates</a>	<a href="#">77</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">77</a>

## What's Changed

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.3R1 for Junos fusion for enterprise.

SEE ALSO

<a href="#">What's New</a>	<a href="#">74</a>
<a href="#">Known Limitations</a>	<a href="#">75</a>
<a href="#">Open Issues</a>	<a href="#">76</a>
<a href="#">Resolved Issues</a>	<a href="#">76</a>
<a href="#">Documentation Updates</a>	<a href="#">77</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">77</a>

## Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.3R1 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

<a href="#">What's New   74</a>
<a href="#">What's Changed   75</a>
<a href="#">Open Issues   76</a>
<a href="#">Resolved Issues   76</a>
<a href="#">Documentation Updates   77</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   77</a>

## Open Issues

There are no known issues in hardware and software in Junos OS Release for 20.3R1 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

<a href="#">What's New   74</a>
<a href="#">What's Changed   75</a>
<a href="#">Known Limitations   75</a>
<a href="#">Resolved Issues   76</a>
<a href="#">Documentation Updates   77</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   77</a>

## Resolved Issues

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



## SEE ALSO

<a href="#">What's New</a>	<a href="#">74</a>
<a href="#">What's Changed</a>	<a href="#">75</a>
<a href="#">Known Limitations</a>	<a href="#">75</a>
<a href="#">Open Issues</a>	<a href="#">76</a>
<a href="#">Documentation Updates</a>	<a href="#">77</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">77</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 for documentation for Junos fusion for enterprise.

## SEE ALSO

<a href="#">What's New</a>	<a href="#">74</a>
<a href="#">What's Changed</a>	<a href="#">75</a>
<a href="#">Known Limitations</a>	<a href="#">75</a>
<a href="#">Open Issues</a>	<a href="#">76</a>
<a href="#">Resolved Issues</a>	<a href="#">76</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">77</a>

## Migration, Upgrade, and Downgrade Instructions

## IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device](#) | [78](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines](#) | [79](#)
- [Preparing the Switch for Satellite Device Conversion](#) | [80](#)
- [Converting a Satellite Device to a Standalone Switch](#) | [81](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [81](#)
- [Downgrading Junos OS](#) | [82](#)



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

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




You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.

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

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

**Downgrading 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

<a href="#">What's New</a>	<a href="#">74</a>
<a href="#">What's Changed</a>	<a href="#">75</a>
<a href="#">Known Limitations</a>	<a href="#">75</a>
<a href="#">Open Issues</a>	<a href="#">76</a>
<a href="#">Resolved Issues</a>	<a href="#">76</a>
<a href="#">Documentation Updates</a>	<a href="#">77</a>



# Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- [What's New | 83](#)
- [What's Changed | 84](#)
- [Known Limitations | 84](#)
- [Open Issues | 85](#)
- [Resolved Issues | 85](#)
- [Documentation Updates | 86](#)
- [Migration, Upgrade, and Downgrade Instructions | 87](#)

These release notes accompany Junos OS Release 20.3R1 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](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

- |   |
|---|
| <a href="#">What's Changed   84</a>                                 |
| <a href="#">Known Limitations   84</a>                              |
| <a href="#">Open Issues   85</a>                                    |
| <a href="#">Resolved Issues   85</a>                                |
| <a href="#">Documentation Updates   86</a>                          |
| <a href="#">Migration, Upgrade, and Downgrade Instructions   87</a> |



# What's Changed

IN THIS SECTION

- [General Routing | 84](#)

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.

## General Routing

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

SEE ALSO

<a href="#">What's New   83</a>
<a href="#">Known Limitations   84</a>
<a href="#">Open Issues   85</a>
<a href="#">Resolved Issues   85</a>
<a href="#">Documentation Updates   86</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   87</a>

# Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.3R1 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



<a href="#">What's New</a>	<a href="#">  83</a>
<a href="#">What's Changed</a>	<a href="#">  84</a>
<a href="#">Open Issues</a>	<a href="#">  85</a>
<a href="#">Resolved Issues</a>	<a href="#">  85</a>
<a href="#">Documentation Updates</a>	<a href="#">  86</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  87</a>

## Open Issues

There are no open issues in the Junos OS Release 20.3R1 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

<a href="#">What's New</a>	<a href="#">  83</a>
<a href="#">What's Changed</a>	<a href="#">  84</a>
<a href="#">Known Limitations</a>	<a href="#">  84</a>
<a href="#">Resolved Issues</a>	<a href="#">  85</a>
<a href="#">Documentation Updates</a>	<a href="#">  86</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  87</a>

## Resolved Issues

### IN THIS SECTION

- [Junos Fusion Provider Edge](#) | 86

This section lists the issues fixed in the Junos OS Release 20.3R1 for Junos fusion for provider edge.



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

**Junos Fusion Provider Edge**

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

SEE ALSO

<a href="#">What's New</a>	<a href="#">  83</a>
<a href="#">What's Changed</a>	<a href="#">  84</a>
<a href="#">Known Limitations</a>	<a href="#">  84</a>
<a href="#">Open Issues</a>	<a href="#">  85</a>
<a href="#">Documentation Updates</a>	<a href="#">  86</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  87</a>

**Documentation Updates**

There are no errata or changes in Junos OS Release 20.3R1 documentation for Junos fusion for provider edge.

SEE ALSO

<a href="#">What's New</a>	<a href="#">  83</a>
<a href="#">What's Changed</a>	<a href="#">  84</a>
<a href="#">Known Limitations</a>	<a href="#">  84</a>
<a href="#">Open Issues</a>	<a href="#">  85</a>
<a href="#">Resolved Issues</a>	<a href="#">  85</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  87</a>



## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 87
- Upgrading an Aggregation Device with Redundant Routing Engines | 90
- Preparing the Switch for Satellite Device Conversion | 90
- Converting a Satellite Device to a Standalone Device | 92
- Upgrading an Aggregation Device | 94
- Upgrade and Downgrade Support Policy for Junos OS Releases | 94
- Downgrading from Junos OS Release 20.1 | 95

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 20.3R1 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/jinstall64-20.3R1.SPIN-domestic-signed.tgz
```



- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.3R1.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/jinstall64-20.3R1.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.3R1.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 20.3R1 **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 20.3R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

## Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.


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

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



## Downgrading from Junos OS Release 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

<a href="#">What's New</a>		<a href="#">83</a>
<a href="#">What's Changed</a>		<a href="#">84</a>
<a href="#">Known Limitations</a>		<a href="#">84</a>
<a href="#">Open Issues</a>		<a href="#">85</a>
<a href="#">Resolved Issues</a>		<a href="#">85</a>
<a href="#">Documentation Updates</a>		<a href="#">86</a>

# Junos OS Release Notes for MX Series

IN THIS SECTION

- [What's New](#) | [96](#)
- [What's Changed](#) | [120](#)
- [Known Limitations](#) | [124](#)
- [Open Issues](#) | [127](#)
- [Resolved Issues](#) | [135](#)
- [Documentation Updates](#) | [150](#)
- [Migration, Upgrade, and Downgrade Instructions](#) | [150](#)



These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- Hardware | 97
- Authentication, Authorization, and Accounting | 100
- Class of Service (CoS) | 100
- EVPN | 100
- High Availability (HA) and Resiliency | 101
- Interfaces and Chassis | 101
- IP Tunneling | 102
- Juniper Extension Toolkit | 103
- Junos OS XML, API, and Scripting | 104
- Junos Telemetry Interface | 104
- Layer 2 Features | 111
- Layer 2 VPN | 111
- Layer 3 Features | 111
- MPLS | 111
- Multicast | 112
- Network Management and Monitoring | 112
- Next Gen Services | 113
- Port Security | 114
- Routing Protocols | 115
- Segment Routing | 117
- Services Applications | 117
- Software Defined Networking (SDN) | 118
- System Management | 119



This section describes the new features and enhancements to existing features in Junos OS Release 20.3R1 for the MX Series routers.

## 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"> <li>• 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.</li> <li>• 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 <a href="#">Port Speed on MX2K-MPC11E Overview</a>.]</li> </ul>
General routing	<ul style="list-style-type: none"> <li>• Support for IP reassembly on GRE tunnel interfaces on: <ul style="list-style-type: none"> <li>• MPC10E-15C-MRATE and MPC10E-10C-MRATE on MX240, MX480, and MX960 routers.</li> <li>• MX2K-MPC11E on MX2010 and MX2020 routers.</li> </ul> [See <a href="#">Configuring Unicast Tunnels</a>.]</li> <li>• Support for Mapping of Address and Port with Encapsulation (MAP-E) and IPv6 rapid deployment (inline 6rd) on: <ul style="list-style-type: none"> <li>• MPC10E-15C-MRATE and MPC10E-10C-MRATE on MX240, MX480, and MX960 routers.</li> <li>• MX2K-MPC11E on MX2010 and MX2020 routers.</li> </ul> [See <a href="#">Configuring Mapping of Address and Port with Encapsulation (MAP-E)</a> and <a href="#">Configuring Inline 6rd</a>.]</li> </ul>
Juniper telemetry interface	<ul style="list-style-type: none"> <li>• Support for resource paths to export traffic statistics from LDP and multipoint LDP sensors with gRPC. [See <a href="#">Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface)</a>.]</li> <li>• Support for LDP and multipoint LDP native sensors. [See <a href="#">sensor (Junos Telemetry Interface)</a>.]</li> </ul>
Layer 3 features	<ul style="list-style-type: none"> <li>• 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 <a href="#">Adaptive Services Overview</a>.]</li> </ul>



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

Feature	Description
Multicast	<ul style="list-style-type: none"> <li>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 <a href="#">Understanding Bidirectional PIM.</a>]</li> </ul> <p><b>NOTE:</b> 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"> <li>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 <a href="#">Understanding AMT.</a>]</li> </ul> <p><b>NOTE:</b> Junos OS Release 20.3R1 does not support AMT gateway.</p>
Network management and monitoring	<ul style="list-style-type: none"> <li>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 <a href="#">Link Degradation Monitoring Overview.</a>]</li> <li>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 <a href="#">Inline Transmission Mode.</a>]</li> </ul>
Security	<ul style="list-style-type: none"> <li>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 <a href="#">Media Access Control Security (MACsec) over WAN.</a>]</li> </ul>
Services applications	<ul style="list-style-type: none"> <li>Support for inline video monitoring using media delivery index (MDI) criteria. [See <a href="#">Understanding Inline Video Monitoring on MX Series Routers.</a>]</li> </ul>



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

Feature	Description
SNMP	<ul style="list-style-type: none"> <li>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"> <li>mplsMldpInterfaceStatsTable</li> <li>mplsMldpFecUpstreamSessPackets</li> <li>mplsMldpFecUpstreamSessBytes</li> <li>mplsMldpFecUpstreamSessDiscontinuityTime</li> </ul> </li> </ul> <p>[See <a href="#">Standard SNMP MIBs Supported by Junos OS</a> and <a href="#">SNMP MIB Explorer</a>.]</p>
Subscriber management and services	<ul style="list-style-type: none"> <li>Support for resource monitoring for broadband edge subscriber management and services. [See <a href="#">Resource Monitoring for Subscriber Management and Services</a>.]</li> </ul>

- **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](#).]

## 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](#).]

## 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
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/conditional/condition/state/dscp[]



Table 5: Supported Sensors (*continued*)

Supported Sensors
/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.
/network-instances/network-instance/afts/next-hop-groups/next-hop-group/id



Table 5: Supported Sensors (*continued*)

Supported Sensors
<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-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 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)

- **Support for static FTI backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (MX Series, PTX Series, and QFX10002)**—Starting in Junos OS Release 20.3R1, we've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to implement underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation on flexible tunnel interfaces (FTIs). Use this feature to engineer effective, loop-free backup paths for core transport networks built with only IP protocols for fast restoration after failures.



We've extended FTIs and existing forwarding constructs to support configuring static IPv4 IP-in-IP tunnels. You can also allow policy matches for routes injected by JET APIs.

[See [policy-statement](#), [tunnel](#), [ipip](#), [show interfaces](#), [show route](#), [Configuring Flexible Tunnel Interfaces](#), and [JET APIs on the Juniper Engineering Network website](#).]

- **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 connection is made, the clksync module gets the notification.

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

## SEE ALSO

[What's Changed | 120](#)

[Known Limitations | 124](#)

[Open Issues | 127](#)

[Resolved Issues | 135](#)

[Documentation Updates | 150](#)

[Migration, Upgrade, and Downgrade Instructions | 150](#)



## What's Changed

### IN THIS SECTION

- [EVPN | 121](#)
- [General Routing | 121](#)
- [High Availability \(HA\) and Resiliency | 122](#)
- [Infrastructure | 122](#)
- [Interfaces and Chassis | 122](#)
- [Junos OS, XML, API, and Scripting | 122](#)
- [Routing Protocols | 123](#)
- [Services Applications | 123](#)
- [Subscriber Management and Services | 123](#)

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 20.3R1 for MX Series routers.



## EVPN

- **New output flag for the show bridge mac-ip table command (MX 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](#).]

- **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**.



## 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.

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

## 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.

## 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 `and` .

SEE ALSO

| [What's New](#) | 96



[Known Limitations | 124](#)[Open Issues | 127](#)[Resolved Issues | 135](#)[Documentation Updates | 150](#)[Migration, Upgrade, and Downgrade Instructions | 150](#)

## Known Limitations

### IN THIS SECTION

- [EVPN | 124](#)
- [General Routing | 124](#)
- [Interfaces and Chassis | 125](#)
- [MPLS | 125](#)
- [Network Management and Monitoring | 125](#)
- [Platform and Infrastructure | 126](#)
- [Routing Protocols | 127](#)
- [Subscriber Management and Services | 127](#)

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

### EVPN

- EVPN service over uncolored scaled SR-TE is not supported. [PR1499719](#)

### General Routing

- Subscriber access facing CPU's utilization of FPC remains 100 percent for 56 minutes after making changes to the service firewall filter configuration. [PR1447003](#)
- On the MPC11E line card, the following error messages are seen when the line card is online: **i2c transaction error (0x00000002)**. [PR1457655](#)



- On the MX10003 routers, during ISSU from Junos OS Release 18.2X75-D430 to 20.3R1, the link flaps and traffic drops for MACsec TIC, [PR1514694](#)
- The MPC11E line card might take additional time to come during the movement from one GNF to another GNF. [PR1469729](#)
- On the MX10003 or MX204 routers, BFD or LACP might flap during the BGP convergence. [PR1472587](#)
- PSX.0 changes that have active subscribers causes the line card to crash. [PR1486665](#)
- Observing traffic drop of around 2.5 seconds on switchover from primary physical interface to backup FTI interface with scaled routes. [PR1490070](#)
- During MBB, a few packets may be dropped while bringing up the FTI logical interface, which is the primary interface. [PR1507779](#)
- Memory leak observed in the JSD process with one or more collector(s) connecting and disconnecting during streaming of the telemetry data. [PR1512296](#)

## Interfaces and Chassis

- Traffic stalled and standby PWS states are not updated on changing to vlan-bridge encapsulation and then back to vlan-circuit-cross-connect. [PR1503102](#)

## MPLS

- On the MX480 router, the following error message is observed: **FPC Resource Monitor: FPC 0 and 1 Heap Memory has crossed free memory watermark of 20.** [PR1513436](#)

## 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

- PIM join message (S,G) might not be created after GRES. [PR1457166](#)
- Unknown unicast filter applied in the EVPN routing instance blocks unexpected traffic. [PR1472511](#)
- Even after subscribing to `/junos/system/linecard/firewall/`, starting the GNMI decoder and performing negative interface triggers the subscription and TCP session remains. [PR1477790](#)
- EVPN aliasing does not work over a SR-TE underlay. [PR1504412](#)



## 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 | 96](#)

[What's Changed | 120](#)

[Open Issues | 127](#)

[Resolved Issues | 135](#)

[Documentation Updates | 150](#)

[Migration, Upgrade, and Downgrade Instructions | 150](#)

## Open Issues

### IN THIS SECTION

- [Class of Service \(CoS\) | 128](#)
- [EVPN | 128](#)
- [Forwarding and Sampling | 128](#)
- [General Routing | 129](#)
- [Infrastructure | 132](#)
- [Interfaces and Chassis | 132](#)
- [Intrusion Detection and Prevention \(IDP\) | 133](#)



- [Layer 2 Ethernet Services | 133](#)
- [MPLS | 133](#)
- [Network Management and Monitoring | 133](#)
- [Platform and Infrastructure | 133](#)
- [Routing Protocols | 134](#)
- [Subscriber Access Management | 134](#)
- [User Interface and Configuration | 134](#)

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

### Class of Service (CoS)

- When an interface attached to the aggregated Ethernet interface is decoupled and an IP address is assigned to it, ARP resolution issues are seen. [PR1504287](#)

### EVPN

- The VXLAN OAM host-bound packets are not throttled with DDoS policers. [PR1435228](#)
- With dynamic list next hop configured, a forwarding problem occurs after performing graceful switchover. [PR1513759](#)
- The **GE LOS** alarm that gets logged on the change in **IFF\_CCCDOWN** are not logged in the syslog message file. [PR1539146](#)
- On the MX480 router, the expected EVPN type 5 routes: 4 is not the same as the actual EVPN type 5 route: 0. [PR1535353](#)
- All the ARP reply packets towards to some address are flooded across the entire fabric. [PR1535515](#)

### Forwarding and Sampling

- The following syslog error message might be seen if the SSD hardware fails: **rp[2191]: krt\_flow\_dfwd\_open,8073: Failed connecting to DFWD, error checking reply - Operation timed out.** [PR1397171](#)
- The **srrd** process might crash in a high route churns scenario or if the process flaps. [PR1517646](#)



- The l2ald process might crash when the device configuration flaps frequently. [PR1529706](#)
- VLAN-ID based firewall match conditions might not work for the VPLS service. [PR1542092](#)

## General Routing

- If a vmhost snapshot is taken on an alternate disk and no further vmhost software image is upgraded, 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](#)
- The output of the **show dynamic-tunnels database statistics** command must have tags for source, destination, tunnel-id, and next hop. [PR1501576](#)
- On the MX150 routers, fpc flaps after loading the **set chassis fpc slot flexible-queuing-mode** configuration and generates the **localhost.bcmod.mpc0** core file sometimes. [PR1534637](#)
- On the MX2010 and MX2020 routers equipped with SFB2, some error logs might be seen. [PR1363587](#)
- FPC generates core files under certain circumstances on addition and deletion of hierarchical CoS from pseudowire devices. [PR1414969](#)
- The FPC might crash when the Packet Forwarding Engine memory exhausts. [PR1439012](#)
- On the MPC11E line card, the **number-of-sub-ports** configuration on the 4x10G channelized ports might cause the channels to go down. [PR1442439](#)
- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card. [PR1444516](#)
- IPv6 packets might get dropped when vMX acts as a VRRPv3 gateway. [PR1449014](#)
- Physical interface policers are not supported in Junos OS Release 19.3 for the MPC11 line card. [PR1452963](#)
- On the MPC11E line card, the FIB download rates are lower than the rates on MPC10E line card by 30 percent. [PR1456816](#)
- The following CDA error message is observed: **LkupAsicClient: Index Dmem block read failed, PFE:0.0.** [PR1459665](#)
- Need to add the Backport jemalloc profiling CLI support to all Junos OS releases where jemalloc is present. [PR1463368](#)
- The following error message is seen 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](#)
- For the MPC10E line card, the IS-IS and micro-BFD sessions do not come up during baseline. [PR1474146](#)
- Dynamic SR-TE tunnels do not get automatically re-created at the new master Routing Engine after the Routing Engine switchover. [PR1474397](#)



- Expected number of 512,000 MAC entries are not re-learned in the bridge table after clearing 512,000 MAC entries from the table. [PR1475205](#)
- On the MX480 router, the following error message is seen after restore or removal with IP/MPLS configurations: **[Error] L2alm : l2alm\_mac\_process\_hal\_delete\_msg:667 Ignoring MAC delete with ifl index 355, fwd\_entry has 7888.** [PR1475785](#)
- Critical syslog error messages at **fpc3 user.crit aftd-trio** are seen during baseline: **[Critical] Em: Possible out of order deletion of AftNode #012#012#012 AftNode details - AftIndirect token:230791 group:0 nodeMask:0xffffffffffffffff indirect:333988 hwInstall:1#012.** [PR1486158](#)
- Login or logout of high scale (around 1 million bearers) causes some sessions not to re-login. [PR1489665](#)
- Backup Routing Engine reboots because of power cycle or failure when the offline and online operations are performed on CB1. [PR1497592](#)
- If MPLS is needed, the crpd container must be instantiated with the MPLS modules that are already installed on the host. [PR1498632](#)
- The MPC11 line card is not supported in Junos OS Release 19.4. [PR1503605](#)
- Traffic loss might be seen under ECMP scenario on the MPC10E or MPC11E line card. [PR1513898](#)
- On the MX960 router, expected traffic is not received with multicast and PIM scaling configurations. [PR1514646](#)
- **yin2tlv** sets unsupported command as **hidden deprecated**. [PR1516910](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- With the HTTP header enrichment function enabled, the processing of the window scaling option significantly reduces the performance of HTTP sessions from 65 Mbps to less than 40 Mbps, which results in decrease of traffic throughput. The download rate also drops. [PR1420894](#)
- Changing the framing modes on a CHE1T1 MIC between E1 and T1 on an MPC3E NG HQoS line card causes the PIC to go offline. [PR1474449](#)
- On the MX2020 and MX2010 routers, the SPMB CPU is elevated when an SFB3 is installed. [PR1516287](#)
- The AMS bundle state toggles momentarily as up, down, and up after configuring commit for a scaled scenario. [PR1521929](#)
- Family IPv6 do not come up for the L2TP subscriber when additional attributes are not passed in the Framed-IPv6-Route VSA. [PR1526934](#)
- In the subscriber management environment, the RADIUS interim accounting records does not get populated with the subscriber statistics. [PR1529602](#)
- The following error messages are observed: **unable to set line-side lane config (err 30).** [PR1492162](#)
- LFM might flap during MX-VC ISSU. [PR1516744](#)
- Subscribers are not logged out after the AGT test stops. [PR1531415](#)



- Multicast traffic might be sent out through unexpected interfaces with distributed IGMP enabled. [PR1536149](#)
- The accounting interim-updates for subscriber does not work after GRES and subsequent reboot of FPCs in the node-slicing setup. [PR1539474](#)
- The npc process generates core file in `igmp_process_wakeup_events,igmp_pfe_thread,thread_detach_tty`. [PR1534542](#)
- On the MPC10 line card, AFT crash is observed at `std::default_delete< AftTermAction>::operator()` (`this=< optimized out>, __ptr=0x7fb0bc5d5910`) at `/volume/evo/files/opt/poky/2.2.1-22/sysroots/core2-64-poky-linux/usr/include/c++/6.2.0/bits/unique_ptr.h:76`. [PR1491527](#)
- The VPLS connection might get stuck in the primary fail status when a dynamic profile is used on the VPLS pseudowire logical interface. [PR1516418](#)
- On the MX960 router, the `show interfaces redundancy rlt0` command shows the current status as primary down as FPC is still in the **Ready** state after the rlt failover. [PR1518543](#)
- The commit error messages comes twice while validating the **physical-cores** commands. [PR1527322](#)
- VRRP synchronization does not occur in the backup Routing Engine with NSR in the **Steady** state. [PR1533357](#)
- Version-alias is missed for the subscribers that are configured with the dynamic profiles after ISSU. [PR1537512](#)
- The `vmxt_lnx` process generates core file at `l2_metro_bd_host_inject_del bd_platform_delete bd_handle_msg`. [PR1538516](#)
- On the MX480 router, after disabling and enabling the primary links in a network with a seamless BFD support (inline) for SR-TE, traffic loss of more than 1000 minutes is observed. [PR1539376](#)
- On the MX960 router, the MLDP Egress statistics does not get populated in the CLI after enabling the sensor-based-statistics under the LPD traffic-statistics. [PR1539450](#)
- With hold time configuration, the GE Interfaces remain down on reboot. [PR1541382](#)
- After changing addresses in the source pool, if the CGNAT traffic does not stop, the NAT translation from the new pool cannot be done. [PR1542202](#)
- Port mirroring with the **maximum-packet-length** configuration does not work over GRE interface. [PR1542500](#)
- On the MPC10 line card, the following error message is observed on the Routing Engine 1 after graceful switchover from the Routing Engine 0 to the Routing Engine 1: **[Error] L2ALIPC : L2AL IPC client failed to connect to l2ald**. [PR1491384](#)
- Next hop are not as expected as a total of nhdb 4MPOST GRES. [PR1539305](#)



- Packets lost is observed during switchover at 100 ms sampling period. [PR1539310](#)
- Subscriber might not come up on some dynamic VLAN ranges in the subscriber management environment. [PR1541796](#)

## Infrastructure

- The following error message is observed continuously in AD with base configurations: **IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) failed.** [PR1485038](#)

## Interfaces and Chassis

- The SFP index in the Packet Forwarding Engine starts at 1, while the port numbering starts at 0. This causes confusion in the log analysis. [PR1412040](#)
- MPLS VPN label can point to the discard next hop after a Routing Engine switchover without NSR if the egress interface is pp0. [PR1488302](#)
- Input and output bytes count mismatch in the IPv6 traffic statistics while issuing the **show interface extensive** command. [PR1505100](#)
- When standby MC-LAG node is rebooted, one time traffic hit of active path traffic is observed and later when the node comes up, the MC-LAG active standby roles are changed to the other device. [PR1505841](#)
- The following error message is observed while removing or adding the configurations: **xolo-fpc0 ppman: [Error] CTRL:RPC:: Cos8021pRwTableCb)::< lambda: RPC to Aftman CoS FC table request failed for key:16783744 iflIndex:23238 status:Invalid argument.** [PR1527032](#)
- The following error message is observed when the CCM PDUs do not take the configured forwarding class: **ppman: [Error] PPM:CTRL\_CFM: PpmCtrlProtoCfm::getFcPlp: CFM interface is not found in intf table.** [PR1534239](#)
- In-line CFM performance monitoring PDU cannot be transported across the EVPN domain. [PR1537381](#)



## Intrusion Detection and Prevention (IDP)

- The CLI provides helpful remarks about the tunable detector parameters of IDP. [PR1490436](#)

## Layer 2 Ethernet Services

- The jdhcpd process crashes while forwarding a malformed DHCP packet. [PR1430874](#)

## MPLS

- The rpd process might crash when the LDP route with indirect next hop is deleted on the aggregated Ethernet interface. [PR1538124](#)
- The RPD scheduler might slip after the link flaps. [PR1516657](#)

## Network Management and Monitoring

- On the MPC11E line card, the following trap message is not observed after a LC reboot when the scaled interfaces are present: **SNMP Link up**. [PR1507780](#)
- Traffic statistics in the **show interface** command is displayed with incorrect cumulative values. [PR1539483](#)
- The MIB value for **ipv6IfStatsInReceives.576** does not match the expected output. [PR1539483](#)

## Platform and Infrastructure

- Packet loss might be observed when the RFC2544 egress reflector session is configured on the non-zero Packet Forwarding Ethernet interface. [PR1538417](#)
- On the MX2010 router, OSPF flapping during ISSU from Junos OS Release 17.3 to Release 17.4 is observed due to InActiveTimer event reason: neighbor was inactive and declared dead. [PR1371879](#)
- The CFM REMOTE MEP does not come up after configuration or if the MEP remains in the **Start** state. [PR1460555](#)
- 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](#)
- A few OAM sessions are not established with the scaled EVPN E-Tree and CFM configurations. [PR1478875](#)
- If the interface is newly added as the CE interface, the existing broadcast, unknown unicast, and multicast (BUM) traffic can be looped. The loop prevention feature is designed to start working whenever a new CE interface is added by configuration. But the existing BUM traffic can be distributed to a new CE interface earlier before enabling the loop prevention feature. [PR1493650](#)



- The MEP session on the aggregated Ethernet interface might not come up if OAM runs with PPM mode by default. [PR1506861](#)
- After performing ISSU with high scale bridge domain configuration, less than 0.0254 percent of traffic loss is observed for a single bridge domain interface. [PR1531051](#)

**Routing Protocols**

- After moving the peer out of the protection group, the path protection is not removed from the PE router. The multipath route are still present. [PR1538956](#)
- The Layer 3 VPN routes might be added to FIB on the route reflector. [PR1532414](#)
- The rpd process might crash on the backup Routing Engine if BGP (standby) receives a route from the peer, which is rejected due to an invalid target community. [PR1508888](#)
- On the MX960 router, the backup path fails to install in the LAN scenario and breaks the SR-MPLS for LAN when more than four end-x SIDs are configured on the interface. [PR1512174](#)
- BFD with authentication for BGP flaps after GRES or NSR switchover on the NG-RE and SCBE2 setup. [PR1522261](#)
- The IS-IS LSP database synchronization issue might be seen while using the flood-group feature. [PR1526447](#)

**Subscriber Access Management**

- Subscriber accounting messages retransmissions exist even after configuring accounting retry 0. [PR1405855](#)

**User Interface and Configuration**

- NETCONF service over SSH might not work on the device that runs Junos OS if in-band management is used. [PR1517160](#)

SEE ALSO

<a href="#">What's New   96</a>
<a href="#">What's Changed   120</a>
<a href="#">Known Limitations   124</a>
<a href="#">Resolved Issues   135</a>
<a href="#">Documentation Updates   150</a>



## Resolved Issues

### IN THIS SECTION

- Application Layer Gateways (ALGs) | 136
- Class of Service (CoS) | 136
- EVPN | 136
- Forwarding and Sampling | 137
- General Routing | 137
- Infrastructure | 143
- Interfaces and Chassis | 143
- Intrusion Detection and Prevention (IDP) | 144
- J-Web | 144
- Juniper Extension Toolkit (JET) | 144
- Junos Fusion Provider Edge | 144
- Layer 2 Ethernet Services | 144
- MPLS | 145
- Network Management and Monitoring | 146
- Platform and Infrastructure | 146
- Routing Protocols | 147
- Services Applications | 148
- Subscriber Access Management | 149
- Subscriber Management and Services | 149
- User Interface and Configuration | 149
- VPNs | 149

This section lists the issues fixed in Junos OS Release 20.3R1 for MX Series routers.

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



## Application Layer Gateways (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 0.0.0.0/0 | 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 cp added 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 nonexisting 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 **DHCPLEASEUNASSIGNED**, 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 **jdhcpd** might be seen if access-profile is configured under the **dhcp-relay** or **dhcp-local-server** statement. [PR1525052](#)
- Receipt of malformed DHCPv6 packets causes **jdhcpd** to crash. [PR1511782](#)
- The **jdhcpd** process crashes when processing a specific DHCPDv6 packet in the DHCPv6 relay configuration. [PR1512765](#)

## MPLS

- The RSVP interface bandwidth calculation rounds up. [PR1458527](#)
- The **rpdp** process might crash in PCEP for the RSVP-TE scenario. [PR1467278](#)
- The **rpdp** 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 **rpdp** process might crash in a rare condition in the SR-TE scenario. [PR1493721](#)
- The **rpdp** 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 **rpdp** 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 **rpdp** process might crash with RSVP configured in a rare timing case. [PR1505834](#)
- Activating or deactivating the LDP-sync under OSPF might cause the LDP neighborhood to go down and stay down. [PR1509578](#)
- The **rpdp** 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 be 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](#)

## 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.

## 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 | 96](#)

---



<a href="#">What's Changed</a>	<a href="#">  120</a>
<a href="#">Known Limitations</a>	<a href="#">  124</a>
<a href="#">Open Issues</a>	<a href="#">  127</a>
<a href="#">Documentation Updates</a>	<a href="#">  150</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  150</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for MX Series routers.

### SEE ALSO

<a href="#">What's New</a>	<a href="#">  96</a>
<a href="#">What's Changed</a>	<a href="#">  120</a>
<a href="#">Known Limitations</a>	<a href="#">  124</a>
<a href="#">Open Issues</a>	<a href="#">  127</a>
<a href="#">Resolved Issues</a>	<a href="#">  135</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  150</a>

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.3R1](#) | [151](#)
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS](#) | [151](#)
- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS](#) | [154](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [156](#)
- [Upgrading a Router with Redundant Routing Engines](#) | [156](#)
- [Downgrading from Release 20.3R1](#) | [157](#)



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 20.3R1

**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-x86-32-20.3R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-20.3R1.9-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-x86-32-20.3R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.3R1.9-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 20.3R1, 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 20.3R1 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-20.3R1.9-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-20.3R1.9-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 20.3R1 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

## Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.

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

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

## Upgrading a Router with Redundant Routing Engines

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

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



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

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

### Downgrading from Release 20.3R1

To downgrade from Release 20.3R1 to another supported release, follow the procedure for upgrading, but replace the 20.3R1 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

<a href="#">What's New   96</a>
<a href="#">What's Changed   120</a>
<a href="#">Known Limitations   124</a>
<a href="#">Open Issues   127</a>
<a href="#">Resolved Issues   135</a>
<a href="#">Documentation Updates   150</a>

## Junos OS Release Notes for NFX Series

#### IN THIS SECTION

- [What's New | 158](#)
- [What's Changed | 159](#)
- [Known Limitations | 160](#)
- [Open Issues | 161](#)
- [Resolved Issues | 163](#)



- Documentation Updates | 164
- Migration, Upgrade, and Downgrade Instructions | 165

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- Application Security | 158
- Wireless WAN | 159

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](#).

### 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

<a href="#">What's Changed   159</a>
<a href="#">Known Limitations   160</a>
<a href="#">Open Issues   161</a>
<a href="#">Resolved Issues   163</a>
<a href="#">Documentation Updates   164</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   165</a>

What's Changed

IN THIS SECTION

- [What's Changed in Release 20.3R1 | 160](#)



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

<a href="#">What's New   158</a>
<a href="#">Known Limitations   160</a>
<a href="#">Open Issues   161</a>
<a href="#">Resolved Issues   163</a>
<a href="#">Documentation Updates   164</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   165</a>

## Known Limitations

### IN THIS SECTION

- [Interfaces | 160](#)

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

### Interfaces

- On NFX150-S1 devices, the MTU for FPC1 interfaces mapped to ovs and sxe interfaces is 1500. [PR1488541](#)

### SEE ALSO



<a href="#">What's New</a>	<a href="#">158</a>
<a href="#">What's Changed</a>	<a href="#">159</a>
<a href="#">Open Issues</a>	<a href="#">161</a>
<a href="#">Resolved Issues</a>	<a href="#">163</a>
<a href="#">Documentation Updates</a>	<a href="#">164</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">165</a>

## Open Issues

### IN THIS SECTION

- [High Availability](#) | [161](#)
- [Interfaces](#) | [161](#)
- [Platform and Infrastructure](#) | [162](#)
- [Virtual Network Functions \(VNFs\)](#) | [162](#)

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.

### High Availability

- For an NFX250 chassis cluster, MAC learning should be disabled on fabric VLANs. We also recommend that you have only one L2 and L3 interface per node as part of the fabric VLAN. [PR1495188](#)

### Interfaces

- When you issue a **show interface** command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. [PR1306191](#)
- On NFX150 and NFX250 NextGen devices, when you add, modify, or delete a VNF interface that is mapped to an L2 or L3 data plane, kernel traces might be observed on the NFX Series device console. [PR1435361](#)



- An error message is not displayed when you configure the **native-vlan-id** option on an access VNF interface though the commit fails. [PR1438854](#)
- The link disable option puts the analyzer interface in an inconsistent state with link state as DOWN and admin state as UP. [PR1442224](#)
- 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

- On NFX150 devices, throughput degradation is noticed in RIOT-OVS-Fortigate-OVS-FlowD and RIOT-OVS-FlowD-OVS-Fortigate-OVS-FlowD cases. [PR1518939](#)
- Login access to JDM through TACACS failed after upgrade to Junos OS Release 18.4R3.  
As a workaround, log in as a local user. [PR1504915](#)
- On NFX150 devices, MAP-E customer edge (CE) configurations do not perform validation to check whether the suffix part is nonzero. The configuration must ensure that the suffix part of configurations involving MAP-E prefixes consists of zeros. [PR1457927](#)
- If you plug an unsupported SFP-T transceiver into an NFX150 device and reboot the device, the FPC1 WAN port does not come online. [PR1411851](#)
- Jumbo frames are not supported through OVS on an NFX250 device. [PR1420630](#)
- Potential security vulnerabilities in Intel firmware that is used in the NFX150 network services platform may allow escalation of privilege, denial of service, or information disclosure. Intel has released firmware updates to mitigate these potential vulnerabilities. For more information, see [NFX150: Multiple vulnerabilities in BIOS firmware \(INTEL-SA-00241\)](#). [PR1480976](#)

## Virtual Network Functions (VNFs)

- On NFX Series devices, while configuring **vmhost vlans** using **vlan-id-list**, the system allows duplicate VLAN IDs in the VLAN ID list. [PR1438907](#).

## SEE ALSO

[What's New | 158](#)

[What's Changed | 159](#)

[Known Limitations | 160](#)



---

[Resolved Issues | 163](#)

---

[Documentation Updates | 164](#)

---

[Migration, Upgrade, and Downgrade Instructions | 165](#)

## Resolved Issues

### IN THIS SECTION

- [Application Security | 163](#)
- [High Availability | 163](#)
- [Interfaces | 163](#)
- [Platform and Infrastructure | 164](#)

Learn which issues were resolved in the Junos OS Release 20.3R1 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.

### 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

<a href="#">What's New   158</a>
<a href="#">What's Changed   159</a>
<a href="#">Known Limitations   160</a>
<a href="#">Open Issues   161</a>
<a href="#">Documentation Updates   164</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   165</a>

## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for NFX Series devices.

### SEE ALSO

<a href="#">What's New   158</a>
<a href="#">What's Changed   159</a>
<a href="#">Known Limitations   160</a>
<a href="#">Open Issues   161</a>
<a href="#">Resolved Issues   163</a>



## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 165
- Basic Procedure for Upgrading to Release 20.3 | 165

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

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

### Upgrade and Downgrade Support Policy for Junos OS Releases

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

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

### Basic Procedure for Upgrading to Release 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 20.3R1:

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 | 158](#)[What's Changed | 159](#)[Known Limitations | 160](#)[Open Issues | 161](#)[Resolved Issues | 163](#)[Documentation Updates | 164](#)

## Junos OS Release Notes for PTX Series

### IN THIS SECTION

- [What's New | 168](#)
- [What's Changed | 180](#)
- [Known Limitations | 183](#)
- [Open Issues | 186](#)
- [Resolved Issues | 187](#)
- [Documentation Updates | 190](#)
- [Migration, Upgrade, and Downgrade Instructions | 190](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).



## What's New

### IN THIS SECTION

- [Hardware | 168](#)
- [Authentication, Authorization, and Accounting | 169](#)
- [IP Tunneling | 169](#)
- [Juniper Extension Toolkit \(JET\) | 170](#)
- [Junos OS XML, API, and Scripting | 171](#)
- [Junos Telemetry Interface | 171](#)
- [Layer 3 Features | 175](#)
- [MPLS | 175](#)
- [Network Management and Monitoring | 175](#)
- [Port Security | 176](#)
- [Routing Protocols | 177](#)
- [Segment Routing | 178](#)
- [Services Applications | 178](#)
- [Software Defined Networking \(SDN\) | 179](#)

Learn about new features introduced in Junos OS Release 20.3R1 for the PTX Series.

### 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

- **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)

- **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
/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



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/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
/network-instances/network-instance/afts/next-hops/next-hop/juniper/state/lsp-id
This leaf is a new augmentation.
/network-instances/network-instance/afts/next-hops/next-hop/state/ip-address
/network-instances/network-instance/afts/next-hops/next-hop/state/mac-address



Table 6: Supported Sensors (continued)

Supported Sensors
<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](#).]

## Software Defined Networking (SDN)

- **Support for static FTI backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (MX Series, PTX Series, and QFX10002)**—Starting in Junos OS Release 20.3R1, we've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to implement underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation on flexible tunnel interfaces (FTIs). Use this feature to engineer effective, loop-free backup paths for core transport networks built with only IP protocols for fast restoration after failures.

We've extended FTIs and existing forwarding constructs to support configuring static IPv4 IP-in-IP tunnels. You can also allow policy matches for routes injected by JET APIs.

[See [policy-statement](#), [tunnel](#), [ipip](#), [show interfaces](#), [show route](#), [Configuring Flexible Tunnel Interfaces](#), and [JET APIs on the Juniper Engineering Network website](#).]

## SEE ALSO

[What's Changed | 180](#)

[Known Limitations | 183](#)

[Open Issues | 186](#)



[Resolved Issues | 187](#)[Documentation Updates | 190](#)[Migration, Upgrade, and Downgrade Instructions | 190](#)

## What's Changed

### IN THIS SECTION

- [Class of Service \(CoS\) | 181](#)
- [General Routing | 181](#)
- [High Availability \(HA\) and Resiliency | 182](#)
- [Junos OS XML, API, and Scripting | 182](#)
- [Juniper Extension Toolkit \(JET\) | 182](#)
- [MPLS | 182](#)
- [Routing Protocols | 183](#)
- [System Management | 183](#)

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 20.3R1 for the PTX Series.



## 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.

## System Management

- **Support for exclude option under file archive (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—The **exclude** option is added under the command **file archive** that specifies the file pattern to exclude. This option helps to exclude files that delay compression or files that do not require compression.

[See [file archive](#).]

### SEE ALSO

<a href="#">What's New   168</a>
<a href="#">Known Limitations   183</a>
<a href="#">Open Issues   186</a>
<a href="#">Resolved Issues   187</a>
<a href="#">Documentation Updates   190</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   190</a>

## Known Limitations

### IN THIS SECTION

- [General Routing | 184](#)
- [MPLS | 185](#)
- [Routing Protocols | 185](#)



Learn about known limitations in Junos OS Release 20.3R1 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](#)

## 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 | 168](#)

[What's Changed | 180](#)

[Open Issues | 186](#)

[Resolved Issues | 187](#)

[Documentation Updates | 190](#)

[Migration, Upgrade, and Downgrade Instructions | 190](#)



## Open Issues

### IN THIS SECTION

- General Routing | [186](#)
- MPLS | [187](#)
- Routing Protocols | [187](#)

Learn about open issues in the Junos OS Release 20.3R1 for the PTX Series.

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

### General Routing

- On a PTX Series platform with the FPC-PTX-P1-A or FPC2-PTX-P1A line card, you might encounter a single event upset (SEU) that can cause a linked-list corruption of the TQ chip. The following syslog error message is reported: **Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt\_min\_free\_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002.** Junos OS chassis management error handling detects such a condition, raises an alarm, and performs the disable-pfe action for the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Soft errors are transient or non-recurring. FPCs experiencing such SEU events do not have any permanent damage. Contact your Juniper Networks support representative if the issue persists even after an FPC restart. [PR1254415](#)
- PTX Series platform drops the wireless access point (WAP) heartbeat packets; as a result, the WAP cannot work. [PR1352805](#)
- The firewall counter for lo0 interface might not increase. As a workaround, set the lo0 filter family inet and family inet6 counters instead of the filter family any. [PR1420560](#)
- Mirrored packets are corrupted when filter is applied with action port-mirror and discarded. [PR1437546](#)
- On PTX1000 and PTX10001 platforms, the port mirror does not work when port mirroring is configured with firewall filter. [PR1491789](#)
- At low timeout values, the flows might not reach the maximum supported scale of 1.2 million flows. Lower timeout configuration increases the number of flow timeouts, resulting in increased load on CPU



for both multi-svcs and uKern processes, and affects the flow creation. We recommend that you configure timeouts above 60 seconds to create the flows successfully. [PR1510150](#)

## MPLS

- At high scale, LSP setup rate might be relatively slower in IPinIP networks. [PR1457992](#)

## Routing Protocols

- During an FRR event, if the backup path is inet table lookup (with backup-ip-forward configuration), then per sid-stats might not work as expected on PTX Series platforms. Traffic loss during FRR switchover is more than 50 ms on some occasions. [PR1491765](#)

### SEE ALSO

[What's New | 168](#)

[What's Changed | 180](#)

[Known Limitations | 183](#)

[Resolved Issues | 187](#)

[Documentation Updates | 190](#)

[Migration, Upgrade, and Downgrade Instructions | 190](#)

## Resolved Issues

### IN THIS SECTION

- [General Routing | 188](#)
- [Interfaces and Chassis | 189](#)
- [MPLS | 189](#)
- [Network Management and Monitoring | 189](#)
- [Routing Protocols | 189](#)

Learn which issues were resolved in the Junos OS Release 20.3R1 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 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 pppmd 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 | 168](#)[What's Changed | 180](#)[Known Limitations | 183](#)[Open Issues | 186](#)[Documentation Updates | 190](#)[Migration, Upgrade, and Downgrade Instructions | 190](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for PTX Series routers.

## SEE ALSO

[What's New | 168](#)[What's Changed | 180](#)[Known Limitations | 183](#)[Open Issues | 186](#)[Resolved Issues | 187](#)[Migration, Upgrade, and Downgrade Instructions | 190](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.3 | 191](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 193](#)
- [Upgrading a Router with Redundant Routing Engines | 194](#)



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 20.3R1:

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-x86-64-20.3R1.9.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-x86-64-20.3R1.9-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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from



Junos OS Release 19.2 to Release 19.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

## Upgrading a Router with Redundant Routing Engines

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

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

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

### SEE ALSO

---

[What's New | 168](#)

---

[What's Changed | 180](#)

---

[Known Limitations | 183](#)

---

[Open Issues | 186](#)

---

[Resolved Issues | 187](#)

---

[Documentation Updates | 190](#)



# Junos OS Release Notes for QFX Series

## IN THIS SECTION

- What's New | 195
- What's Changed | 209
- Known Limitations | 212
- Open Issues | 214
- Resolved Issues | 218
- Documentation Updates | 223
- Migration, Upgrade, and Downgrade Instructions | 224

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- Hardware | 197
- EVPN | 198
- High Availability (HA) and Resiliency | 201
- IP Tunneling | 201
- Juniper Extension Toolkit (JET) | 202
- Junos OS XML, API, and Scripting | 203
- Junos Telemetry Interface | 203
- Layer 3 Features | 204
- MPLS | 204
- Network Management and Monitoring | 204



- Routing Policy and Firewall Filters | 206
- Routing Protocols | 206
- Security | 207
- Services Applications | 207
- Software Defined Networking (SDN) | 207
- Software Licensing | 207
- Virtual Chassis | 208

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 20.3R1: 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.



## 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"> <li>• 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 <a href="#">Overview of MPLS Firewall Filters on Loopback Interface</a>.]</li> <li>• 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 <a href="#">Firewall Filter Flexible Match Conditions</a>.]</li> </ul>
Timing and synchronization	<ul style="list-style-type: none"> <li>• Precision Time Protocol (PTP) transparent clock is supported on the QFX5120-48T. [See <a href="#">Transparent Clock Overview</a>.]</li> </ul>

- **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](#).]

## 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 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 Defined Networking (SDN)

- **Support for static FTI backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (MX Series, PTX Series, and QFX10002)**—Starting in Junos OS Release 20.3R1, we've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to implement underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation on flexible tunnel interfaces (FTIs). Use this feature to engineer effective, loop-free backup paths for core transport networks built with only IP protocols for fast restoration after failures.

We've extended FTIs and existing forwarding constructs to support configuring static IPv4 IP-in-IP tunnels. You can also allow policy matches for routes injected by JET APIs.

[See [policy-statement](#), [tunnel](#), [ipip](#), [show interfaces](#), [show route](#), [Configuring Flexible Tunnel Interfaces](#), and [JET APIs on the Juniper Engineering Network website](#).]

## 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 208](#) 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	<b>Advanced 1:</b> BGP, FBF, GRE, IS-IS, JTI, MC-LAG, OSPF, sFlow, VRF, and VRRP
	<b>Advanced 2:</b> Includes <b>Advanced 1</b> 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 <b>Advanced 2</b> 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

<a href="#">What's Changed</a>	<a href="#">  209</a>
<a href="#">Known Limitations</a>	<a href="#">  212</a>
<a href="#">Open Issues</a>	<a href="#">  214</a>
<a href="#">Resolved Issues</a>	<a href="#">  218</a>
<a href="#">Documentation Updates</a>	<a href="#">  223</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#">  224</a>

# What's Changed

IN THIS SECTION

- [Class of Service \(CoS\)](#) | [210](#)
- [General Routing](#) | [210](#)
- [High Availability \(HA\) and Resiliency](#) | [211](#)
- [Interfaces and Chassis](#) | [211](#)
- [Junos OS XML, API, and Scripting](#) | [211](#)
- [Routing Protocols](#) | [211](#)



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

## 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.
- **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**.



## 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 | 195](#)[Known Limitations | 212](#)[Open Issues | 214](#)[Resolved Issues | 218](#)[Documentation Updates | 223](#)[Migration, Upgrade, and Downgrade Instructions | 224](#)

## Known Limitations

### IN THIS SECTION

- [Layer 2 Ethernet Services | 212](#)
- [Platform and Infrastructure | 213](#)
- [Routing Protocols | 213](#)

Learn about known limitations in Junos OS Release 20.3R1 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.

### Layer 2 Ethernet Services

- If a configuration or image filename has nonallowed special characters (such as #, %, and @) in it, ZTP over HTTP/HTTPS might not work. When HTTP/HTTPS URL is formed to download the file, the URL contains the filename in it. HTTP/HTTPS does not expect any special characters in the URL. If special characters are present, the HTTP/HTTPS protocol returns "Bad request". In order to avoid the issue, do not use any nonallowed special characters in the filename. [PR1503588](#)
- If you configure image and script on the DHCP server as part of the ZTP, DHCPv6 client binding does not happen after image upgrade and reboot of image and script. [PR1532304](#)



## Platform and Infrastructure

- The 100-Gigabit Ethernet interface goes down after you configure and delete the Ethernet loopback configuration. [PR1353734](#)
- On QFX10000 line of devices, if the analyzer is configured to a mirror traffic of an input aggregated Ethernet interface and a new member is added to the same aggregated Ethernet interface, then the analyzer might not provide sample packets that flow through the newly added child interface. [PR1417694](#)
- TCAM calculation issue is found in Junos OS Release 18 and Junos OS Release 19 codes after introducing the IPACL VXLAN filters. [PR1469515](#)
- On QFX Series platform with **set routing-options resolution preserve-nexthop-hierarchy** statement configured, reaching tunnel destination out-going route via BGP-over-BGP route recursive resolution is not supported. [PR1498085](#)
- On a fully scaled system where all the slices are utilized by different families of CLI filters, if you try to delete one family and change another family with higher number of filter terms that requires expansion of the filter, the Packet Forwarding Engine fails to add the new changed filter as out of sequence messages are generated, that is, change of filter is called earlier than deletion of another filter. [PR1512242](#)

## Routing Protocols

- Higher than expected loss and traffic drops and discards silently during node failures with node protection on FTI interfaces for RSVP LSPs. [PR1456350](#)
- Third party vendor SDK does not support variable mask for destination IP address in tunnel termination table, so firewall terms for de-encapsulation action should always have destination address as /32 address. Source IP address can be variable mask or optional. [PR1511893](#)
- On the QFX5200 switch, hierarchical ECMP supports only two levels of ECMP. If a BGP route is resolved over multiple PRPD routes, we've unilist 1 (BGP route), unilist 2 (PRPD route), ucast1, ucast 2 (FTI underlay can be unilist). There will be three unilist in the hierarchy. Because of this, we've flattened the unilist paths for FTI NHs as well. Traffic distribution behavior is the same across QFX5100 and QFX5200. Hierarchal ECMP is not supported on QFX5100, even for routes pointing to non-FTI. For unilist of unilist, QFX5100 flattens all the unicast paths, so traffic will be equally distributed to the final list of unicast NHs, not at top-level unilist. [PR1517519](#)

## SEE ALSO

[What's New | 195](#)

[What's Changed | 209](#)

[Open Issues | 214](#)



---

[Resolved Issues | 218](#)

---

[Documentation Updates | 223](#)

---

[Migration, Upgrade, and Downgrade Instructions | 224](#)

---

## Open Issues

### IN THIS SECTION

- [High Availability \(HA\) and Resiliency | 214](#)
- [Layer 2 Features | 214](#)
- [Platform and Infrastructure | 215](#)
- [Routing Protocols | 217](#)
- [Virtual Chassis | 218](#)

Learn about open issues in Junos OS Release 20.3R1 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.

### High Availability (HA) and Resiliency

- The QFX5200-32C reboot time is degraded. A flush cache issue is seen because of the reliable SSD disk input/output change made for this platform. [PR1511607](#)

### Layer 2 Features

- In case of QFX5000 Virtual Chassis or Virtual Chassis Fabric setups, when IGMP snooping is enabled, multicast traffic is forwarded based on IGMP joins or reports. But when the IGMP report times out, traffic should be dropped; instead it floods in the VLAN. This happens only in case of QFX5000 Virtual Chassis or Virtual Chassis Fabric; this issue is not seen on stand-alone QFX5000 switches. [PR1431893](#)
- On QFX5110 and QFX5120 platforms, changing lo0 IP address might sometimes result either in stale entry of IP in mpls\_entry table or in a missing IP entry, which results in traffic drop for VXLAN traffic. [PR1472333](#)
- QFX5100 switches that send routed traffic (either transit or locally originated) out to an interface configured for Q-in-Q traffic fail to correctly add the two VLAN tags. [PR1481648](#)



## Platform and Infrastructure

- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED remains unlit. [PR1317750](#)
- The QFX10000 line of devices drop the wireless access point (WAP) heartbeat packets; as a result, the WAP cannot work. [PR1352805](#)
- USB upgrade of NOS image is not supported. [PR1373900](#)
- On QFX5110 and QFX5120 platforms, unicast RPF check in strict mode might not work properly. [PR1417546](#)
- The issue occurs because of a PECHIP limitation when underlay is tagged. After de-encapsulation, when the inner packet is recirculated, it retains the VLAN tag property from outer header because the outer header was tagged. Thus 4 bytes of inner tag is overwritten in the inner packet and the packet got corrupted, which results in EGP chksum trap seen in PECHIP. Fixing PECHIP limitation in software has high risk. As a workaround, enable **encapsulate-inner-vlan** configuration. [PR1435864](#)
- The unified ISSU is not supported on QFX5200 switches and fails from Junos OS Release 17.2X75-D43.2 to some target versions. Also, dcpfe crash might be seen. [PR1438690](#)
- On the QFX10000 line of devices, in an EVPN-VXLAN (spine-leaf) scenario, the QFX10000 spine switches are configured with VXLAN Layer 3 gateway (utilizing the virtual gateway) on an IRB interface. If you enable and then subsequently remove the VXLAN Layer 3 gateway on this IRB interface on one or some of these spine switches, traffic drop might be observed. As a workaround, configure all virtual gateways with unique IPv4 or IPv6 MAC address. [PR1446291](#)
- After changing the VLAN name on the trunk interface, while port is receiving continuous traffic for that VLAN, local host MAC learning holds for more than 30 seconds. In case of trunk port, when VLAN name is changed, bridge domain entry is deleted from hardware and a new entry is installed in the hardware. When the new entry is yet to be installed in hardware, port keeps receiving traffic for that VLAN and learn source MAC and notifies to Packet Forwarding Engine with old bridge domain ID. When Packet Forwarding Engine software receives this MAC drops it as bridge domain and port mapping will not be present in software which is a must criteria for a source MAC received on a bridge domain. Once Packet Forwarding Engine drops the MAC, upper layers (L2ALD) does not get this MAC info and aging thread marks the hash index in hardware as stale. Until that hash index is not cleared in the hardware, same source MAC cannot be learnt on the same hash index. Ageing thread periodically scans one MAC table out of 4 tables at a time in intervals of 10 seconds and checks for stale entries and clear the hardware hash stale entry, and this time is almost 40-50 seconds based on the number of Packet Forwarding Engine chips in an FPC. In case of access port, default bridge domain is installed in the hardware to receive untagged traffic and does not get deleted while changing VLAN name associated to that access port. So this issue is not seen for access port. [PR1454274](#)
- On QFX5110 switches, VXLAN VNI (mcast) scaling traffic issue from VXLAN tunnel to Layer 2 interface is observed. [PR1462548](#)



- BGP route addition and deletion time and BGP, OSPF, and IS-IS link flap convergence time are increased. [PR1464572](#)
- Dynamic IP-over-IP tunnels and filter-based IP-over-IP de-encapsulation filter on loopback interface cannot coexist together. If dynamic IP-over-IP tunnels were configured earlier, then FPC needs a reboot before it can be used for loopback IP-over-IP de-encapsulation filter. Also, the loopback interface might contain implicit filter. If these implicit filters get hit, the de-encapsulation filter might not get hit. [PR1479613](#)
- On QFX Series platforms running Junos VM instance (excluding QFX10000 Series platforms), the laser signal might be transmitted on the disabled interfaces with QSFP and QSFP28 optics after device reboot. [PR1487554](#)
- If the interface is newly added as a CE interface, the existing broadcast, unicast, and multicast (BUM) traffic can be looped. The loop prevention feature is designed to start working whenever a new CE interface is added by configuration. But the existing BUM traffic can be distributed to a new CE interface earlier, before enabling the loop prevention feature. [PR1493650](#)
- Storage full message is displayed for the /var/tmp directory when the file is copied. The **file copy** command uses a default staging-directory (a temporary directory) and this staging directory is /var/tmp for the root user and /var/home/<user-dir> for all non-root users. If you face storage system full issues while copying the file, you can use the **staging-directory** command line option to choose a different staging directory (other than the default). [PR1494489](#)
- On the QFX5210 switches, unexpected behavior for port LEDs lights is observed after the upgrade. [PR1498175](#)
- In the l2circuit termination scenario with input-vlan-map/output-vlan-map and family ccc, the output-vlan-map push operation might not work. It has a traffic impact. [PR1510629](#)
- In an EVPN-VXLAN scenario, multicast traffic might not reach to spine to form (S,G) in PIM enabled spines. Issue might happen due to various triggers including multiple rollback of configurations on spine, interface flap, and clear BGP. [PR1510794](#)
- On QFX5100 Virtual Chassis, degradation is observed at the time of system reboot and FPC online. [PR1513540](#)
- On QFX5000 platforms with QFX-5e image, if the 100-Gigabit Ethernet port is enabled with auto-channelization (which is by default) and the AOC non-breakout transceiver is used on it, the 100-Gigabit Ethernet port might be detected as breakout and auto-channelized to other speed (for example, 50-Gigabit Ethernet). This impacts the interface connection. [PR1515487](#)
- Disruptive switchover (no GRES or NSR configured) can lead to stale PPM entries programmed on the new master Routing Engine. If both GRES and NSR are activated after disruptive switchover, and then a GRES is performed, BFD sessions might flap continuously. [PR1518106](#)
- Some inter VLAN traffic flows do not converge after rebooting spine device (QFX10002) in EVPN VXLAN non-collapsed scaled scenario when traffic is already flowing. [PR1522585](#)



- On the QFX5120 platform, when there is a traffic that is being routed between two VXLANs from an access interface of one VXLAN to an access interface of another VXLAN, traffic might get untagged. [PR1527939](#)
- On the QFX5000 platform, aggregated Ethernet port shut down with storm control trigger shall not be up after recovery timeout. [PR1534642](#)
- 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](#)
- Traffic loss is observed when multicast over GRE is configured. [PR1536886](#)
- If the Layer 3 filter is applied on a VXLAN IRB, then VXLANs and IRBs are deactivated in a separate commits, IRB filters attached does not go away. [PR1537108](#)
- With an EVPN VXLAN configuration, when restart of I2-learning command is executed, BFD sessions on IRB interface might not come up. [PR1538600](#)
- False management Ethernet link down alarm is seen. [PR1538674](#)
- Inter VRF traffic drop might be seen in QFX10008 platform with EVPN VXLAN configurations. [PR1540200](#)
- Under certain circumstances, inactive VTEP timers are not stopped though there is MAC posted for learning, resulting in VTEP deletion and total loss of traffic. As a workaround, restart I2-learning within 5 minutes of RVTEP creation. [PR1540208](#)
- On the QFX5000 platform, inter VLAN traffic drop might be seen on a leaf node when EVPN-VXLAN is configured along with the LAG. [PR1541406](#)

## Routing Protocols

- If DDoS protection is disabled on QFX5100 Virtual Chassis and multicast traffic is being sent, the Virtual Chassis might become unstable, with high CPU usage and it might crash eventually, creating FXPC core files. Disabling DDoS protection will disable rate limiting for all host-bound traffic. We do not recommend disabling DDoS protection on the device, because a high amount of control traffic can overwhelm the system, causing system instability. [PR1238875](#)
- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, when the **mini-PDT-base** configuration is issued, the following error message is seen in the hardware: **BRCM\_NH-,brcm\_nh\_bdvlan\_ucast\_uninstall(), 128:I3 nh 6594 uninstall failed**. There is no functionality impact because of this error message. [PR1407175](#)
- IPv6 routes pointing to the BGP LU path are not programmed in hardware with the **preserve-nexthop-hierarchy** statement. This results in traffic drop for IPv6 routes. The **preserve-nexthop-hierarchy** statement is required for IP-in-IP, in this case, the IPv6 route pointing to an IP tunnel has no issues. [PR1510053](#)
- On the QFX10000 line of devices, if multiple sub-interfaces of the same aggregated Ethernet interface belong to different routing instances, and these sub-interfaces are configured with the same IP address



and separate BFD sessions, the remaining BFD sessions flap continuously if one of these BFD sessions is deleted. [PR1516556](#)

- On the QFX5120 platform, when traffic is being routed between two VXLANs from an access interface of one VXLAN to an access interface of another VXLAN, traffic might go out untagged. [PR1536608](#)

### Virtual Chassis

- The QFX5110-48S reports false parity error messages like `soc_mem_array_sbusdma_read` and SDK can raise false alarms for such parity error messages. This is a false positive error message. [PR1276970](#)
- On QFX5000 Virtual Chassis, DDoS violations that happen on the backup are not reported to Routing Engine. [PR1490552](#)

### SEE ALSO

<a href="#">What's New</a>	<a href="#"> </a>	<a href="#">195</a>
<a href="#">What's Changed</a>	<a href="#"> </a>	<a href="#">209</a>
<a href="#">Known Limitations</a>	<a href="#"> </a>	<a href="#">212</a>
<a href="#">Resolved Issues</a>	<a href="#"> </a>	<a href="#">218</a>
<a href="#">Documentation Updates</a>	<a href="#"> </a>	<a href="#">223</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions</a>	<a href="#"> </a>	<a href="#">224</a>

## Resolved Issues

### IN THIS SECTION

- [Class of Service \(CoS\)](#) | [219](#)
- [EVPN](#) | [219](#)
- [Infrastructure](#) | [219](#)
- [Interfaces and Chassis](#) | [219](#)
- [Layer 2 Features](#) | [220](#)
- [Layer 2 Ethernet Services](#) | [220](#)
- [MPLS](#) | [220](#)
- [Platform and Infrastructure](#) | [220](#)



- Routing Protocols | 222
- User Interface and Configuration | 223

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.

### 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 OVSDB 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

<a href="#">What's New   195</a>
<a href="#">What's Changed   209</a>
<a href="#">Known Limitations   212</a>
<a href="#">Open Issues   214</a>
<a href="#">Documentation Updates   223</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   224</a>

Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for the QFX Series Switches.

SEE ALSO

<a href="#">What's New   195</a>
<a href="#">What's Changed   209</a>
<a href="#">Known Limitations   212</a>



[Open Issues | 214](#)[Resolved Issues | 218](#)[Migration, Upgrade, and Downgrade Instructions | 224](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 224](#)
- [Installing the Software on QFX10002-60C Switches | 227](#)
- [Installing the Software on QFX10002 Switches | 227](#)
- [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 | 228](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 230](#)
- [Performing a Unified ISSU | 234](#)
- [Preparing the Switch for Software Installation | 235](#)
- [Upgrading the Software Using Unified ISSU | 235](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 237](#)

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-x86-64-20.3R1.9.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-x86-64-20.3R1.9.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-x86-64-20.3R1.n-secure-signed.tgz reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> reboot** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-20.3R1.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

**Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches**



**NOTE:** Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add <pathname><source>** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add <pathname><source> re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```



After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

### Installing the Software on QFX10008 and QFX10016 Switches



Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

**NOTE:** Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at <https://www.juniper.net/support>.



**WARNING:** If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

```
user@switch# delete chassis redundancy
```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```



After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.3R1.n-secure-signed.tgz
```

For more information about the **request system software add** command, see the [CLI Explorer](#).

9. Reboot the switch to start the new software using the **request system reboot** command:

```
user@switch> request system reboot
```

**NOTE:** You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete <package-name>** command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

```
user@switch> show version
```

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.



11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
Slot 0:
  Current state           Backup
  Election priority       Master (default)
Routing Engine status:
Slot 1:
  Current state           Master
  Election priority       Backup (default)
```

14. Install the new software package using the **request system software add** command:

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.3R1.n-secure-signed.tgz
```

For more information about the **request system software add** command, see the [CLI Explorer](#).



15. Reboot the Routing Engine using the **request system reboot** command:

```
user@switch> request system reboot
```

**NOTE:** You must reboot to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your system. Instead, finish the installation and then issue the **request system software delete jinstall <package-name>** command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the **show version** command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the [CLI Explorer](#).

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
  Slot 0:
    Current state           Master
    Election priority       Master (default)
Routing Engine status:
  Slot 1:
    Current state           Backup
    Election priority       Backup (default)
```

## Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.



**NOTE:** Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 235](#)
- [Upgrading the Software Using Unified ISSU on page 235](#)

## 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-x86-64-20.3R1.n-secure-signed.tgz*.

**NOTE:** During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-host-qfx-5-f-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R1.n-secure-signed.tgz completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
```



```

ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item                Status                Reason
  FPC 0                Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff

```

**NOTE:** A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

**NOTE:** If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

## Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases



provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.

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

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

#### SEE ALSO

<a href="#">What's New   195</a>
<a href="#">What's Changed   209</a>
<a href="#">Known Limitations   212</a>
<a href="#">Open Issues   214</a>
<a href="#">Resolved Issues   218</a>
<a href="#">Documentation Updates   223</a>

## Junos OS Release Notes for SRX Series

#### IN THIS SECTION

- [What's New | 239](#)
- [What's Changed | 249](#)
- [Known Limitations | 251](#)
- [Open Issues | 252](#)
- [Resolved Issues | 254](#)
- [Documentation Updates | 258](#)
- [Migration, Upgrade, and Downgrade Instructions | 258](#)



These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- Application Security | 240
- Authentication and Access Control | 240
- Chassis Clustering | 240
- Flow-Based and Packet-Based Processing | 241
- Installation and Upgrade | 241
- Interfaces and Chassis | 241
- Intrusion Detection and Prevention (IDP) | 242
- Junos Telemetry Interface | 242
- Junos OS XML API and Scripting | 243
- J-Web | 243
- Layer 2 Features | 245
- Logical Systems and Tenant Systems | 245
- Network Management and Monitoring | 245
- Routing and Forwarding Options | 247
- Security | 247
- Unified Threat Management (UTM) | 247
- VPNs | 247

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



## 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 on new IKED platforms (SRX5000 line of devices with SPC3 and vSRX)**—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 \(ESN\)](#).]

SEE ALSO

<a href="#">What's Changed   249</a>
<a href="#">Known Limitations   251</a>
<a href="#">Open Issues   252</a>
<a href="#">Resolved Issues   254</a>
<a href="#">Documentation Updates   258</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   258</a>



## What's Changed

### IN THIS SECTION

- Authentication and Access Control | 249
- Junos OS XML API and Scripting | 249
- J-Web | 250
- Network Address Translation (NAT) | 250
- System Logs | 250

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

### 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 \(System Logging\).](#)]

SEE ALSO

<a href="#">What’s New   239</a>
<a href="#">Known Limitations   251</a>
<a href="#">Open Issues   252</a>
<a href="#">Resolved Issues   254</a>



Documentation Updates | [258](#)

Migration, Upgrade, and Downgrade Instructions | [258](#)

## Known Limitations

### IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 251](#)
- [J-Web | 251](#)
- [VPNs | 252](#)

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.

### Flow-Based and Packet-Based Processing

- Due to internal message failures between 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 `show services application-identification commit-status`. [PR1493127](#)

### J-Web

- For a spoke device in a hub-and-spoke topology, J-Web shows the VPN topology as Site to Site. [PR1495973](#)
- The J-Web IPsec VPN workflow only supports route-based VPNs. Policy-based VPNs are not supported. [PR1498169](#)



# VPNs

- When multiple traffic selectors are configured on a particular VPN, the iked process checks for a maximum of 1 DPD probe that is sent to the peer for the configured DPD interval. The DPD probe will be 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 st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)

## SEE ALSO

<a href="#">What's New   239</a>
<a href="#">What's Changed   249</a>
<a href="#">Open Issues   252</a>
<a href="#">Resolved Issues   254</a>
<a href="#">Documentation Updates   258</a>
<a href="#">Migration, Upgrade, and Downgrade Instructions   258</a>

## Open Issues

### IN THIS SECTION

- [J-Web | 253](#)
- [VPNs | 253](#)

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.



## J-Web

- Configuration of global setting options of IPsec VPN such as TCP encapsulation profile, IPsec power mode, and IKE package installation are not supported from J-Web. [PR1496439](#)
- 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 SRX5000 line of devices, J-Web can take up to 60 seconds to 90 seconds to load 60,000 security policies. [PR1521841](#)
- J-Web does not support disabling or enabling the security firewall or global policy rules. The policy rules that are deactivated through CLI are also not visible in the J-Web UI. [PR1522128](#)

## VPNs

- On the SRX5000 line of 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](#)
- In an IPsec VPN scenario on the SRX5000 line of devices, theiked process treats retransmission of IKE\_INIT request packets as new connections when the SRX Series device acts as a responder of IKE negotiation. This causes IKE tunnel negotiation to fail, and IPsec VPN traffic might be impacted. [PR1460907](#)
- 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](#)

## SEE ALSO

[What's New | 239](#)

[What's Changed | 249](#)

[Known Limitations | 251](#)

[Resolved Issues | 254](#)

[Documentation Updates | 258](#)

[Migration, Upgrade, and Downgrade Instructions | 258](#)



## 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.

### Application Security

- AppQoS 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 srpxfe 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 srxpfe 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](#)

## SEE ALSO

---

[What's New | 239](#)

---

[What's Changed | 249](#)

---

[Known Limitations | 251](#)

---



---

[Open Issues | 252](#)

---

[Documentation Updates | 258](#)

---

[Migration, Upgrade, and Downgrade Instructions | 258](#)

---

## Documentation Updates

There are no errata or changes in Junos OS Release 20.3R1 documentation for the SRX Series.

### SEE ALSO

---

[What's New | 239](#)

---

[What's Changed | 249](#)

---

[Known Limitations | 251](#)

---

[Open Issues | 252](#)

---

[Resolved Issues | 254](#)

---

[Migration, Upgrade, and Downgrade Instructions | 258](#)

---

## Migration, Upgrade, and Downgrade Instructions

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

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

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4.



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

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

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

For information about ISSU, see the [Chassis Cluster User Guide for Security Devices](#).

#### SEE ALSO

<a href="#">What's New   239</a>
<a href="#">What's Changed   249</a>
<a href="#">Known Limitations   251</a>
<a href="#">Open Issues   252</a>
<a href="#">Resolved Issues   254</a>
<a href="#">Documentation Updates   258</a>

## Junos OS Release Notes for vMX

#### IN THIS SECTION

- [What's New | 260](#)
- [Open Issues | 261](#)
- [Resolved Issues | 262](#)
- [Licensing | 262](#)
- [Upgrade Instructions | 263](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

### IN THIS SECTION

- [Installation and Upgrade](#) | 260
- [Juniper Extension Toolkit \(JET\)](#) | 260
- [System Management](#) | 261

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

### 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.](#)]

## Open Issues

### IN THIS SECTION

- [General Routing | 262](#)

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.



## General Routing

- 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](#)

## 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.

## Platform and Infrastructure

- `core.vmx.mpc0` is seen at `5 0x096327d5` in `l2alm_sync_entry_in_pfes` (`context=0xd92e7b28`, `sync_info=0xd92e7a78`) at `../src/pfe/common/applications/l2alm/l2alm_common_hw_api.c:1727`. [PR1430440](#)
- On MX150 and vMX platforms, if flow caching is enabled, VXLAN packet might be discarded. This is because flow caching does not support VXLAN. [PR1466470](#)
- In vMX instances, after every commit, the following error message is displayed in the log message: **chassisd[7836]: %DAEMON-3-CHASSISD\_IOCTL\_FAILURE: acb\_get\_fpga\_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device)**. [PR1477941](#)
- On MX150 and vMX platforms, the input errors might be displayed as zero when you use the `show interfaces extensive` output when there are CRC/Align errors present on the interface. [PR1485706](#)
- Configuring the statement ranges for auto-sensed VLANs may not work on the vMX. [PR1503538](#)

## 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 | 264](#)
- [What's Changed | 265](#)
- [Known Limitations | 265](#)
- [Open Issues | 266](#)
- [Resolved Issues | 266](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).



## What's New

### IN THIS SECTION

- [Routing Protocols](#) | 264

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 20.3R1, see [What's New](#) for MX Series routers.

### 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

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 20.3R1, see [What's Changed](#) for MX Series routers.

## Known Limitations

### IN THIS SECTION

- [Routing Protocols](#) | 266

Learn about known limitations in this release for vRR.



To learn more about common BGP or routing known limitation in Junos OS 20.3R1, 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 20.3R1, 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.

## 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 20.3R1, 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.



# Junos OS Release Notes for vSRX

## IN THIS SECTION

- [What's New | 267](#)
- [Known Limitations | 271](#)
- [Open Issues | 272](#)
- [Resolved Issues | 273](#)
- [Migration, Upgrade, and Downgrade Instructions | 275](#)

These release notes accompany Junos OS Release 20.3R1 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](https://www.juniper.net/documentation/product/en_US/junos-os).

## What's New

## IN THIS SECTION

- [Interfaces and Chassis | 268](#)
- [Juniper ATP Cloud | 268](#)
- [Junos Telemetry Interface | 268](#)
- [Management | 270](#)
- [Performance and Scaling | 270](#)
- [VPNs | 270](#)

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



## 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 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](#).]

## Known Limitations

### IN THIS SECTION

- [Intrusion Detection and Prevention \(IDP\) | 272](#)
- [J-Web | 272](#)

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 | 273](#)
- [User Access and Authentication | 273](#)

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](#)

## 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.

## 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 | 276](#)
- [Validating the OVA Image | 281](#)

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 20.3R1 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 20.3R1 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 Sep 25 14:15 /var/jail/tmp/alarmd.ts
259.7K Sep 25 14:11 /var/log/hostlogs/vjunos0.log.1.gz
494B Sep 25 14:15 /var/log/interactive-commands.0.gz
20.4K Sep 25 14:15 /var/log/messages.0.gz
27B Sep 25 14:15 /var/log/wtmp.0.gz
27B Sep 25 14:14 /var/log/wtmp.1.gz
3027B Sep 25 14:13 /var/tmp/BSD.var.dist
0B Sep 25 14:14 /var/tmp/LOCK_FILE
666B Sep 25 14:14 /var/tmp/appidd_trace_debug
0B Sep 25 14:14 /var/tmp/eedebug_bin_file
34B Sep 25 14:14 /var/tmp/gksdchk.log
46B Sep 25 14:14 /var/tmp/kmdchk.log
57B Sep 25 14:14 /var/tmp/krt_rpf_filter.txt
42B Sep 25 14:13 /var/tmp/pfe_debug_commands
0B Sep 25 14:14 /var/tmp/pkg_cleanup.log.err
30B Sep 25 14:14 /var/tmp/policy_status
0B Sep 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 20.3R1 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-x86-64-20.3-2020-9-10.0_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-x86-64-20.3-2020-9-10.0_RELEASE_20.3_THROTTLE.tgz
no-copy no-validate reboot
Verified junos-vsrx-x86-64-20.3-2020-9-10.0_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-20.3-2020-9-10.0_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-20.3-2020-9-10.0_RELEASE_20.3_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31
junos-srx-mr-vsrx-20.3-2020-9-10.0_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-20.3-2020-9-10.0_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-20.3-2020-9-10.0_RELEASE_20.3_THROTTLE-linux.tgz ...
upgrade_platform: Input package
/var/tmp/junos-srx-mr-vsrx-20.3-2020-9-10.0_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-20.3-2020-9-10.0_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-20.3-2020-9-10.0_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 20.3R1 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 20.3-2020-9-10.0_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: 20.3-2020-9-10.0_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 have been developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information about the list of supported products, see [Juniper Flex Program](#).

## Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

## Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. <https://apps.juniper.net/feature-explorer/>
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. [prsearch.juniper.net](https://prsearch.juniper.net).
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](https://apps.juniper.net/hct/home)




**NOTE:** To obtain information about the components that are supported on the devices, and the special compatibility guidelines with the release, see the Hardware Guide for the product.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products. [apps.juniper.net/compliance/](https://apps.juniper.net/compliance/).

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).



# Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

## Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>



## Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit <https://myjuniper.juniper.net>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <https://support.juniper.net/support/requesting-support/>.

## Revision History

30 June 2022—Revision 17, 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 May 2022—Revision 16, 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.

5 May 2022—Revision 15, 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.

7 October 2021—Revision 14, 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.

16 September 2021—Revision 13, 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.

15 July 2021—Revision 12, 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.

3 June 2021—Revision 11, 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.

6 May 2021—Revision 10, 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.



22 April 2021—Revision 9, 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.

1 March 2021—Revision 8, 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.

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.