

# Release Notes

Published  
2021-04-01

## Junos<sup>®</sup> OS 20.2R3 Release Notes

---

### SUPPORTED ON

- ACX Series, cSRX, EX Series, JRR Series, fusion for enterprise, fusion for provider edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

# Release Notes: Junos<sup>®</sup> OS Release 20.2R3 for the ACX Series, cSRX, EX Series, JRR Series, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX

1 April 2021

<b>Contents</b>	<b>Introduction   12</b>
	<b>Junos OS Release Notes for ACX Series   12</b>
	<b>What's New   13</b>
	What's New in Release 20.2R3   13
	What's New in Release 20.2R2   13
	What's New in Release 20.2R1   14
	<b>What's Changed   22</b>
	What's Changed in Release 20.2R3   23
	What's Changed in Release 20.2R2   24
	What's Changed in Release 20.2R1-S2   25
	What's Changed in Release 20.2R1   25
	<b>Known Limitations   27</b>
	General Routing   27
	<b>Open Issues   31</b>
	General Routing   32
	Virtual Chassis   34

**Resolved Issues | 35****Resolved Issues: 20.2R3 | 35****Resolved Issues: 20.2R2 | 36****Resolved Issues: 20.2R1 | 39****Documentation Updates | 41****Migration, Upgrade, and Downgrade Instructions | 41****Upgrade and Downgrade Support Policy for Junos OS Releases | 41****Junos OS Release Notes for cSRX | 42****What's New | 43****What's New in Release 20.2R3 | 43****What's New in Release 20.2R2 | 43****What's Changed | 43****What's Changed in Release 20.2R3 | 44****What's Changed in Release 20.2R2 | 44****Known Limitations | 44****Open Issues | 44****Resolved Issues | 44****Resolved Issues: 20.2R3 | 44****Resolved Issues: 20.2R2 | 44****Junos OS Release Notes for EX Series | 44****What's New | 45****What's New in Release 20.2R3 | 45****What's New in Release 20.2R2 | 46****What's New in Release 20.2R1-S1 | 46****What's New in Release 20.2R1 | 46****What's Changed | 54****What's Changed in Release 20.2R3 | 54****What's Changed in Release 20.2R2 | 56****What's Changed in Release 20.2R1 | 56****Known Limitations | 58****EVPN | 58****General Routing | 58****Infrastructure | 58****Layer 2 Ethernet Services | 59**

**Open Issues | 59****General Routing | 60****Infrastructure | 61****Interfaces and Chassis | 62****Layer 2 Features | 62****Layer 2 Ethernet Services | 62****Platform and Infrastructure | 62****Routing Protocols | 63****Resolved Issues | 63****Resolved Issues: 20.2R3 | 63****Resolved Issues: 20.2R2 | 65****Resolved Issues: 20.2R1 | 67****Documentation Updates | 70****Migration, Upgrade, and Downgrade Instructions | 71****Upgrade and Downgrade Support Policy for Junos OS Releases | 71****Junos OS Release Notes for JRR Series | 72****What's New | 73****What's New in Release 20.2R3 | 73****What's New in Release 20.2R2 | 73****What's New in Release 20.2R1 | 73****What's Changed | 74****Known Limitations | 74****Open Issues | 74****Resolved Issues | 75****Resolved Issues: 20.2R3 | 76****Resolved Issues: 20.2R2 | 76****Resolved Issues: 20.2R1 | 76****Documentation Updates | 76****Migration, Upgrade, and Downgrade Instructions | 77****Upgrade and Downgrade Support Policy for Junos OS Releases | 77****Junos OS Release Notes for Junos Fusion for Enterprise | 78****What's New | 79****What's Changed | 79****Known Limitations | 80**

Open Issues | **80**

Resolved Issues | **81**

Resolved Issues: Release 20.2R3 | **81**

Resolved Issues: Release 20.2R2 | **81**

Resolved Issues: Release 20.2R1 | **81**

Documentation Updates | **82**

Migration, Upgrade, and Downgrade Instructions | **82**

Basic Procedure for Upgrading Junos OS on an Aggregation Device | **83**

Upgrading an Aggregation Device with Redundant Routing Engines | **84**

Preparing the Switch for Satellite Device Conversion | **85**

Converting a Satellite Device to a Standalone Switch | **86**

Upgrade and Downgrade Support Policy for Junos OS Releases | **86**

Downgrading Junos OS | **87**

Junos OS Release Notes for Junos Fusion for Provider Edge | **88**

What's New | **88**

What's New in Release 20.2R3 | **89**

What's New in Release 20.2R2 | **89**

What's New in Release 20.2R1 | **89**

What's Changed | **90**

Known Limitations | **90**

Open Issues | **91**

Resolved Issues | **91**

Resolved Issues: 20.2R3 | **92**

Resolved Issues: 20.2R2 | **92**

Resolved Issues: 20.2R1 | **92**

Documentation Updates | **92**

Migration, Upgrade, and Downgrade Instructions | **93**

Basic Procedure for Upgrading an Aggregation Device | **93**

Upgrading an Aggregation Device with Redundant Routing Engines | **96**

Preparing the Switch for Satellite Device Conversion | **96**

Converting a Satellite Device to a Standalone Device | **98**

Upgrading an Aggregation Device | **100**

Upgrade and Downgrade Support Policy for Junos OS Releases | **101**

Downgrading from Junos OS Release 20.1 | **101**

**Junos OS Release Notes for MX Series | 102****What's New | 102****What's New in Release 20.2R3 | 103****What's New in Release 20.2R2-S3 | 103****What's New in Release 20.2R2-S2 | 103****What's New in Release 20.2R2 | 103****What's New in Release 20.2R1-S1 | 104****What's New in Release 20.2R1 | 104****What's Changed | 129****What's Changed in Release 20.2R3 | 129****What's Changed in Release 20.2R2 | 131****What's Changed in Release 20.2R1 | 133****Known Limitations | 136****General Routing | 137****Infrastructure | 138****Interfaces and Chassis | 138****MPLS | 138****Network Management and Monitoring | 138****Platform and Infrastructure | 138****Routing Protocols | 139****Open Issues | 139****Class of Service (CoS) | 140****EVPN | 140****Forwarding and Sampling | 140****General Routing | 141****High Availability (HA) and Resiliency | 145****Infrastructure | 145****Interfaces and Chassis | 145****Layer 2 Ethernet Services | 146****MPLS | 146****Platform and Infrastructure | 146****Routing Policy and Firewall Filters | 147****Routing Protocols | 147****Services Applications | 148**

Subscriber Access Management | **148**

User Interface and Configuration | **148**

VPNs | **149**

Resolved Issues | **149**

Resolved Issues: 20.2R3 | **150**

Resolved Issues: 20.2R2 | **160**

Resolved Issues: 20.2R1 | **169**

Documentation Updates | **185**

Advanced Subscriber Management Provider | **185**

Migration, Upgrade, and Downgrade Instructions | **186**

Basic Procedure for Upgrading to Release 20.2R3 | **187**

Procedure to Upgrade to FreeBSD 11.x-Based Junos OS | **187**

Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | **190**

Upgrade and Downgrade Support Policy for Junos OS Releases | **191**

Upgrading a Router with Redundant Routing Engines | **192**

Downgrading from Release 20.2R3 | **192**

Junos OS Release Notes for NFX Series | **193**

What's New | **193**

What's New in Release 20.2R3 | **194**

What's New in Release 20.2R2 | **194**

What's New in Release 20.2R1 | **194**

What's Changed | **195**

What's Changed in Release 20.2R3 | **196**

What's Changed in Release 20.2R2 | **196**

What's Changed in Release 20.2R1 | **196**

Known Limitations | **196**

Open Issues | **197**

High Availability | **198**

Interfaces | **198**

Platform and Infrastructure | **198**

Virtual Network Functions (VNFs) | **198**

Resolved Issues | **199**

Resolved Issues: 20.2R3 | **199**

Resolved Issues: 20.2R2 | **200**

Resolved Issues: 20.2R1 | 200

Documentation Updates | 201

Migration, Upgrade, and Downgrade Instructions | 202

Upgrade and Downgrade Support Policy for Junos OS Releases | 202

Basic Procedure for Upgrading to Release 20.2 | 203

Junos OS Release Notes for PTX Series | 204

What's New | 205

What's New in Release 20.2R3 | 205

What's New in Release 20.2R2 | 205

What's New in Release 20.2R1 | 205

What's Changed | 213

What's Changed in Release 20.2R3 | 214

What's Changed in Release 20.2R2 | 215

Known Limitations | 216

General Routing | 217

MPLS | 217

Routing Protocols | 217

Open Issues | 218

General Routing | 218

Infrastructure | 220

Interfaces and Chassis | 220

Layer 2 Ethernet Services | 220

MPLS | 220

Routing Protocols | 220

Resolved Issues | 221

Resolved Issues: 20.2R3 | 221

Resolved Issues: 20.2R2 | 223

Resolved Issues: 20.2R1 | 224

Documentation Updates | 226

Migration, Upgrade, and Downgrade Instructions | 226

Basic Procedure for Upgrading to Release 20.2 | 227

Upgrade and Downgrade Support Policy for Junos OS Releases | 229

Upgrading a Router with Redundant Routing Engines | 230



## Junos OS Release Notes for the QFX Series | 231

### What's New | 231

What's New in Release 20.2R3 | 232

What's New in Release 20.2R2 | 232

What's New in Release 20.2R1-S1 | 232

What's New in Release 20.2R1 | 234

### What's Changed | 256

What's Changed in Release 20.2R3 | 256

What's Changed in Release 20.2R2 | 257

What's Changed in Release 20.2R1 | 258

### Known Limitations | 260

Class of Service (CoS) | 260

Layer 2 Features | 260

Layer 2 Ethernet Services | 260

Platform and Infrastructure | 260

Routing Protocols | 261

### Open Issues | 262

EVPN | 262

High Availability (HA) and Resiliency | 263

Infrastructure | 263

Interfaces and Chassis | 263

Layer 2 Features | 263

Layer 2 Ethernet Services | 263

Platform and Infrastructure | 264

Routing Protocols | 266

Virtual Chassis | 267

### Resolved Issues | 267

Resolved Issues: 20.2R3 | 268

Resolved Issues: 20.2R2-S2 | 272

Resolved Issues: 20.2R2 | 272

Resolved Issues: 20.2R1 | 275

### Documentation Updates | 280

**Migration, Upgrade, and Downgrade Instructions | 281****Upgrading Software on QFX Series Switches | 281****Installing the Software on QFX10002-60C Switches | 284****Installing the Software on QFX10002 Switches | 284****Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches | 285****Installing the Software on QFX10008 and QFX10016 Switches | 287****Performing a Unified ISSU | 291****Preparing the Switch for Software Installation | 292****Upgrading the Software Using Unified ISSU | 292****Upgrade and Downgrade Support Policy for Junos OS Releases | 294****Junos OS Release Notes for SRX Series | 295****What's New | 296****What's New in Release 20.2R3 | 296****What's New in Release 20.2R2 | 296****What's New in Release 20.2R1 | 296****What's Changed | 306****What's Changed in Release 20.2R3 | 306****What's Changed in Release 20.2R2 | 308****What's Changed in Release 20.2R1-S1 | 310****What's Changed in Release 20.2R1 | 310****Known Limitations | 315****Flow-Based and Packet-Based Processing | 315****J-Web | 316****VPNs | 316****Open Issues | 317****Flow-Based and Packet-Based Processing | 317****J-Web | 318****Routing Policy and Firewall Filters | 318****VPNs | 318****Resolved Issues | 319****Resolved Issues: 20.2R3 | 319****Resolved Issues: 20.2R2 | 322****Resolved Issues: 20.2R1 | 324**

Documentation Updates | **328**

Migration, Upgrade, and Downgrade Instructions | **328**

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases | **328**

Junos OS Release Notes for vMX | **329**

What's New | **330**

What's New in Release 20.2R3 | **330**

What's New in Release 20.2R2 | **330**

What's Changed | **331**

What's Changed in Release 20.2R3 | **331**

What's Changed in Release 20.2R2 | **331**

Known Limitations | **331**

Open Issues | **331**

Resolved Issues | **331**

Resolved Issues: 20.2R3 | **332**

Resolved Issues: 20.2R2 | **332**

Licensing | **332**

Upgrade Instructions | **332**

Junos OS Release Notes for vRR | **333**

What's New | **333**

What's New in Release 20.2R3 | **333**

What's New in Release 20.2R2 | **334**

What's Changed | **334**

What's Changed in Release 20.2R3 | **334**

What's Changed in Release 20.2R2 | **334**

Known Limitations | **334**

Open Issues | **335**

Resolved Issues | **335**

Resolved Issues: 20.2R3 | **335**

Junos OS Release Notes for vSRX	<b>335</b>
What's New	<b>336</b>
What's New in Release 20.2R3	<b>336</b>
What's New in Release 20.2R2	<b>336</b>
What's Changed	<b>337</b>
What's Changed in Release 20.2R3	<b>337</b>
What's Changed in Release 20.2R2	<b>337</b>
Known Limitations	<b>338</b>
J-Web	<b>338</b>
Open Issues	<b>338</b>
Intrusion Detection and Prevention (IDP)	<b>339</b>
J-Web	<b>339</b>
Platform and Infrastructure	<b>339</b>
Resolved Issues	<b>339</b>
Resolved Issues: 20.2R3	<b>340</b>
Resolved Issues: 20.2R2	<b>340</b>
Migration, Upgrade, and Downgrade Instructions	<b>341</b>
Upgrading Software Packages	<b>342</b>
Validating the OVA Image	<b>348</b>
Upgrading Using ISSU	<b>348</b>
Licensing	<b>348</b>
Compliance Advisor	<b>349</b>
Finding More Information	<b>349</b>
Documentation Feedback	<b>349</b>
Requesting Technical Support	<b>351</b>
Self-Help Online Tools and Resources	<b>351</b>
Creating a Service Request with JTAC	<b>352</b>
Revision History	<b>352</b>

# Introduction

Junos OS runs on the following Juniper Networks<sup>®</sup> hardware: ACX Series, cSRX, EX Series, JRR Series, Junos fusion for enterprise, Junos Fusion for provider edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

These release notes accompany Junos OS Release 20.2R3 for the ACX Series, cSRX, EX Series, JRR Series, Junos fusion for enterprise, Junos fusion for provider edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and 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 348](#)
  - [Licensing on page 348](#)
  - [Compliance Advisor on page 349](#)
  - [Finding More Information on page 349](#)
  - [Documentation Feedback on page 349](#)
  - [Requesting Technical Support on page 351](#)

## Junos OS Release Notes for ACX Series

### IN THIS SECTION

- [What's New | 13](#)
- [What's Changed | 22](#)
- [Known Limitations | 27](#)
- [Open Issues | 31](#)
- [Resolved Issues | 35](#)

- Documentation Updates | 41
- Migration, Upgrade, and Downgrade Instructions | 41

These release notes accompany Junos OS Release 20.2R3 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

### IN THIS SECTION

- What's New in Release 20.2R3 | 13
- What's New in Release 20.2R2 | 13
- What's New in Release 20.2R1 | 14

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

### What's New in Release 20.2R3

There are no new features or enhancements to existing features for ACX Series routers in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features or enhancements to existing features for ACX Series routers in Junos OS Release 20.2R2.

## What's New in Release 20.2R1

### Hardware

- **New ACX710 Universal Metro Routers (ACX Series)**—In Junos OS Release 20.2R1, we introduce the ACX710 router. The ACX710 is a compact 1-U router that provides system throughput of up to 320 Gbps through the following port configurations:
  - Twenty-four 10GbE or 1GbE ports (ports 0 through 23) that operate at 10-Gbps speed when you use small form-factor pluggable plus (SFP+) transceivers or at 1-Gbps speed when you use small form-factor pluggable (SFP) optics. Ports 0 through 15 also support 1000 Mbps speeds when you use tri-rate SFP optics. Ports 16 through 23 support 100 Mbps and 1000 Mbps speeds when you use tri-rate SFP optics.
  - Four 100GbE ports (ports 0 through 3) that support quad small form-factor pluggable 28 (QSFP28) transceivers. You can channelize these ports into four 25-Gbps interfaces using breakout cables and channelization configuration. These ports also support 40-Gbps speed when you use quad small form-factor pluggable plus (QSFP+) optics. You can channelize these 40-Gbps ports into four 10-Gbps interfaces using breakout cables and channelization configuration. [See [Channelize Interfaces on ACX710 Routers.](#)]

The ACX710 router is a DC-powered device that is cooled using a fan tray with five high-performance fans to cool the chassis.

To install the ACX710 router hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the [ACX710 Universal Metro Router Hardware Guide](#).

[Table 1 on page 14](#) summarizes the ACX710 features supported in Junos OS Release 20.2R1.

**Table 1: Features Supported by the ACX710 Routers**

Feature	Description
Class of service (CoS)	<ul style="list-style-type: none"> <li>• Standard CoS feature support, including configuring classification, rewrite, shaping, buffering, and scheduling parameters for traffic management. [See <a href="#">CoS on ACX Series Routers Features Overview.</a>]</li> </ul>
DHCP	<ul style="list-style-type: none"> <li>• DHCP server, DHCP client, and DHCP relay configuration for IPv4 and IPv6 services. [See <a href="#">Understanding DHCP Client Operation on ACX Series.</a>]</li> </ul>
EVPN	<ul style="list-style-type: none"> <li>• EVPN-VPWS. [See <a href="#">Overview of VPWS with EVPN Signaling Mechanisms EVPN-VPWS with flexible cross-connect (FXC).</a>]</li> <li>• EVPN-VPWS with flexible cross-connect (FXC). [See <a href="#">Overview of Flexible Cross-Connect Support on VPWS with EVPN.</a>]</li> <li>• EVPN with ELAN services over MPLS. [See <a href="#">EVPN Overview.</a>]</li> </ul>

Table 1: Features Supported by the ACX710 Routers (*continued*)

Feature	Description
Firewalls and policers	<ul style="list-style-type: none"> <li>Configure firewall filters on packets (families such as bridge domain, IPv4, IPv6, CCC, and MPLS) based on packet match conditions. Along with the match conditions, actions such as count, discard, log, syslog, policer are performed on the packets that match the filter. You can configure policers and attach them to a firewall term. [See <a href="#">Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview</a>.]</li> </ul>
High availability (HA) and resiliency	<ul style="list-style-type: none"> <li>VRRP protocol support with Broadcom's DNX chipset. [See <a href="#">Understanding VRRP Overview</a>.]</li> <li>Configure alarm input and output, manage FRUs, and monitor environment. The router also supports field-replaceable unit (FRU) management and environmental monitoring. [See <a href="#">alarm-port</a>.]</li> <li>Platform resiliency to handle failures and faults of the components such as fan trays, temperature sensors, and power supplies. The router also supports firmware upgrade for FPGA and U-boot. [See <a href="#">show chassis alarms</a> and <a href="#">show system firmware</a>.]</li> </ul>
Layer 2 features	<ul style="list-style-type: none"> <li>Layer 2 support: bridging, bridge domain with no vlan-id, with vlan-id none, or with single vlan-id, single learning domain support, Q-in-Q service for bridging, MAC limit feature support, no local switching support for bridge domain, and E-LINE from a bridge with no MAC learning. [See <a href="#">Layer 2 Bridge Domains on ACX Series Overview</a>.]</li> <li>Layer 2 support for bridge interfaces for vlan-map push operation, swap operation, pop operation, and swap-swap operation. [See <a href="#">Layer 2 Bridging Interfaces Overview</a>.]</li> <li>Layer 2 support for control protocols (L2CP): RSTP, MSTP, LLDP, BPDU guard/protection, loop protection, root protection, Layer 2 protocol tunneling, storm control, IRB interface, LAG support with corresponding hashing algorithm, E-LINE, E-LAN, E-ACCESS, and E-Transit service over L2/Bridge with the following AC interface types: Port, VLAN, Q-in-Q, VLAN range and VLAN list. [See <a href="#">Layer 2 Control Protocols on ACX Series Routers</a>.]</li> <li>Layer 2 circuit cross-connect (L2CCC) support for Layer 2 switching cross-connects. You can leverage the hardware support available for cross-connects on the ACX710 device with the Layer 2 local switching functionality using certain models. With this support, you can provide the EVP and EVPL services. [See <a href="#">Configuring MPLS for Switching Cross-Connects</a>.]</li> <li>Reflector function support in RFC 2544. [See <a href="#">RFC 2544-Based Benchmarking Tests Overview</a>.]</li> </ul>



Table 1: Features Supported by the ACX710 Routers (*continued*)

Feature	Description
Layer 3 features	<ul style="list-style-type: none"> <li>• Layer 3 VPN and Layer 3 IPv6 VPN Provider Edge router (6VPE) support over MPLS. The router uses MPLS as a transport mechanism with support for label-switching router (LSR), label edge routers (LERs), and pseudowire services. These protocols are also supported: ECMP, OSPF, IS-IS, and BGP. [See <a href="#">Understanding Layer 3 VPNs.</a>]</li> <li>• Basic Layer 3 services over segment routing infrastructure. The segment routing features supported are: segment routing with OSPF through MPLS, segment routing with IS-IS through MPLS, segment routing traffic engineering (SR-TE), segment routing global block (SRGB) range label used by source packet routing in networking (SPRING), anycast segment identifiers (SIDs) and prefix SIDs in SPRING, and segment routing with topology independent (TI)-loop-free alternate (LFA) provides fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. [See <a href="#">Segment Routing LSP Configuration.</a>]</li> <li>• Enhanced timing and synchronization support using Synchronous Ethernet with ESMC and BITS-Out. [See <a href="#">Synchronous Ethernet Overview</a> and <a href="#">synchronization (ACX Series).</a>]</li> <li>• Supports full-mesh VPLS domain deployment. The router supports interworking of both BGP as well as LDP-based VPLS. BGP can be used only for auto-discovery of the VPLS PEs, while LDP signaling for VPLS connectivity. [See <a href="#">Introduction to VPLS.</a>]</li> </ul>
MPLS	<ul style="list-style-type: none"> <li>• Supports the Path Computation Element Protocol (PCEP). You can configure the PCEP implementation for both RSVP-TE and segment routing label-switched paths (LSPs). [See <a href="#">PCEP Configuration.</a>]</li> <li>• Support for MPLS fast reroute (FRR) and unicast reverse-path forwarding (uRPF). [See <a href="#">fast-reroute (Protocols MPLS)</a> and <a href="#">Guidelines for Configuring Unicast RPF on ACX Series Routers.</a>]</li> <li>• Provides MPLS ping and traceroute support. [See <a href="#">MPLS Connectivity Verification and Troubleshooting Methods.</a>]</li> </ul>
Multicast	<ul style="list-style-type: none"> <li>• Multicast support for IPv4 and IPv6 PIM-SM, SSM, IGMP snooping and proxy support, IGMP, IGMPv1/v2/v3 snooping, IGMP snooping support for LAG, global multicast support, MLD, and multicast support on IRB. [See <a href="#">Multicast Overview.</a>]</li> </ul>

Table 1: Features Supported by the ACX710 Routers (*continued*)

Feature	Description
Network management and monitoring	<ul style="list-style-type: none"> <li>• TWAMP support. [See <a href="#">Two-Way Active Measurement Protocol on ACX Series.</a>]</li> <li>• NETCONF sessions over TLS. [See <a href="#">NETCONF Sessions over Transport Layer Security (TLS).</a>]</li> <li>• Support for adding custom YANG data models to the Junos OS schema [See <a href="#">Understanding the Management of Non-Native YANG Modules on Devices Running Junos OS.</a>]</li> <li>• Secure boot support in U-boot phase to authenticate and verify the loaded software image while also preventing software-based attack. [See <a href="#">Software Installation and Upgrade Guide.</a>]</li> </ul>
OAM	<ul style="list-style-type: none"> <li>• IEEE 802.3ah standard for operation, administration, and management (OAM) connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM. [See <a href="#">IEEE 802.1ag OAM Connectivity Fault Management Overview.</a>]</li> </ul>
System management	<ul style="list-style-type: none"> <li>• Zero-touch provisioning (ZTP) can automate the provisioning of the device configuration and software image. [See <a href="#">Software Installation and Upgrade Guide.</a>]</li> </ul>

Table 1: Features Supported by the ACX710 Routers (*continued*)

Feature	Description
---------	-------------

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

### **Authentication, Authorization, and Accounting**

- **Support for LDAP authentication and authorization over TLS (ACX710)**— Starting in Junos OS Release 20.2R1, we support LDAP authentication and authorization for Junos OS user login. Through the use of LDAP over TLS (LDAPS), we've implemented the LDAP authentication and authorization support for Junos OS user login user by providing TLS security between the device running Junos OS (which is the LDAPS client) and the LDAPS server.

To enable LDAPS support, you can configure the **ldaps-server** option at the **[edit system authentication-order]** hierarchy level. LDAPS ensures the secure transmission of data between a client and a server with better privacy, confidentiality, data integrity and higher scalability.

[See [Understanding LDAP Authentication over TLS](#).]

### **Class of Service (CoS)**

- **Support for hierarchical class of service (HCoS) (ACX5448)**—Starting with Junos OS Release 20.2R1, ACX5448 devices support up to four levels of hierarchical scheduling (physical interfaces, logical interface sets, logical interfaces, and queues). By default, all interfaces on the ACX5448 use port-based scheduling (eight queues per physical port). To enable hierarchical scheduling, set **hierarchical-scheduler** at the **[edit interfaces interface-name]** hierarchy level.

[See [Hierarchical Class of Service in ACX Series Routers](#).]

### **EVPN**

- **Noncolored SR-TE LSPs with EVPN-MPLS (ACX5448, EX9200, MX Series, and vMX)**—Starting in Junos OS Release 20.2R1, ACX5448, EX9200, MX Series, and vMX routers support noncolored static segment routing-traffic engineered (SR-TE) label-switched paths (LSPs) with an EVPN-MPLS core network and the following Layer 2 services running at the edges of the network:

- E-LAN
- EVPN-ETREE
- EVPN-VPWS with E-Line

Without color, all LSPs resolve using a BGP next hop only.

The Juniper Networks routers support noncolored SR-TE LSPs in an EVPN-MPLS core network with the following configurations:

- EVPN running in a virtual switch routing instance

- Multihoming in active/active and active/standby modes

The Juniper Networks routers also support noncolored SR-TE LSPs when functioning as a Data Center Interconnect (DCI) device that handles EVPN Type 5 routes.

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

### **Interfaces and Chassis**

- **Port speeds and channelization (ACX710 routers)**—Starting in Junos OS Release 20.2R1, you can configure multiple speeds and interface channelization on our new ACX710 router. The router has 28 ports, which support the following speeds:
  - Ports **0** through **23** on PIC 0 support 1-Gbps speed (with SFP transceivers) and 10-Gbps speed (with SFP+ transceivers).
  - Ports **0** through **3** on PIC 1 support the default 100-Gbps speed (with QSFP28 transceivers) or the configured 40-Gbps speed (with QSFP+ transceivers). You can use the **set chassis fpc slot-number pic pic-number port port-number speed speed** CLI command and breakout cables to channelize each:
    - 100-Gbps port into four 25-Gbps interfaces
    - 40-Gbps port into four 10-Gbps interfaces

[See [Channelize Interfaces on ACX710 Routers.](#)]

- **Ethernet OAM and BFD support (ACX710)**—Starting in Junos OS Release 20.2R1, the ACX710 routers support IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM) connectivity fault management (CFM), BFD, and the ITU-T Y.1731 standard for Ethernet service OAM.

[See [Introduction to OAM Connectivity Fault Management \(CFM\).](#)]

- **Alarm port configuration, FRU management, and environmental monitoring (ACX710)**—Starting in Junos OS Release 20.2R1, you can configure the alarm port on the ACX710 router. You can use the alarm input to connect the router to external alarm sources such as security sensors so that the router receives alarms from these sources and displays those alarms. You can use the alarm output to connect the router to an external alarm device that gives audible or visual alarm signals based on the configuration. You can configure three alarm inputs and one alarm output by using the **alarm-port** statement at the **[edit chassis]** hierarchy level. You can view the alarm port details by using the **show chassis craft-interface** command.

The ACX710 also supports FRU management and environmental monitoring.

[See [alarm-port.](#)]

- **Multichassis link aggregation groups, configuration synchronization, and configuration consistency check (ACX5448 routers)**—Starting in Junos OS Release 20.2R1, multichassis link aggregation (MC-LAG) includes support of Layer 2 circuit functionality with **ether-ccc** and **vlan-ccc** encapsulations.

MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running spanning-tree protocols (STPs).

[See [Multichassis Link Aggregation Features, Terms, and Best Practices.](#)]

### *Juniper Extension Toolkit (JET)*

- **JET Clang toolchain supports cross-compiling JET applications for use on ARM platforms (ACX710)**—Starting in Junos OS Release 20.2R1, you can use the Clang toolchain to compile JET applications written in C, Python, or Ruby to run on the ARM architecture as well as Junos OS with FreeBSD and upgraded FreeBSD. The Clang toolchain for ARM is included in the JET software bundle. After you have downloaded the JET software bundle, you can access the Clang toolchain at `/usr/local/junos-jet/toolchain/llvm/`. Use the `mk-arm,bsd` command to use the Clang toolchain to compile your application.

[See [Develop On-Device JET Applications.](#)]

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the `set system scripts language python3` command.

[See [language \(Scripts\)](#), [Develop Off-Device JET Applications](#), and [Develop On-Device JET Applications.](#)]

### *Junos Telemetry Interface*

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models `openconfig-local-routing.yang` and `openconfig-network-instance.yang`.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [Mapping OpenConfig Network Instance Commands to Junos Operation.](#)]

### *MPLS*

- **Support for MPLS ping and traceroute for segment routing (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.2R1, we extend the MPLS ping and traceroute support for all types segment routing--traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

We also support the following features:

- FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
- ECMP traceroute support for all types of SR-TE paths.

We do not support the following:

- Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
- OAM for IPv6 prefix.

- BFD

[See [traceroute mpls segment-routing spring-te](#) and [ping mpls segment routing spring-te](#).]

### **Multicast**

- **Support for IPv6 multicast using MLD (ACX5448)**—Starting with Junos OS Release 20.2R1, ACX5448 routers support Multicast Listener Discovery (MLD) snooping with MLDv1 and MLDv2 for both any source multicast and SSM. Support for MLD snooping in EVPN was introduced in Junos OS Release 19.4R2.

MLD snooping for IPv6 is used to optimize Layer 2 multicast forwarding. It works by checking the MLD messages sent between hosts and multicast routers to identify which hosts are interested in receiving IPv6 multicast traffic, and then forwarding the multicast streams to only those VLAN interfaces that are connected to the interested hosts (rather than flooding the traffic to all interfaces). You can enable or disable MLD snooping per VLAN at the `[edit protocols mld-snooping vlan vlan-ID]` hierarchy level. Note, however, that you cannot use ACX Series routers to connect to a multicast source.

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

### **Network Management and Monitoring**

- **NETCONF sessions over TLS (ACX710)**—Starting in Junos OS Release 20.2R1, ACX710 routers support establishing Network Configuration Protocol (NETCONF) sessions over Transport Layer Security (TLS) to manage devices running Junos OS. TLS uses mutual X.509 certificate-based authentication and provides encryption and data integrity to establish a secure and reliable connection. NETCONF sessions over TLS enable you to remotely manage devices using certificate-based authentication and to more easily manage networks on a larger scale than when using NETCONF over SSH.

[See [NETCONF Sessions over Transport Layer Security \(TLS\)](#).]

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **Support for port mirroring (ACX5448)**—Starting in Junos OS Release 20.2R1, you can use analyzers to mirror copies of packets to a configured destination. Mirroring helps in debugging network problems and also in defending the network against attacks. You can mirror all ingress traffic to a configured port (or port list), using a protocol analyzer application that passes the input to mirror through a list of ports configured through the logical interface. You configure the analyzer at the `[edit forwarding-options analyzer]` hierarchy level.

Configuration guidelines and limitations:

- Maximum of four default analyzer sessions

- LAGs supported as mirror output; a maximum of eight child members
- Not supported:
  - Egress mirroring
  - Mirroring on IRB, Virtual Chassis, or management interfaces
  - Nondefault analyzers

[See [show forwarding-options analyzer.](#)]

### **Routing Policy and Firewall Filters**

- **Support for firewall filters and policers (ACX710)**—Starting with Junos OS Release 20.2R1, the ACX710 router supports configuring firewall filters on packets (families such as bridge domain, IPv4, IPv6, CCC, and MPLS) based on packet match conditions. Along with the match conditions, actions such as count, discard, log, syslog, and policer are performed on the packets that match the filter. You can configure policers and attach them to a firewall term.

[See [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview.](#)]

### SEE ALSO

[What's Changed | 22](#)

[Known Limitations | 27](#)

[Open Issues | 31](#)

[Resolved Issues | 35](#)

[Documentation Updates | 41](#)

[Migration, Upgrade, and Downgrade Instructions | 41](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 23](#)
- [What's Changed in Release 20.2R2 | 24](#)
- [What's Changed in Release 20.2R1-S2 | 25](#)
- [What's Changed in Release 20.2R1 | 25](#)

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

## What's Changed in Release 20.2R3

### *Junos OS XML API and Scripting*

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

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

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

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

### *Network Management and Monitoring*

- Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:
  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.
  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and emits only an `<ok>` or `<rpc-error>` element.

[See [Configuring RFC-Compliant NETCONF Sessions](#).]

- Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—When you configure the `rfc-compliant` statement at the `edit system services netconf` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:



- If a successful **<commit>** operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
- The NETCONF server response emits the **<source-daemon>** element as a child of the **<error-info>** element instead of the **<rpc-error>** element.
- If you also configure the **flatten-commit-results** statement at the **edit system services netconf** hierarchy level, the NETCONF server suppresses any **<commit-results>** XML subtree in the response and only emits an **<ok>** or **<rpc-error>** element.

[See [Configuring RFC-Compliant NETCONF Sessions.](#)]

### *User Interface and Configuration*

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

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

## What's Changed in Release 20.2R2

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

See [gigheter-options](#) and [ether-options](#)

### *Routing Protocols*

- **Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the **Isdist.0** and **Isdist.1** routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the **Isdist.0** and **Isdist.1** routing tables as part of node characteristics and advertised them as the router ID.
- **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.

## What's Changed in Release 20.2R1-S2

### General Routing

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

See [gigether-options](#) and [ether-options](#)

## What's Changed in Release 20.2R1

### 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

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the `persist-groups-inheritance` option at the `[edit system commit]` hierarchy level is enabled by default. To disable this option, use `no-persist-groups-inheritance`.

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

- **New major alarms (ACX-710)** —We have introduced the following major alarms:
  - PTP No Foreign Master—Indicates that the external Precision Time Protocol (PTP) master is not sending announce packets.
  - PTP Sync Fail—Indicates that the PTP lock-status is not in Phase Aligned state.
  - Chassis Loss of all Equipment Clock Synch References—Indicates that both the primary and secondary SyncE references have failed and the chassis PLL is in holdover.
  - Chassis Loss of Equipment Clock Synch Reference 1—Indicates that the primary SyncE reference has failed, and no secondary SyncE reference is configured or present.
  - Chassis Loss of Equipment Clock Synch Reference 2—Indicates that you have configured at least two or more SyncE sources and the secondary SyncE source has failed.

NOTE: These alarms get cleared when the system recovers from the error condition.

See [show chassis alarms](#).

### *Juniper Extension Toolkit (JET)*

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

[See [Develop Off-Device JET Applications](#) and [Develop On-Device JET Applications](#).]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH\_USAGE\_INVALID**.

[See [Juniper EngNet](#).]

### *Network Management and Monitoring*

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

### SEE ALSO

---

[What's New | 13](#)

---

[Known Limitations | 27](#)

---

[Open Issues | 31](#)

---

[Resolved Issues | 35](#)

---

[Documentation Updates | 41](#)

---

[Migration, Upgrade, and Downgrade Instructions | 41](#)

## Known Limitations

### IN THIS SECTION

- [General Routing | 27](#)

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

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

### General Routing

- If Layer 2 VPN sessions have the OAM **control-channel** option set to **router-alert-label**, the **no-control-word** option in the Layer 2 VPN should not be used for BFD sessions to come up. [PR1432854](#)
- In case of Dot1P, CFI rewrite based on TC or DP classification is not possible on the ACX5448 and ACX710 routers. As a workaround to preserve or control the incoming packet CFI bit at egress side (rewrite), configure 802.1ad, which has the control over the CFI rewrite as well. [PR1435966](#)
- The time consumed on 1-Gigabit performance is not equal to that on 10-Gigabit performance. Compensation is done to bring the mean value under class A but the peak-to-peak variations are high and can go beyond 100 ns. It has a latency variation with peak-to-peak variations of around 125–250 ns without any traffic (for example, 5–10 percent of the mean latency introduced by each phy which is of around 2.5us). [PR1437175](#)
- With an asymmetric network connection, EX: 10G MACsec port connected to a 10-Gigabit Ethernet channelized port, high and asymmetric T1 and T4 time errors introduce a high two-way time error. This introduces different CF updates in 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. Find the maximum and mean values of the time errors with different traffic rates (for example, two router scenario). The maximum value can jump as high as 1054 ns with 95 percent traffic, 640 ns with 90 percent traffic, and 137 ns with no traffic. [PR1441388](#)
- On the ACX710 router, a variable amount of time is taken to reflect the TWAMP packets. Because of this, the packet latency is not uniform. [PR1477329](#)
- On the ACX710 router, as per current design and BCOM input, load balancing does not work on any packet which is injected from host path. [PR1477797](#)
- On the ACX710 router, OSPF neighbors are not learned via VPLS connections because the **vlan-tags outer *vlan-id1* inner *vlan-id2*** statement is not supported in VPLS routing instance. [PR1477957](#)

- On the ACX710 router, sequential increment of both SRC and DST MAC do not provide better load balance as per HASH result. [PR1477964](#)
- On the ACX710 router, load balancing does not happen based on inner IP address when MPLS labelled traffic is received on NNI interface. [PR1478945](#)
- On the ACX710 router, for TCP protocol as well as for non-TCP protocol, loss-priority medium-low is not supported. [PR1479164](#)
- For ethernet-vpls encapsulation, if both DST IP and SRC IP are identically varied at the same octet, then hashing might not happen and leads to undefined behavior in load balancing on the ACX710 router. [PR1479767](#)
- For bridge LB with vlan-bridge encapsulation, if both SRC IP and DST IP are incremented or decremented by the same order (such as DIP = 10.1.1.1 (increment by 1 up to 100) and SIP = 20.2.3.1 (increment by 1 upto 100)), then hashing does not happen on the ACX710 router. [PR1479986](#)
- For vlan-ccc encapsulation, if both SRC IP and DST IP are incremented or decremented by the same order (such as DIP = 10.1.1.1 (increment by 1 upto 100) and SIP = 20.2.3.1 (increment by 1 upto 100)), then hashing does not happen on the ACX710 router. [PR1480228](#)
- On the ACX710 router, the input packet statistics for the **show interfaces** command represents the input packets at the MAC. The error packets which get dropped by MAC and that do not reach PHY will not be accounted. [PR1480413](#)
- The accounting-profile statement is not supported on any of the ACX platforms. Therefore, the cli configuration for accounting-profile is hidden. [PR1480546](#)
- On ACX710 routers, temperature threshold of fire shutdown and high fan speed are same. [PR1481248](#)
- MRU field is not shown in the **show interface output** command, The behavior is same across all the ACX platforms. Configuration commit does not show any error as no platform checks exist at that CLI level. [PR1481585](#)
- Fragmentation or reassembly is not supported on ACX710 platforms due to the lack of hardware support. [PR1481867](#)
- On ACX5448 and ACX710 routers, each traffic stream is measured independently per port. Storm control is initiated only if one of the streams exceeds the storm control level. For example, if you set a storm control level of 100 Megabits and the broadcast and unknown unicast streams on the port are each flowing at 80 Mbps, storm control is not triggered. [PR1482005](#)
- System lands in loader prompt when power cycle is done with faulty USB plugged in. [PR1482658](#)
- VLAN MAP operations for VPLS/L2circuit/EVPN will support only with TPID 8100. [PR1483023](#)
- On the ACX710 router, RFC2544 reports high latency and throughput loss when the packet size is 64 bytes at 100 percent line rate on the ASIC. The ASIC has low threshold value due to which packets are moved to DRAM from SRAM. When packets are moved to DRAM, high latency and packet drop are observed. [PR1483370](#)

- On the Packet Forwarding Engine shell, diagnostics are displayed for 100 G DAC cable under **show diagnostics info** command. This is because the DAC cable has its diags page populated which is all zeroes. The diagnostics under CLI are displayed correctly as N/A. [PR1483416](#)
- ACX710 supports the maximum term/match up to 4000 ingress and 3000 egress entries. Scaling is unidimensional between ingress and egress as TCAM banks are shared. [PR1483560](#)
- On the ACX710 router, VRRP over aggregated Ethernet interface is not supported. [PR1483594](#)
- On the ACX710 router, traffic loss is seen for segment routing, if protection (FRR) is enabled for 128 IPv6 prefix route. [PR1484234](#)
- Counters for PCS bit errors are not supported because of hardware limitations. Hence "Bit errors" and "Errored blocks" are not supported on an ACX710. [PR1484766](#)
- If any queue is configured with high priority, it is expected that accuracy of traffic distribution might vary for normal queues because of chip limitation. [PR1485405](#)
- Tagged LACP packets are not terminated by the device but flooded in the bridge domain. This is because tagged LACP packets are considered data packets as LACP is supposed to be untagged. [PR1486274](#)
- For Layer 3 VPN configuration, sequential increment of both SRC IP and DST IP address would not provide better load balance as per hash result on the ACX710 router. [PR1486406](#)
- On the ACX710 router, double tagged interfaces implicit normalization to VLAN ID none is not supported. [PR1486515](#)
- On the ACX710 router, double tagged interfaces implicit normalization to VLAN ID none, ingress VLAN map operation, and pop-pop are not supported. [PR1486520](#)
- On the ACX710 router, packet priority at egress is derived from the internal priority. This internal priority is derived from the outer VLAN priority at ingress. Thus, the exiting packet retains the same priority as the ingress outer VLAN priority. [PR1486571](#)
- When you add or delete a configuration or a LAG member link flaps, configuration updates happen for all other members of the LAG too. This results in transient traffic drop on the ACX710 devices. [PR1486997](#)
- On the ACX710 router, double tagged ELMI and LLDP PDUs are dropped when L2PT is enabled for these protocols on the ingress interface. These PDUs are supposed to be untagged/native VLAN tagged and hence the drop. [PR1487931](#)
- On the ACX710 router, VLAN map operations like swap/swap does not work because the **vlan-tags outer *vlan-id1* inner *vlan-id2*** statement is not supported in VPLS routing instance. [PR1488084](#)
- On the ACX710 router, whenever the 100-Gigabit Ethernet interface is disabled, the alarm is not shown in the **jnxDomMib jnxDomCurrentLaneWarnings** and **jnxDomCurrentLaneAlarms**. [PR1489940](#)
- On the ACX710 router, in case of Layer 2 circuit, load balancing does not occur based on inner MAC address when MPLS labelled traffic is received on an NNI interface. [PR1490441](#)
- EVPN-VPWS, L3VPN, and L2VPN FRR convergence time with aggregated Ethernet as the Active core interface is not meeting <50 ms and might be 100 ms to 150 ms. [PR1492730](#)

- On the ACX710 router, unable to scale 1000 CFM sessions at 3 ms intervals; an error message is observed. [PR1495753](#)
- On ACX5448 routers, aggregated Ethernet LACP toggles with host path traffic with MAC rewrite configuration enabled. [PR1495768](#)
- The **traceroute mpls ldp** command does not work in case **explicit-null** is configured. It does not affect data path traffic. [PR1498339](#)
- On the ACX710 router, the convergence time for the traffic to switch over from the primary to the secondary link during link flap could be expected to be around 60 to 200 ms with the basic link aggregation configuration. [PR1499965](#)
- The MAC learning rate in ACX710 is measured as 2621 entries per second in software when there is no intervention of polling the MAC table entry from CLI periodically. When there is periodic polling in retrieving the MAC table entries through **show command output** CLI command or through script, during MAC learning in progress, the number of MAC entries learnt is around 1730 per second. Because this will take the CPU time and have an impact in the number of MAC learnt entries in software table. [PR1500523](#)
- On ACX710 routers, the PTP clock recovery is re-started when the clksyncd process is restarted. This results in the PTP lock state moving to freerun on the clksyncd process restart. [PR1502162](#)
- On the ACX710 router, not able to scale BFD to 1024 sessions with IPv4 and IPv6. [PR1502170](#)
- On the ACX710 router, GPS satellites do not track intermittently with GPS-only constellation. [PR1505325](#)
- On ACX710 routers, PTP does not work with vlan-map operations. [PR1507809](#)
- On ACX710 routers, unexpected delay counter values are seen in the output for show ptp statistics detail when upstream master stops sending the PTP packets. [PR1508031](#)
- On ACX710 routers, if the ukern is restarted with the chassis-control restart command, the state of the PTP lock status on the Routing Engine will transition among holdover/acquiring/phase locked. The clock data is displayed accordingly. Once the Packet Forwarding Engine is up and running after restart, clock data is stable and correct. During the time the Packet Forwarding Engine is not up, the clock display is inconsistent but eventually it becomes valid once the Packet Forwarding Engine is up and the clock is created and announce packets are being generated. [PR1508385](#)
- On ACX710 routers, servo status toggles to free-run/holdover-in-spec/acquiring on doing ABMCA change from virtual port to PTP. [PR1510880](#)
- Whenever we switch from one server to other server, HOLD-OVER-IN state expected for sometime with current implementation until it switches to other server(using warm reset API provided). This state cannot be avoided and it does not impact any functionality. HOLD-OVER-IN state some intermediate state expected from servo, since this is state comes from hardware while switch to other reference. [PR1513659](#)
- On ACX710 routers, local repair can be in seconds (>50 ms) during FRR convergence. If explicit NULL is configured on the PHP node and on the PHP node of the backup path, the link failure is observed at PHP node. Global repair resumes the traffic flow. [PR1515512](#)

- The maximum FIB route scale supported in an ACX710 router are as below:  
FIB IPv6 route scale - 80,000  
FIB IPv4 route scale - 170,000  
If routes are added above this scale, an error indicating **lpm route add** failure is reported. [PR1515545](#)
- 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](#)
- On the ACX5048 router, **queue-counters-trans-bytes-rate** are more than expected while configuring the physical interface and logical interface shaping with the transmit rate and scheduler-map. [PR1538934](#)

#### SEE ALSO

[What's New | 13](#)

[What's Changed | 22](#)

[Open Issues | 31](#)

[Resolved Issues | 35](#)

[Documentation Updates | 41](#)

[Migration, Upgrade, and Downgrade Instructions | 41](#)

## Open Issues

#### IN THIS SECTION

● [General Routing | 32](#)

● [Virtual Chassis | 34](#)

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

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



## General Routing

- Loopback status is not shown for OT interfaces on CLI (available from vty only). [PR1358017](#)
- The SD (Signal Degrade) threshold is normally lower than the SF threshold (that is, so that as errors increase, SD condition is encountered first). For the ACX6360 optical links there is no guard code to prevent the user from setting the SD threshold above the SF threshold, which would cause increasing errors to trigger the SF alarm before the SD alarm. This will not cause any issues on systems with correctly provisioned SD/SF thresholds. [PR1376869](#)
- On the ACX6360-OR router, enhancement is needed for the FRR BER threshold SNMP support. [PR1383303](#)
- On ACX6360 router, Tx power cannot be configured using + sign. [PR1383980](#)
- The switchover time observed was more than 50 minutes under certain soak test conditions with an increased scale with a multiprotocol multirouter topology. [PR1387858](#)
- The ccc logs are not compressed after rotation. [PR1398511](#)
- A jnxIfOtnOperState trap notification is sent for all ot-interfaces. [PR1406758](#)
- The em2 interface configuration causes FPC to crash during initialization and FPC does not come online. After deleting the em2 configuration and restarting the router, FPC comes online. [PR1429212](#)
- DHCP clients are not able to scale to 96,000. [PR1432849](#)
- Protocols get forwarded when using a non-existing SSM map source address in IGMPv3 instead of pruning. [PR1435648](#)
- Memory leaks are expected in this release. [PR1438358](#)
- Drop profile maximum threshold might not be reached when the packet size is other than 1000 bytes. This is due to the current design limitation. [PR1448418](#)
- The IPv6 BFD sessions flap when configured below 100 ms. [PR1456237](#)
- The CFM remote MEP does not come up after configuration or remains in start state. [PR1460555](#)
- On ACX710 routers, packet drop is observed after changing ALT port cost for RSTP. [PR1482566](#)
- On ACX710 routers, VRRP over dual tagged interface is not supported. [PR1483759](#)
- On ACX710 routers, FEC of channel 0 in a channelized 25-Gigabit Ethernet interface is set to **None** while channels 1, 2, and 3 have FEC74 as the default value for 100-Gigabit Ethernet LR4 optics. The desired FEC value can be set through the CLI command **set interfaces et-x/y/z: channel no gigheter-options fec fec value**. [PR1488040](#)
- Commit check error might be found when members of different speed added to aggregated bundle when mixed mode is not set. [PR1490373](#)
- The following syslog error message is observed: **ACX\_DFW\_CFG\_FAILED**. [PR1490940](#)

- On ACX6360 platforms, port mirroring does not work when the port mirroring is configured with the firewall filter. [PR1491789](#)
- On ACX710 routers, the **ping mpls l2ckt/l2vpn** command does not work if the **no-control-word** statement is configured. [PR1492963](#)
- On ACX710 routers, the **ping mpls l2circuit** command does not work if the **explicit-null** is configured. It does not affect the data path traffic. [PR1494152](#)
- On ACX710 routers, the PTP clock recovery is re-started when the **clksyncd** process is restarted. This results in the PTP lock state moving to freerun on the **clksyncd** process restart. [PR1502162](#)
- On ACX710 routers, if we configure DHCP option 012 host-name in DHCP server and the actual base configuration file also has the host-name in it, then overwriting of the base configuration file's host-name with the DHCP option 012 host-name is happening. [PR1503958](#)
- On the ACX6360 platform, the core file **core-ripsaw-node-aftd-expr** is generated and you are unable to back trace the file. [PR1504717](#)
- On ACX710 routers, when the following steps are done for PTP, chassis does not lock:
  1. Use one or two ports as source for chassis synchronization and lock both PTP and SyncE locked.
  2. Disable both logical interfaces.
  3. Restart **clksyncd**.
  4. Rollback 1.

As a workaround, you can avoid this issue by deleting the PTP configuration, restarting **clksyncd**, and then reconfiguring PTP. [PR1505405](#)
- MPLS LSP check fails while verifying basic LSP retry limit. Reset the src-address of the LSP to 0 (if src-address is not configured) whenever it changes its state from up to down. So when the ingress LSP goes to down state, reset it to 0. The script fails because the script checks for src-address to be present for the ingress LSP session. [PR1505474](#)
- On ACX710 routers, PTP does not seem to work with vlan-map operations. [PR1507809](#)
- On ACX710 routers, unexpected delay counter values are seen under **show ptp statistics detail** when upstream master stops sending the PTP packets. [PR1508031](#)
- On ACX710 routers, if the ukern is restarted with the **chassis-control restart** command, the state of the PTP lock status on the Routing Engine changes among holdover/acquiring/phase locked. The clock data is displayed accordingly. Once the Packet Forwarding Engine is up and runs after restart, clock data is stable and correct. During the time the Packet Forwarding Engine is not up, the clock display is inconsistent but eventually it becomes valid once the Packet Forwarding Engine is up and the clock is created and announce packets are being generated. [PR1508385](#)
- On ACX710 routers, EXP re-marking is supported only for a single MPLS label packet. [PR1509627](#)

- On ACX710 routers, if the console cable is plugged in and the terminal connection is active and sending characters to the interface, the system boot might be interrupted and boot will be stalled at the uboot# prompt. [PR1513553](#)
- On ACX710 routers, local repair can be in seconds (>50 ms) during FRR convergence. If the explicit NULL is configured on the PHP node and on the PHP node of the backup path, the link failure is observed at PHP node. Global repair resumes the traffic flow. [PR1515512](#)
- Alarm might not be seen on ACX710 routers when the system is booted with recovery snapshot. [PR1517221](#)
- On ACX710 routers, SyncE to 1PPS transient test results do not meet G.8273.2 SyncE to 1PPS transient metric. [PR1522796](#)
- Even though enhanced-ip is active, the following alarm is observed during ISSU: **RE0 network-service mode mismatch between configuration and kernel setting**. [PR1546002](#)
- On ACX5448 and ACX710 routers, the start session ack is delayed by 10 seconds when configured as TWAMP server. [PR1556829](#)
- CoS remarking might not work as expected when three color policer is applied. [PR1559665](#)
- ACX Series does not delete a MAC address from the MAC table if there is traffic destined to the MAC address or traffic sourced from the MAC address or both. The fix will allow ACX to only look at traffic sourced from MAC address before deleting the MAC address entry from MAC table. So, if there is no traffic sourced from the MAC for an interval of MAC aging timer, the MAC would be deleted from the MAC table at the end of MAC aging timer with out taking into account the traffic destined to the MAC address. [PR1565642](#)
- Console and auxiliary ports provide out-of-band remote access to a device. When the console and auxiliary ports are configured as insecure, root login is not allowed to establish terminal connections, and superusers and anyone with a user identifier (UID) of 0 are not allowed to establish terminal connections in multiuser mode. However, ACX710 router has no auxiliary port, out-of-band access is through console port always. By configuring **set system ports auxiliary insecure** statement, ACX710 router reboots with boot reason due to watchdog timeout. [PR1580016](#)

## Virtual Chassis

- On the ACX5000 router, the following false positive parity error message is observed: **soc\_mem\_array\_sbusdma\_read**. The SDK can raise false alarms for parity error messages like this. [PR1276970](#)

SEE ALSO

---

[What's Changed | 22](#)

---

[Known Limitations | 27](#)

---

[Resolved Issues | 35](#)

---

[Documentation Updates | 41](#)

---

[Migration, Upgrade, and Downgrade Instructions | 41](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 35](#)
- [Resolved Issues: 20.2R2 | 36](#)
- [Resolved Issues: 20.2R1 | 39](#)

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

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

### Resolved Issues: 20.2R3

#### General Routing

- The vpls-oam sessions are detected with error (RDI sent by some MEP) after changing VLANs. [PR1478346](#)
- The hardware FRR for EVPN-VPWS, EVPN-FXC, and Layer 3 VPN with a composite next hop are not supported. [PR1499483](#)
- ACX1100, ACX2100, ACX2200, ACX2000, and ACX4000 routers might stop forwarding transit and control traffic. [PR1508534](#)
- On the ACX5448 router, the transit DHCPv4 and DHCPv6 packets drop in a Layer 2 domain. [PR1517420](#)
- On ACX500-I routers, the **show services session count** command does not work as expected. [PR1520305](#)
- Interface does not come up with the **auto-negotiation** setting between ACX1100 router and the other ACX Series routers, MX Series routers and QFX Series switches as the other end. [PR1523418](#)
- With the ACX5448 router with 1000 CFM, the CCM state does not go in the Ok state after loading the configuration or restarting the Packet Forwarding Engine. [PR1526626](#)

- The l2cpd memory leak might be observed with aggregated Ethernet interface flap. [PR1527853](#)
- FEC field is not displayed when the interface is down. [PR1530755](#)
- Packets dropped might be seen after configuring PTP transparent clock. [PR1530862](#)
- The **show class-of-service routing-instance** command does not show the configured classifier. [PR1531413](#)
- On ACX710 routers, the rpd process generates core file at **l2ckt\_vc\_adv\_recv, l2ckt\_adv\_rt\_flash (taskptr=0x4363b80, rtt=0x4418100, rti=< optimized out>, data=< optimized out>, opcode=< optimized out>)** at **../../../../../../../../src/junos/usr.sbin/rpd/l2vpn/l2ckt.c:7982**. [PR1537546](#)
- Management Ethernet link down alarm is observed while verifying system alarms in a Virtual Chassis setup. [PR1538674](#)
- On the ACX5448 router, unexpected behavior of the **show chassis network-services** command is observed. [PR1538869](#)
- The ACX5448 router as transit for the BGP labeled unicast drops traffic. [PR1547713](#)
- PTP slave might discard the PTP packets from the primary when MPLS explicit-null is configured. [PR1547901](#)
- The ARP packets from the CE device are added with VLAN tag if the VLAN-ID is configured in the EVPN routing instance. [PR1555679](#)
- On the ACX5448 router, the unicast packets from the CE devices might be forwarded by the PE devices with additional VLAN tag if IRB is used. [PR1559084](#)
- On the ACX5048 router, the fxpc process generates core file on the analyzer configuration. [PR1559690](#)
- The lo0 firewall is not programmed to Packet Forwarding Engine and **ACX\_DFW\_CFG\_FAILED: ACX Error (dfw):dnx\_policer\_create : Policer Creation Failed No resources for operation** message is seen. [PR1566417](#)
- The lcksyncd log file will show empty eventually. [PR1567687](#)
- On the ACX5448 router, untagged traffic is being incorrectly queued and marked. [PR1570899](#)
- RFC2544 reflector feature might not work on higher port (for example, port 46). [PR1571975](#)

### **Layer 2 Features**

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

## **Resolved Issues: 20.2R2**

### **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](#)

- While performing repeated power-off or power-on of the device, SMBUS transactions timeout is observed. [PR1463745](#)
- On the ACX5048 router, the egress queue statistics do not work for the aggregated Ethernet interfaces. [PR1472467](#)
- On ACX710 routers, VPLS OAM sessions are detected with error (remote defect indication sent by some MEPs) after changing VLANs. [PR1478346](#)
- 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 ACX5048 and ACX5096 routers, the LACP control packets might get dropped due to high CPU utilization. [PR1493518](#)
- 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](#)
- On the ACX710 router, high convergence is observed with the EVPN-ELAN service in a scaled scenario during FRR switchover. [PR1497251](#)
- Outbound SSH connection flaps or memory leaks occur 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](#)
- On the ACX5448 router, the EXP rewrite for the Layer 3 VPN sends all traffic with incorrect EXP. [PR1500928](#)
- SFP-T is unrecognized after FPGA upgrade and power cycle. [PR1501332](#)
- The error message **mpls\_extra NULL** might be seen when you add, change, or delete MPLS route. [PR1502385](#)
- 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 clients 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\_I3\_egress\_create : Unable to create egress object.** [PR1513644](#)
- The loopback filter cannot take more than two TCAM slices. [PR1513998](#)

- 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](#)
- Tagged traffic matching the vlan-list configuration in the vlan-ccc logical interface gets dropped in the ingress interface. [PR1519568](#)
- 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](#)
- The `vlan-id-list` statement might not work as expected on the ACX5448 and ACX710 platforms. [PR1527085](#)
- The `show class-of-service routing-instance` command does not show configured classifier on ACX Series platforms. [PR1531413](#)
- Memory leak in local OutLif in VPLS and CCC topology. [PR1532995](#)
- Management Ethernet link down alarm is seen while verifying system alarms in a Virtual Chassis setup. [PR1538674](#)

### **Interfaces and Chassis**

- The FPC crash might be observed with inline mode CFM configured. [PR1500048](#)

### **Routing Protocols**

- The rpd process might report 100 percent CPU usage with BGP route damping enabled. [PR1514635](#)

## **Resolved Issues: 20.2R1**

### **General Routing**

- Drift messages in ACX2200, which is a PTP hybrid (PTP + Synchronous Ethernet) device. [PR1426910](#)
- ACX5448-D interfaces support: The input bytes value for the **show interfaces extensive** command is not at par with older ACX Series or MX Series devices. [PR1430108](#)
- On an ACX5448 device, DHCP packets are not transparent over Layer 2 circuit. [PR1439518](#)
- On an ACX5048 device, SNMP polling stops after the link is flapped or the SFP transceiver is replaced, and **ACX\_COS\_HALP(acx\_cos\_gport\_sched\_set\_strict\_priority:987): Failed to detach** logs might be seen. [PR1455722](#)
- ACX5448-D and ACX5448-M devices do not display airflow information and temperature sensors as expected. [PR1456593](#)
- Unable to get shared buffer count as expected. [PR1468618](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- On an ACX710 device, MPLS packet load balancing is done without hashing enabled. [PR1475363](#)
- FPC might continuously crash after deactivating or activating loopback filter or reboot the system after configuring the loopback filter. [PR1477740](#)
- The dcpfe core file is generated when disabling or enabling MACsec through Toby scripts. [PR1479710](#)
- Link does not come up when a 100-Gigabit Ethernet port is channelized into four port 25-Gigabit Ethernet interfaces. [PR1479733](#)
- Memory utilization enhancement on ACX platforms to reduce the memory foot print. [PR1481151](#)
- On ACX5448 devices, **dnx\_nh\_mpls\_tunnel\_install** logs are seen. [PR1482529](#)
- ACX AUTHD process memory usage is 15 percent. [PR1482598](#)
- FPC crash is seen on ACX5448 platform. [PR1485315](#)
- On an ACX5448 device, Layer 2 VPN with interface ethernet-ccc **input-vlan-map/output-vlan-map** can cause traffic to be discarded silently. [PR1485444](#)
- On the ACX710 router, VPLS flood group results in IPv4 traffic drop after core interface flap. [PR1491261](#)
- On the ACX710 routers, LSP (primary and standby) does not Act/Up after routing or rpd restart. [PR1494210](#)



- During speed mismatch, QSFP28/QSFP+ optics/cables might or might not work. [PR1494600](#)
- ACX710 BFD sessions are in initialization state with CFM scale of 1000 on reboot or chassis control restart. [PR1503429](#)
- On an ACX500-i router, SFW sessions are not getting updated on ms- interfaces. [PR1505089](#)
- On an ACX710 router, wavelength changed from CLI does not take effect in tunable optics. [PR1506647](#)
- PIC slot might be shut down in less than 240 seconds due to the over-temperature start time is handled incorrectly. [PR1506938](#)
- BFD flaps with the error **ACX\_OAM\_CFG\_FAILED: ACX Error (oam):dnx\_bfd\_l3\_egress\_create: Unable to create egress object** after random time interval. [PR1513644](#)

### ***Interfaces and Chassis***

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

### ***Layer 2 Ethernet Services***

- Member links state might be asynchronized on a connection between a PE device and a CE device in an EVPN active/active scenario. [PR1463791](#)

### ***MPLS***

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

### ***Routing Protocols***

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

### ***VPNs***

- The Layer 2 circuit neighbor might be stuck in RD state at one end of the MG-LAG peer. [PR1498040](#)
- The rpd core files are generated while disabling Layer 2 circuit with connection protection, backup neighbor configuration, and Layer 2 circuit trace logs enabled. [PR1502003](#)

### SEE ALSO

---

[What's New | 13](#)

---

[What's Changed | 22](#)

---

[Known Limitations | 27](#)

---

[Open Issues | 31](#)

---

[Documentation Updates | 41](#)

[Migration, Upgrade, and Downgrade Instructions | 41](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 documentation for ACX Series routers.

### SEE ALSO

[What's New | 13](#)

[What's Changed | 22](#)

[Known Limitations | 27](#)

[Open Issues | 31](#)

[Resolved Issues | 35](#)

[Migration, Upgrade, and Downgrade Instructions | 41](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 41](#)

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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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

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

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

#### SEE ALSO

[What's New | 13](#)

[What's Changed | 22](#)

[Known Limitations | 27](#)

[Open Issues | 31](#)

[Resolved Issues | 35](#)

[Documentation Updates | 41](#)

## Junos OS Release Notes for cSRX

#### IN THIS SECTION

- [What's New | 43](#)
- [What's Changed | 43](#)
- [Known Limitations | 44](#)
- [Open Issues | 44](#)
- [Resolved Issues | 44](#)

These release notes accompany Junos OS Release 20.2R3 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

- [What's New in Release 20.2R3 | 43](#)
- [What's New in Release 20.2R2 | 43](#)

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

### What's New in Release 20.2R3

There are no new features for cSRX in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features for cSRX in Junos OS Release 20.2R2.

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 44](#)
- [What's Changed in Release 20.2R2 | 44](#)

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

### What's Changed in Release 20.2R3

There are no changes in behavior or syntax for cSRX in Junos OS Release 20.2R3.

### What's Changed in Release 20.2R2

There are no changes in behavior or syntax for cSRX in Junos OS Release 20.2R2.

## Known Limitations

There are no known behavior or limitation for cSRX in Junos OS Release 20.2R3.

## Open Issues

There are no known issues for cSRX in Junos OS Release 20.2R3.

## Resolved Issues

### Resolved Issues: 20.2R3

There are no resolved issues for cSRX in Junos OS Release 20.2R3.

### Resolved Issues: 20.2R2

There are no resolved issues for cSRX in Junos OS Release 20.2R2.

# Junos OS Release Notes for EX Series

## IN THIS SECTION

- [What's New | 45](#)
- [What's Changed | 54](#)

- Known Limitations | 58
- Open Issues | 59
- Resolved Issues | 63
- Documentation Updates | 70
- Migration, Upgrade, and Downgrade Instructions | 71

These release notes accompany Junos OS Release 20.2R3 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

### IN THIS SECTION

- What's New in Release 20.2R3 | 45
- What's New in Release 20.2R2 | 46
- What's New in Release 20.2R1-S1 | 46
- What's New in Release 20.2R1 | 46

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

**NOTE:** The following EX Series switches are supported in Release 20.2R3: EX2300, EX2300-C, EX3400, EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

### What's New in Release 20.2R3

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 20.2R3.

## What's New in Release 20.2R2

There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 20.2R2.

## What's New in Release 20.2R1-S1

### *Software Installation and Upgrade*

- **Zero touch provisioning (ZTP) with IPv6 support (EX3400, EX4300, QFX5100 and QFX5200 switches, MX-Series routers)**—Starting in Junos OS Release 20.2R1-S1, ZTP supports the DHCPv6 client. During the bootstrap process, the device first uses the DHCPv4 client to request for information regarding image and configuration file from the DHCP server. The device checks the DHCPv4 bindings sequentially. If there is a failure with one of the DHCPv4 bindings, the device will continue to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device will check for DHCPv6 bindings and follow the same process as for DHCPv4 until the device can be provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

**NOTE:** Only HTTP and HTTPS transport protocols are supported EX3400, EX4300, QFX5100, and QFX5200 devices.

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

## What's New in Release 20.2R1

### *Authentication, Authorization, and Accounting*

- **Retain the authentication session based on DHCP or SLAAC snooping entries (EX2300, EX3400, and EX4300)**—Starting in Junos OS Release 20.2R1, you can configure the authenticator to check for a DHCP, DHCPv6, or SLAAC snooping entry before terminating the authentication session when the MAC address ages out. If a snooping entry is present, the authentication session for the end device with that MAC address remains active. This ensures that the end device will be reachable even if the MAC address ages out.

[See [Authentication Session Timeouts](#).]

### *EVPN*

- **802.1X authentication with EVPN-VXLAN (EX4300-48MP and EX4300-48MP Virtual Chassis)**—Starting in Junos OS Release 20.2R1, EX4300-48MP switches that act as access switches can use 802.1X authentication to protect an EVPN-VXLAN network from unauthorized end devices. EX4300-48MP switches support the following 802.1X authentication features on access and trunk ports:

- Access ports: single, single-secure, and multiple supplicant modes
- Trunk ports: single and single-secure supplicant modes
- Guest VLAN
- Server fail
- Server reject
- Dynamic VLAN
- Dynamic firewall filters
- RADIUS accounting
- Port bounce with Change of Authorization (CoA) requests
- MAC RADIUS client authentication
- Central Web Authentication (CWA) with redirect URL
- Captive portal client authentication
- Flexible authentication with fallback scenarios

[See [802.1X Authentication](#).]

- **Support for firewall filtering on EVPN-VXLAN traffic (EX4300-MP)**—Starting with Junos OS Release 20.2R1, you can configure firewall filters and policers on the VXLAN traffic in an EVPN network (EVPN-VXLAN traffic). You set the rules that the devices uses to accept or discard packets by defining the terms for a firewall filter. For filters that you would apply to a port or VLAN, configure firewall filters at the **[edit firewall family ethernet-switching]** hierarchy level. For filters that you would apply to an IRB interface, configure firewall filters at the **[edit firewall family inet]** hierarchy level. After a firewall filter is defined, you can then apply it at an interface.

[See [Firewall Filtering and Policing Support for EVPN-VXLAN](#).]

- **Noncolored SR-TE LSPs with EVPN-MPLS (ACX5448, EX9200, MX Series, and vMX)**—Starting in Junos OS Release 20.2R1, ACX5448, EX9200, MX Series, and vMX routers support noncolored static segment routing-traffic engineered (SR-TE) label-switched paths (LSPs) with an EVPN-MPLS core network and the following Layer 2 services running at the edges of the network:

- E-LAN
- EVPN-ETREE
- EVPN-VPWS with E-Line

Without color, all LSPs resolve using a BGP next hop only.

The Juniper Networks routers support noncolored SR-TE LSPs in an EVPN-MPLS core network with the following configurations:

- EVPN running in a virtual switch routing instance



- Multihoming in active/active and active/standby modes

The Juniper Networks routers also support noncolored SR-TE LSPs when functioning as a Data Center Interconnect (DCI) device that handles EVPN Type 5 routes.

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

- **MAC filtering, storm control, and port mirroring support in EVPN-VXLAN overlay networks (EX4300-48MP)**—Starting with Junos OS Release 20.2R1, EX4300-48MP switches support the following features in an EVPN-VXLAN overlay network:

- MAC filtering
- Storm control
- Port mirroring and analyzers

[See [MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.](#)]

- **Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface (EX4600)**—Starting in Junos OS Release 20.2R1, you can configure and successfully commit the following on a physical interface of an EX4600 switch in an EVPN-VXLAN environment:

- Layer 2 bridging (**family ethernet-switching**) on any logical interface unit number (unit 0 and any nonzero unit number).
- VXLAN on any logical interface unit number (unit 0 and any nonzero unit number).
- Layer 2 bridging (**family ethernet-switching** and **encapsulation vlan-bridge**) on different logical interfaces (unit 0 and any nonzero unit number).
- Layer 3 IPv4 routing (**family inet**) and VXLAN on different logical interfaces (unit 0 and any nonzero unit number).

For these configurations to be successfully committed and work properly, you must specify the **encapsulation flexible-ethernet-services** configuration statements at the physical interface level—for example, **set interfaces xe-0 /0/5 encapsulation flexible-ethernet-services**.

[See [Understanding Flexible Ethernet Services Support With EVPN-VXLAN.](#)]

### *High Availability (HA) and Resiliency*

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes roles. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

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

### *Juniper Extension Toolkit (JET)*

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

[See [language \(Scripts\)](#), [Develop Off-Device JET Applications](#), and [Develop On-Device JET Applications](#).]

### *Junos OS XML, API, and Scripting*

- **Support for Rest API (EX2300, EX2300-MP, EX3400, EX4300, EX4300-MP, EX4600, EX4650, and EX9200)**—Starting in Release 20.2R1, Junos OS supports the REST API on EX2300, EX2300-MP, EX3400, EX4300, EX4300-MP, EX4600, EX4650, and EX9200 switches. The REST API enables you to securely connect to the Junos OS devices, execute remote procedure calls (RPC) commands, use REST API explorer GUI to conveniently experiment with any of the REST APIs, and use a variety of formatting and display options including JavaScript Object Notation (JSON).

[See [REST API Guide](#).]

### *Junos Telemetry Interface*

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

- **Support for OpenConfig configuration model version 4.0.1 for BGP with JTI (EX2300, EX3400, EX4300, EX4600, and EX9200)**— Junos OS Release 20.2R1 provides support for the OpenConfig version 4.0.1 data models **openconfig-bgp-neighbor.yang** and **openconfig-bgp-policy.yang** using Junos telemetry

interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream telemetry statistics to an outside collector.

The following major resource paths are supported with gRPC and JTI:

- `/network-instances/network-instance/protocols/protocol/bgp/global/`
- `/network-instances/network-instance/protocols/protocol/bgp/global/afi-safis/afi-safi/`
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/`
- `/network-instances/network-instance/protocols/protocol/bgp/peer-groups/peer-group/`

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

- **Support for OpenConfig configuration model version 1.0.0 for local routing with JTI (EX2300, EX3400, EX4300, EX4600, and EX9200)**— Junos OS Release 20.2R1 provides support for the OpenConfig version 1.0.0 data model `openconfig-local-routing.yang` using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream telemetry statistics to an outside collector.

The following major resource paths are supported with gRPC and JTI:

- `/local-routes/static-routes/static/`
- `/local-routes/local-aggregates/aggregate/`

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

- **Packet Forwarding Engine and Routing Engine sensor support with JTI (EX2300, EX2300-MP, and EX3400)**— Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) with remote procedure call (gRPC) services to export Packet Forwarding Engine statistics and Routing Engine statistics from EX2300, EX2300-MP, and EX3400 switches to an outside collector. These statistics can also be exported through UDP (native) sensors.

Supported Packet Forwarding Engine sensors are:

- Sensor for CPU (ukernel) memory (resource path `/junos/system/linecard/cpu/memory/`)
- Sensor for firewall filter statistics (resource path `/junos/system/linecard/firewall/`)
- Sensor for physical interface traffic (resource path `/junos/system/linecard/interface/`)
- Sensor for logical interface traffic (resource path `/junos/system/linecard/interface/logical/usage/`). Not supported on EX2300 or 2300-MP switches.
- Sensor for software-pollled queue-monitoring statistics (resource path `/junos/system/linecard/qmon-sw/`). Not supported on EX2300 or 2300-MP switches.

Supported Routing Engine sensors are:

- Sensor for LACP state export (resource path `/lACP/`)
- Sensor for chassis environmentals export (resource path `/junos/system/components/component/`)
- Sensor for chassis components export (resource path `/components/`)
- Sensor for LLDP statistics export (resource path `/lldp/interfaces/interface[name='name']/`)
- Sensor for BGP peer information export (resource path `/network-instances/network-instance/protocols/protocol/bgp/`). Not supported on EX2300 or 2300-MP switches.
- Sensor for RPD task memory utilization export (resource path `/junos/task-memory-information/`)
- Sensor network discovery ARP table state (resource path `/arp-information/`)
- Sensor for network discovery NDP table state (resource path `/nd6-information/`)

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

### Layer 2 Features

- **L2PT support (EX4650 and QFX5120-48Y switches, and QFX5100 and QFX5110 switches and Virtual Chassis)**—Starting in Junos OS Release 20.2R1, you can configure Layer 2 protocol tunneling (L2PT) to tunnel any of the following Layer 2 protocols: CDP, E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, STP (including RSTP and MSTP), UDLD, VSTP, and VTP.

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

### Multicast

- **Static multicast route leaking for VRF and virtual router instances (EX4650 and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can configure the switch to statically share (leak) IPv4 multicast routes for IGMPv3 (S,G) traffic among different virtual router or virtual routing and forwarding (VRF) instances. You can only leak static multicast routes per group, not per source and group. The destination prefix length must be 32.

To configure multicast route leaking to the VRF or virtual router instance *routing-instance-name*, configure the **next-table *routing-instance-name.inet.0*** statement at the **[edit routing-instances *routing-instance-name* routing-options static route destination-prefix/32]** hierarchy level.

[See [Understanding Multicast Route Leaking for VRF and Virtual Router Instances](#).]

- **Multicast-only fast reroute (MoFRR) (EX4650 and QFX5120-48Y)**—Starting in Junos OS Release 20.2R1, you can configure MoFRR to minimize multicast packet loss in PIM domains when link failures occur. With MoFRR enabled, the switch maintains primary and backup traffic paths, forwarding traffic from the primary path and dropping traffic from the backup path. If the primary path fails, the switch can quickly start forwarding the backup path stream (which becomes the primary path). The switch creates a new backup path if it detects available alternative paths. MoFRR applies to all multicast (S,G) streams by default, or you can configure a policy for the (S,G) entries where you want MoFRR to apply.

[See [Understanding Multicast-Only Fast Reroute](#).]

### **Network Management and Monitoring**

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

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

### **Routing Policy and Firewall Filters**

- **Support for MPLS firewall filter on loopback interface (EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can apply an MPLS firewall filter to a loopback interface on a Label switching router (LSR). For example, you can configure an MPLS packet with `ttl=1` along with MPLS qualifiers such as `label`, `exp`, and Layer 4 `tcp/udp` port numbers. Supported actions include `accept`, `discard`, and `count`.

You configure this feature at the `[edit firewall family mpls]` hierarchy level. You can only apply a loopback filters on `family mpls` in the ingress direction.

[See [Overview of MPLS Firewall Filters on Loopback Interface.](#)]

### **Routing Protocols**

- **Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast (MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices)**—Starting with Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices support BGP PIC Edge protection for Layer 2 circuit, Layer 2 VPN, and VPLS (BGP VPLS, LDP VPLS and FEC 129 VPLS) services with BGP labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

A prerequisite for BGP PIC Edge protection is to program the Packet Forwarding Engine (PFE) with expanded next-hop hierarchy.

To enable BGP PIC Edge protection, use the following CLI configuration statements:

- Expand next-hop hierarchy for BGP labeled unicast family:

```
[edit protocols]
user@host#set bgp group group-name family inet labeled-unicast nexthop-resolution
preserve-nexthop-hierarchy;
```

- BGP PIC for MPLS load balance nexthops:

```
[edit routing-options]
user@host#set rib routing-table-name protect core;
```

- Fast convergence for Layer 2 circuit and LDP VPLS:

```
[edit protocols]
user@host#set l2circuit resolution preserve-nexthop-hierarchy;
```

- Fast convergence for Layer 2 VPN, BGP VPLS, and FEC129:

```
[edit protocols]
user@host#set l2vpn resolution preserve-nexthop-hierarchy;
```

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

## SEE ALSO

[What's Changed | 54](#)

[Known Limitations | 58](#)

[Open Issues | 59](#)

[Resolved Issues | 63](#)

[Documentation Updates | 70](#)

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

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 54](#)
- [What's Changed in Release 20.2R2 | 56](#)
- [What's Changed in Release 20.2R1 | 56](#)

Learn about what changed in this release for EX Series Switches.

### What's Changed in Release 20.2R3

#### *Junos OS XML API and Scripting*

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

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

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

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

#### *Network Management and Monitoring*

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the **edit system services netconf hello-message yang-module-capabilities** hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the

appropriate statements at the **edit system services netconf netconf-monitoring netconf-state-schemas** hierarchy level.

[See [hello-message](#).]

[See [netconf-monitoring](#).]

- **Changes to <commit> RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—When you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level, the NETCONF server's response for **<commit>** operations includes the following changes:
  - If a successful **<commit>** operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the **<source-daemon>** element as a child of the **<error-info>** element instead of the **<rpc-error>** element.
  - If you also configure the **flatten-commit-results** statement at the **[edit system services netconf]** hierarchy level, the NETCONF server suppresses any **<commit-results>** XML subtree in the response and emits only an **<ok>** or **<rpc-error>** element.

[See [Configuring RFC-Compliant NETCONF Sessions](#).]



### *User Interface and Configuration*

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

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

### **What's Changed in Release 20.2R2**

#### *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.

#### *Routing Protocols*

- **Advertising /32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the **Isdist.0** and **Isdist.1** routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the **Isdist.0** and **Isdist.1** routing tables as part of node characteristics and advertised as router IDs.

### **What's Changed in Release 20.2R1**

#### *General Routing*

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the **[edit system commit]** hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

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

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

[See [show system resource-monitor](#) and [Resource Monitoring for Subscriber Management and Services](#).]

### *Juniper Extension Toolkit (JET)*

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the PASS keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

[See [Develop Off-Device JET Applications](#) and [Develop On-Device JET Applications](#).]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH\_USAGE\_INVALID**.

[See [Juniper EngNet](#).]

### *Network Management and Monitoring*

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

### SEE ALSO

---

[What's New | 45](#)

---

[Known Limitations | 58](#)

---

[Open Issues | 59](#)

---

[Resolved Issues | 63](#)

---

[Documentation Updates | 70](#)

---

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

## Known Limitations

### IN THIS SECTION

- [EVPN | 58](#)
- [General Routing | 58](#)
- [Infrastructure | 58](#)
- [Layer 2 Ethernet Services | 59](#)

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

- When only one link is present between the leaf devices, it goes down, resulting in traffic drop. [PR1480847](#)
- InterVNI multicast is not supported in EVPN-VXLAN edge routing model on EX4650. [PR1517082](#)

### General Routing

- Junos OS might hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. As a workaround, you can power cycle the device. [PR1385970](#)
- The interfaces on certain EX9251 line of switches might get stuck in a down state, if the remote interface sends invalid code to the local interface. Link might not come up even after the remote peer has begun sending a good signal. The "Failed to complete DFE tuning" syslog might appear. This syslog message has no functional impact. [PR1473280](#)
- On all Junos OS platforms, in a QinQ environment, xSTP is enabled on the interface having logical interface with vlan-id-list configured, then it will only run on those logical interfaces whose vlan-id range includes native-vlan-id configured and all others will in discarding state. This might lead to traffic drop. [PR1532992](#)

### Infrastructure

- Depending on the actual traffic pattern and the order in which the MACs are learned, the actual MAC DB scale may vary. This is due to the way the MACs are internally stored in the hardware. [PR1485319](#)

- On EX-4300MP, 9000 IPv6 MC routes can be installed. If you try to add more IPv6 MC routes, error messages will be seen. [PR1493671](#)
- EX4650 ASIC uses a static hashing and RTAG7 hash algorithm that might be alike on each chipset. Hence, we recommend that you fine-tune hash parameters based on the traffic profile used when deviation in load balance is observed. On TD3 chipset based platforms, the following configuration is required to fine-tune hashing deviation; 1. set forwarding-options enhanced-hash-key hash-parameters ecmp offset 29. 2. set forwarding-options enhanced-hash-key hash-parameters ecmp preprocess. [PR1516883](#)

## Layer 2 Ethernet Services

- Sometimes image upgrade through ZTP might fail because of the insufficient space on EX3400. For information on how to free up the space see [KB31198](#). [PR1515013](#)

### SEE ALSO

[What's New | 45](#)

[What's Changed | 54](#)

[Open Issues | 59](#)

[Resolved Issues | 63](#)

[Documentation Updates | 70](#)

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

## Open Issues

### IN THIS SECTION

- [General Routing | 60](#)
- [Infrastructure | 61](#)
- [Interfaces and Chassis | 62](#)
- [Layer 2 Features | 62](#)
- [Layer 2 Ethernet Services | 62](#)
- [Platform and Infrastructure | 62](#)
- [Routing Protocols | 63](#)

Learn about open issues in Junos OS Release 20.2R3 for EX Series switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

## General Routing

- On the MX204 and MX10003 routers, the following garbage value on syslog messages from craftd demon is observed: **craftd[xxxx]: fatal error, failed to open smb device: JÏÈ**. [PR1359929](#)
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. [PR1362609](#)
- On EX2300, when watchdog is induced, the last reboot reason is shown as Swizzle Reboot. [PR1369924](#)
- On an EX9208 switch, a few xe- interfaces go down with the error message **if\_msg\_ifd\_cmd\_tlv\_decode ifd xe-0/0/0 #190 down with ASIC Error**. [PR1377840](#)
- On EX4300-48MP, EX2300-24T, and EX4650 platforms, either unicast RPF in strict mode or ICMP redirect does not work properly. [PR1417546](#)
- On the EX9214 device, if the MACsec-enabled link flaps after reboot, the error **errorlib\_set\_error\_log(): err\_id(-1718026239)** is observed. [PR1448368](#)
- On Junos OS platforms with next generation Routing Engine installed, the vhostd process might crash without generating a core file and automatic restart might fail. [PR1448413](#)
- In overall commit time, the evaluation of mustd constraints is taking two seconds more than usual. This is because the persist-group-inheritance feature has been made a default feature in the latest Junos OS releases. Eventually, this feature helps improve the subsequent commit times for scaled configurations significantly. The persist-group-inheritance feature is useful in customer scenarios where groups and nested groups are used extensively. In those scenarios, the group inheritance paths are not built every time, thus subsequent commits are faster. [PR1457939](#)
- EX2300-48MP Virtual Chassis is rebooted silently and randomly without generating a core file. Syslogs and console logs are not generated before rebooting the switch, because the reboot reason is shown as a normal reboot. [PR1463583](#)
- On EX4300 switches, when packets entering a port exceed a size of 144 bytes, they might get dropped in few cases. [PR1464365](#)
- On EX4650 platform, after using force reboot, the output of CLI command 'show version' might show the model as QFX5120-48y-8c and after committing the http services, J-Web of the device might be inaccessible due to model issue. [PR1480252](#)
- On BCM Packet Forwarding Engine-based EX Series platforms frame higher than MTU+4 and lesser than MTU+8 bytes, with invalid FCS, code error, or IEEE length check error, is treated as Jabber frame. [PR1487709](#)
- On EX Series platforms using chipset with SFP+ implemented, interface on the platforms might be in active status when TX or RX connector is removed. As a result, traffic might drop. [PR1495564](#)

- SNMP POE MIB walk produce withers no results or sometimes result from the master Virtual Chassis whenever the Virtual Chassis is renamed as one. [PR1503985](#)
- On the EX4300-48MP device, the reboot time, FPC uptime, and interface uptime are degraded by 20 percent when compared with Junos OS Releases 19.1R3, 19.2R2, and 19.4R2. [PR1514364](#)
- Traffic not load balanced by EX4300-48MP and EX4300-VC over ESI links with evpn\_vxlan configured. [PR1550305](#)
- On the EX4300 device, script fails while committing the IPsec authentication configuration due to the missing algorithm statement. [PR1557216](#)
- When dot1x **server-fail-voip vlan-name** is configured, ensure that both server-fail-voip vlan-name and voip vlan are configured using vlan name and not by using vlan-id. [PR1561323](#)
- On EX4600 platform, internal comment 'Placeholder for QFX platform config' might be seen on **show config** output. [PR1567037](#)

## Infrastructure

- On EX Series switches except EX4300/EX4600/EX9200, an interface is configured for single VLAN or multiple VLANs, if all these VLANs of this interface have igmp-snooping enabled, then this interface will drop hot standby router protocol for IPv6 (HSRPv2) packets. But, if some VLANs do not have igmp-snooping enabled, then this interface works fine. [PR1232403](#)
- On EX Series switches, If you are configuring a large-scale number of firewall filters on some interfaces, the FPC might crash and generate core files. [PR1434927](#)
- **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](#)
- Power loss during software install can leave artifacts that consume space. These need to be included in package cleanup procedure. [PR1544222](#)

## Interfaces and Chassis

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

## Layer 2 Features

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

## Layer 2 Ethernet Services

- If **forward-only** is set within **dhcp-reply** in a Juniper Networks device as a DHCP relay agent, the DHCP DECLINE packets that are broadcasted from the DHCP client are dropped and not forwarded to the DHCP server. [PR1429456](#)
- OSPF and OSPF3 adjacency uptime is more than expected after NSSU upgrade and outage is higher than the expected. [PR1551925](#)

## Platform and Infrastructure

- On the EX9208 device, 33 percent degradation with MAC learning rate is observed in Junos OS Release 19.3R1 compared to Junos OS Release 18.4R1. [PR1450729](#)
- On EX4300 platforms configured with ERP, after multiple devices reboot/restart at the same time, ERP might not revert back to the IDLE state. This issue might be seen in situations where the ERP node-id is not configured manually and after the restart, the default node-id (switch base MAC address) might get reset to 00:00:00:00:00:00, effectively causing multiple devices to have the same node-id. [PR1461434](#)
- The pfex\_junos process generates core file at 0x01847994 in pfeman\_watchdog (arg=< optimized out>) at `../../../../src/pfe/common/applications/pfeman/pfeman_rt_pfex.c:1411`. [PR1535178](#)
- Upgrading satellite devices might lead to some SDs in SyncWait state. Cascade port flap not causing the issue. [PR1556850](#)
- "Last flapped" timestamp for interface fxp0 gets reset every time "monitor traffic interface fxp0" is executed. [PR1564323](#)
- On all EX9200 platforms with EVPN-VXLAN configured, the next-hop memory leak in MX Series ASIC happens whenever there is a route churn for remote MAC-IP entries learned bound to the IRB interface in EVPN-VXLAN routing instance. When the ASIC's next-hop memory partition exhausted the FPC might reboot. [PR1571439](#)

## Routing Protocols

- Verifying loader only uses ECDSA256+SHA256 for integrity checks but does not say so. [PR1504211](#)

### SEE ALSO

[What's New | 45](#)

[What's Changed | 54](#)

[Known Limitations | 58](#)

[Resolved Issues | 63](#)

[Documentation Updates | 70](#)

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

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 63](#)
- [Resolved Issues: 20.2R2 | 65](#)
- [Resolved Issues: 20.2R1 | 67](#)

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

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

### Resolved Issues: 20.2R3

#### *General Routing*

- IRB MAC will not be programmed in hardware when MAC persistence timer expires. [PR1484440](#)
- While verifying the last-change op-state value through XML, the rpc-reply message is inappropriate. [PR1492449](#)
- The mge interface might still stay up while the far end of the link goes down. [PR1502467](#)



- The output VLAN push might not work. [PR1510629](#)
- DHCP traffic might not be forwarded correctly when sending DHCP unicast packets. [PR1512175](#)
- EX4300-48MP device might go out of service during a software upgrade operation. [PR1526493](#)
- On the EX2300 device, the following PoE message is observed **poe\_get\_dev\_class: Failed to get PD class info.** [PR1536408](#)
- The LLDP neighborship with the VoIP phones cannot be established. [PR1538482](#)
- On the EX3400 and EX2300 switches, the upgrade fails due to the lack of available storage. [PR1539293](#)
- FPC might not be recognized after power cycle (hard reboot). [PR1540107](#)
- DHCP discover packet might be dropped if DHCP inform packet is received first. [PR1542400](#)
- Slaac-Snoopd child process generates a core file upon multiple switchovers on the Routing Engine. [PR1543181](#)
- In every software upgrade host needs to get upgrade. [PR1543890](#)
- On EX4300-48MP line of switches with Linux TVP architecture and Junos OS as VM, the Junos CLI outputs do not confirm if the Junos OS and the host kernel are compatible with each other. [PR1543901](#)
- The chip on FPC linecard might crash when the system reboots. [PR1545455](#)
- "show pfe route summary hw" shows random high free and 'Used' column for 'IPv6 LPM(< 64)' routes. [PR1552623](#)
- The statement 'action-shutdown' of storm control does not work for ARP broadcast packets. [PR1552815](#)
- Traffic might be dropped when a firewall filter rule uses 'then vlan' as the action. [PR1556198](#)
- On EX3400VC line of switches, the DAEMON-7-PVIDB throws syslog messages for every 12 to 14 minutes after you upgrade to Junos OS Release 19.1R3-S3. [PR1563192](#)
- Client authentication is failing after performing GRES. [PR1563431](#)

### **Infrastructure**

- On the EX4600 and EX4300 Virtual Chassis or Virtual Chassis fabric, the VSTP configurations device goes unreachable and becomes nonresponsive after commit. [PR1520351](#)
- EX 4300 VC/VCF : Observing HEAP malloc(0) detected. [PR1546036](#)
- Traffic related to IRB interface might be dropped when mac-persistence-timer expires [PR1557229](#)

### **Platform and Infrastructure**

- DHCP binding is not happening after graceful switchover. [PR1515234](#)
- lldp-receive-packet-count is not getting exchanged properly in l2pt operation for lldp after configuring protocols. [PR1532721](#)
- LLDP neighborship might not come up on EX4300 non-aggregated Ethernet interfaces. [PR1538401](#)

- The targeted-broadcast feature might not work after a reboot. [PR1548858](#)
- The BGP session replication might fail to start after the session crashes on a backup Routing Engine. [PR1552603](#)
- The targeted-broadcast feature may send out duplicate packets. [PR1553070](#)

#### ***Routing Protocols***

- The OSPFv3 adjacency should not be established when IPsec authentication is enabled. [PR1525870](#)
- DCPFE crash might be observed while updating VRF for multicast routes during irb uninit. [PR1546745](#)
- Sending multicast traffic to downstream receiver on MX Series-based Virtual Chassis platforms might fail. [PR1555518](#)
- The untagged packets might not work on EX Series platforms. [PR1568533](#)

#### ***User Interface and Configuration***

- The license errors may get returned on backup RE when trying to commit the configuration. [PR1543037](#)

#### ***Virtual Chassis***

- EX4600/EX4300 mixed VC : Error message 'ex\_bcm\_pic\_eth\_uint8\_set' is seen when changing configuration related to interface. [PR1573173](#)

### **Resolved Issues: 20.2R2**

#### ***Authentication and Access Control***

- The DOT1XD\_AUTH\_SESSION\_DELETED event is not triggered with a single supplicant mode. [PR1512724](#)
- The dot1x client won't be moved to held state when the authenticated PVLAN is deleted. [PR1516341](#)

#### ***EVPN***

- Unable to create a new VTEP interface. [PR1520078](#)

#### ***General Routing***

- Virtual Chassis split after network topology is changed. [PR1427075](#)
- EX2300 Series: High CPU load due to receipt of specific multicast packets on Layer 2 interface (CVE-2020-1668). [PR1491905](#)
- Authentication session might be terminated if PEAP request is retransmitted by the authenticator. [PR1494712](#)
- The fxpc might crash when renumbering the master member id value of the EX2300/EX3400 Virtual Chassis. [PR1497523](#)

- Outbound SSH connection flaps or memory leaks occur during the 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](#)
- 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 in restart. [PR1498903](#)
- On the EX4300, EX3400, and EX2300 Virtual Chassis with NSB and xSTP enabled, continuous traffic loss might be observed while performing GRES. [PR1500783](#)
- The mge interface might still stay up while the far end of its link goes down. [PR1502467](#)
- LLDP is not acquired when native-vlan-id and tagged VLAN-ID are the same on a port. [PR1504354](#)
- The output VLAN push might not work. [PR1510629](#)
- LLDP might not work when PVLAN is configured on EX Series and QFX Series Virtual Chassis. [PR1511073](#)
- Traffic might not flow as per configured policer parameters. [PR1512433](#)
- LACP goes down after performing Routing Engine switchover if MACsec is enabled on the LAG members on EX4300. [PR1513319](#)
- The 100M SFP-FX is not supported on satellite device in Junos fusion setup. [PR1514146](#)
- A "dot1x" memory leak is observed. [PR1515972](#)
- The dcpfe (PFE) process might crash due to memory leak. [PR1517030](#)
- MPPE-Send or Recv-key attribute is not extracted correctly by dot1xd. [PR1522469](#)
- "Drops" and "Dropped packets" counters in the output for "show interface extensive" are double-counted. [PR1525373](#)

### **Infrastructure**

- The qmon-sw sensor is not supported in EX3400. [PR1506710](#)
- The IP communication between directly connected interfaces on EX4600 might fail. [PR1515689](#)
- OID ifOutDiscards reports zero and sometimes shows valid value. [PR1522561](#)

### **Layer 2 Features**

- On the QFX5000 line of switches, traffic imbalance might be observed if hash-params is not configured. [PR1514793](#)
- The MAC address in the hardware table might become out of synchronization between the master and member in Virtual Chassis after the MAC flaps. [PR1521324](#)

### **Platform and Infrastructure**

- Packets get dropped when next hop is IRB over an It interface. [PR1494594](#)

- LLDP neighborship might not come up on EX4300 non-AE interfaces. [PR1538401](#)
- Redirected IP traffic is duplicated. [PR1518929](#)

### ***Routing Protocols***

- On EX4300-MP and EX4600, high CPU load occurs due to receipt of specific Layer 2 frames in EVPN-VXLAN deployment. (CVE-2020-1687) & High CPU load occurs due to receipt of specific Layer 2 frames when deployed in a Virtual Chassis configuration (CVE-2020-1689). [PR1495890](#)
- The rpd might report 100 percent CPU usage with BGP route damping enabled. [PR1514635](#)
- Packet loss might be observed while verifying traffic from access to core network for IPv4/IPv6 interfaces. [PR1520059](#)
- OSPFv3 adjacency should not be established when IPsec authentication is enabled. [PR1525870](#)

### ***User Interface and Configuration***

- Installing J-Web application package might fail on the EX2300/EX3400 platforms. [PR1513612](#)
- The J-Web does not display the correct flow-control status on EX Series devices. [PR1520246](#)

### ***Virtual Chassis***

- EX4650: "kldload: an error occurred while loading the module" during booting. [PR1527170](#)

## **Resolved Issues: 20.2R1**

### ***Authentication and Access Control***

- EX2300-48MP: Client did not receive captive-portal success page by downloading the ACL parameter as Authentication failed. [PR1504818](#)

### ***EVPN***

- The ESI of IRB interfaces does not get updated after an autonomous-system number change if the interface is down. [PR1482790](#)
- The VXLAN function might be broken due to a timing issue after the change in PR 1495098. [PR1502357](#)

### ***Infrastructure***

- Kernel core files might be observed if you deactivate the daemon on EX2300/EX3400 platforms. [PR1483644](#)

### ***Interfaces and Chassis***

- FRU has no connection arguments `fru_send_msg Global FPC x` is observed after MX Series Virtual Chassis local or global switchover. [PR1428254](#)
- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. [PR1459201](#)

- Executing commit might hang up due to a stuck dcd process. [PR1470622](#)
- A stale IP address might be seen after a specific order of configuration changes under a logical-systems scenario. [PR1477084](#)

### ***Junos Fusion for Enterprise***

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

### ***Junos Fusion Satellite Software***

- Temperature sensor alarm is seen on EX4300 in a Junos fusion scenario. [PR1466324](#)

### ***Layer 2 Ethernet Services***

- Member links state might be asynchronized on a connection between PE and CE devices in an EVPN active/active scenario. [PR1463791](#)
- Issues with DHCPv6 relay processing Confirm and Reply packets. [PR1496220](#)

### ***Layer 2 Features***

- The LLDP function might fail when a Juniper device connects to a non-Juniper one. [PR1462171](#)
- EX4650/QFX5120: QinQ: The third VLAN tag is not pushed onto the stack and SWAP is being done instead. [PR1469149](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)

### ***MPLS***

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

### ***Platform and Infrastructure***

- The IRB traffic might get dropped after mastership switchover. [PR1453025](#)
- The switch might not be able to learn MAC addresses with `dot1x` and `interface-mac-limit` configured. [PR1470424](#)
- EX4300: Input firewall filter attached to isolated or community VLANs not matching 802.1p bits on the VLAN header. [PR1478240](#)
- MAC learning under bridge-domain stops after an MC-LAG interface flap. [PR1488251](#)
- The NSSU upgrade might fail on EX4300 switches due to a storage issue in the `/var/tmp` directory. [PR1494963](#)
- Traffic loss might be seen with framing errors or runts if MACsec is configured on EX4300. [PR1502726](#)
- The MAC Pause frames will be incrementing in the Receive direction if half-duplex mode on 10-Mbps or 100-Mbps speed is configured. [PR1452209](#)

- Link up delay and traffic drop might be seen on mixed SP L2/L3 and EP L2 type configurations. [PR1456336](#)
- MAC addresses learned on RTG may not be aged out after the aging time. [PR1461293](#)
- RTG link faces nearly 20 seconds down during backup node rebooting. [PR1461554](#)
- The jdhcpd process might consume high CPU and no further subscribers can be brought up if there are more than 4000 DHCP relay clients in the MAC move scenario. [PR1465277](#)
- FPCs might get disconnected from the EX3400 Virtual Chassis briefly after a reboot or an upgrade. [PR1467707](#)
- Traffic loss might be seen with framing errors or runts if MACsec is configured on EX4600 or QFX5100 platforms. [PR1469663](#)
- SSH session closes while checking for the **show configuration | display set** command for both local and nonlocal users. [PR1470695](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- CoS 802.1p bits rewrite might not happen in Q-in-Q mode. [PR1472350](#)
- DSCP marking might not work as expected if the fixed classifiers are applied to interfaces on QFX5000 or EX4600 platforms. [PR1472771](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- The RIPv2 packets forwarded across a Layer 2 circuit connection might be dropped. [PR1473685](#)
- On EX4300, the output of **show security macsec statistics** shows high values incorrectly. [PR1476719](#)
- EX3400 me0 interface might remain down. [PR1477165](#)
- The dhcpd process may crash in a Junos fusion environment. [PR1478375](#)
- Trio based linecard might crash when there is bulk route update failure in a corner case. [PR1478392](#)
- TFTP installation from loader prompt may not succeed on the EX Series devices. [PR1480348](#)
- ARP request packets for an unknown host might get dropped in remote PE in EVPN-VXLAN scenario. [PR1480776](#)
- On EX2300 switches, SNMP traps are not generated when the MAC addresses limit threshold is reached. [PR1482709](#)
- Incorrect 'frame length' of 132 bytes might be shown in packet header. [PR1487876](#)
- Virtual Chassis ports might go down in a mixed Virtual Chassis setup of QFX5100-24Q-2P/EX4300 and EX4600/EX4300. [PR1489985](#)
- DHCP binding fails while you verify DHCPv4 snooping functionality with P-VLAN with a firewall to block or allow certain IPv4 packets. [PR1490689](#)
- Traffic loss could be observed in a mixed-Virtual Chassis setup of QFX5100 and EX4300. [PR1493258](#)

- Traffic loss could be seen in an MC-LAG scenario on QFX5120 and EX4650. [PR1494507](#)
- Traffic might get dropped if AE member interface is deleted/added or a SFP of the AE member interface is unplugged/plugged. [PR1497993](#)

### **Routing Protocols**

- BGP IPv4/IPv6 convergence and RIB install and delete time is degraded in Junos OS Releases 19.1R1, 19.2R1, 19.3R1, and 19.4R1. [PR1414121](#)
- MUX State in LACP interface does not go to **collecting and distributing** and remains **attached** after enabling the ae interface. [PR1484523](#)
- FPC might go to "NotPrsnt" state after upgrading with non-TVP image in VC/VCF setup. [PR1485612](#)
- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- Firewall filter could not work in certain conditions in an Virtual Chassis setup. [PR1497133](#)

### **User Interface and Configuration**

- **umount: unmount of /.mount/var/val/chroot/packages/mnt/jweb-ex32-d2cf6f6b failed: Device busy** message is seen when Junos OS is upgraded with the validate option. [PR1478291](#)

### SEE ALSO

[What's New | 45](#)

[What's Changed | 54](#)

[Known Limitations | 58](#)

[Open Issues | 59](#)

[Documentation Updates | 70](#)

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

## **Documentation Updates**

There are no errata or changes in Junos OS Release 20.2R2 documentation for EX Series switches.

### SEE ALSO

[What's New | 45](#)

---

[What's Changed | 54](#)

---

[Known Limitations | 58](#)

---

[Open Issues | 59](#)

---

[Resolved Issues | 63](#)

---

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

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 71](#)

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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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

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



## SEE ALSO

---

[What's New | 45](#)

---

[What's Changed | 54](#)

---

[Known Limitations | 58](#)

---

[Open Issues | 59](#)

---

[Resolved Issues | 63](#)

---

[Documentation Updates | 70](#)

## Junos OS Release Notes for JRR Series

### IN THIS SECTION

- [What's New | 73](#)
- [What's Changed | 74](#)
- [Known Limitations | 74](#)
- [Open Issues | 74](#)
- [Resolved Issues | 75](#)
- [Documentation Updates | 76](#)
- [Migration, Upgrade, and Downgrade Instructions | 77](#)

These release notes accompany Junos OS Release 20.2R3 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

- [What's New in Release 20.2R3 | 73](#)
- [What's New in Release 20.2R2 | 73](#)
- [What's New in Release 20.2R1 | 73](#)

Learn about what changed in Junos OS main and maintenance releases for JRR Series Route Reflectors.

### What's New in Release 20.2R3

There are no new features or enhancements to existing features for JRR Series in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features or enhancements to existing features for JRR Series in Junos OS Release 20.2R2.

### What's New in Release 20.2R1

#### *Layer 2 Features*

- **Support for Link Layer Discovery Protocol (JRR200)**—Starting in Junos OS Release 20.2R1, JRR Series devices support Link Layer Discovery Protocol (LLDP) is supported both on the management port em0 and on the WAN ports em2 through em9. LLDP is a link-layer protocol defined in IEEE 802.1AB that allows network devices to advertise their identity, capabilities, and configuration to other devices on the LAN.

[See [Understanding LLDP.](#)]

### SEE ALSO

---

[What's Changed | 74](#)

---

[Known Limitations | 74](#)

---

[Open Issues | 74](#)

---

[Resolved Issues | 75](#)

---

[Documentation Updates | 76](#)

---

[Migration, Upgrade, and Downgrade Instructions | 77](#)

## What's Changed

There are no changes in behavior and syntax in Junos OS Release 20.2R3 for JRR Series Route Reflectors.

SEE ALSO

[What's New | 73](#)

[Known Limitations | 74](#)

[Open Issues | 74](#)

[Resolved Issues | 75](#)

[Documentation Updates | 76](#)

[Migration, Upgrade, and Downgrade Instructions | 77](#)

## Known Limitations

There are no known limitations in Junos OS Release 20.2R3 for JRR Series Route Reflectors.

SEE ALSO

[What's New | 73](#)

[What's Changed | 74](#)

[Open Issues | 74](#)

[Resolved Issues | 75](#)

[Documentation Updates | 76](#)

[Migration, Upgrade, and Downgrade Instructions | 77](#)

## Open Issues

There are no open issues in Junos OS Release 20.2R3 for JRR Series Route Reflectors.

## SEE ALSO

[What's New | 73](#)[What's Changed | 74](#)[Known Limitations | 74](#)[Resolved Issues | 75](#)[Documentation Updates | 76](#)[Migration, Upgrade, and Downgrade Instructions | 77](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 76](#)
- [Resolved Issues: 20.2R2 | 76](#)
- [Resolved Issues: 20.2R1 | 76](#)

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

## Resolved Issues: 20.2R3

### General Routing

- On the JRR200 device, four out of eight fans might not work after upgrading to Junos OS Release 19.4R1 and later. This might cause high temperature of the device eventually impacting the traffic. [PR1534706](#)
- On the JRR200 device, four out of eight fans might not work after upgrading to Junos OS Release 19.4R1 and later. This might cause high temperature of the device eventually impacting the traffic. [PR1534706](#)

## Resolved Issues: 20.2R2

### General Routing

- On the JRR200 routers, the firewall filter with non-zero TTL value might cause a commit error. [PR1531034](#)

## Resolved Issues: 20.2R1

### General Routing

- USB install image is not working for JRR200 platform. [PR1471986](#)
- Link state of virtual em interfaces in Junos OS might not reflect the true link status of corresponding physical interfaces in the Linux host. [PR1492087](#)

## SEE ALSO

[What's New | 73](#)

[What's Changed | 74](#)

[Known Limitations | 74](#)

[Open Issues | 74](#)

[Documentation Updates | 76](#)

[Migration, Upgrade, and Downgrade Instructions | 77](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 documentation for JRR200 Route Reflectors.

## SEE ALSO

---

[What's New | 73](#)

---

[What's Changed | 74](#)

---

[Known Limitations | 74](#)

---

[Open Issues | 74](#)

---

[Resolved Issues | 75](#)

---

[Migration, Upgrade, and Downgrade Instructions | 77](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 77](#)

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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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 EOL releases and to review a list of EOL releases, see <https://www.juniper.net/support/eol/junos.html>.

#### SEE ALSO

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

---

- [What's Changed | 74](#)

---

- [Known Limitations | 74](#)

---

- [Open Issues | 74](#)

---

- [Resolved Issues | 75](#)

---

- [Documentation Updates | 76](#)

## Junos OS Release Notes for Junos Fusion for Enterprise

#### IN THIS SECTION

- [What's New | 79](#)
- [What's Changed | 79](#)
- [Known Limitations | 80](#)
- [Open Issues | 80](#)
- [Resolved Issues | 81](#)
- [Documentation Updates | 82](#)
- [Migration, Upgrade, and Downgrade Instructions | 82](#)

These release notes accompany Junos OS Release 20.2R3 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 20.2R3 for Junos fusion for enterprise.

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

### SEE ALSO

---

[What's Changed | 79](#)

---

[Known Limitations | 80](#)

---

[Open Issues | 80](#)

---

[Resolved Issues | 81](#)

---

[Documentation Updates | 82](#)

---

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

## 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 20.2R3 for Junos fusion for enterprise.

### SEE ALSO

---

[What's New | 79](#)

---

[Known Limitations | 80](#)

---

[Open Issues | 80](#)

---

[Resolved Issues | 81](#)

---

[Documentation Updates | 82](#)

---

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



## Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.2R3 for Junos fusion for enterprise.

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

### SEE ALSO

---

[What's New | 79](#)

---

[What's Changed | 79](#)

---

[Open Issues | 80](#)

---

[Resolved Issues | 81](#)

---

[Documentation Updates | 82](#)

---

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

## Open Issues

There are no known issues in hardware and software in Junos OS Release for 20.2R3 Junos fusion for enterprise.

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

### SEE ALSO

---

[What's New | 79](#)

---

[What's Changed | 79](#)

---

[Known Limitations | 80](#)

---

[Resolved Issues | 81](#)

---

[Documentation Updates | 82](#)

---

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

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: Release 20.2R3 | 81](#)
- [Resolved Issues: Release 20.2R2 | 81](#)
- [Resolved Issues: Release 20.2R1 | 81](#)

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

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

### Resolved Issues: Release 20.2R3

There are no resolved issues in Junos OS Release 20.2R3 for Junos fusion for enterprise.

### Resolved Issues: Release 20.2R2

- The 100M SFP-FX is not supported as a satellite device in a Junos fusion setup. [PR1514146](#)

### Resolved Issues: Release 20.2R1

- Observing duplicate ECID values for cluster and extended ports on member ports of same cluster. [PR1408947](#)
- The SDPD process generates a core file at `vfpc_all_eports_deletion_complete` `vfpc_dampen_fpc_timer_expiry`. [PR1454335](#)
- Loop detection might not work on extended ports in a Junos fusion scenario. [PR1460209](#)
- The temperature sensor alarm is seen on EX4300 in a Junos fusion scenario. [PR1466324](#)

### SEE ALSO

[What's New | 79](#)

[What's Changed | 79](#)

---

[Known Limitations | 80](#)

---

[Open Issues | 80](#)

---

[Documentation Updates | 82](#)

---

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

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 for documentation for Junos fusion for enterprise.

### SEE ALSO

---

[What's New | 79](#)

---

[What's Changed | 79](#)

---

[Known Limitations | 80](#)

---

[Open Issues | 80](#)

---

[Resolved Issues | 81](#)

---

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

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading Junos OS on an Aggregation Device | 83](#)
- [Upgrading an Aggregation Device with Redundant Routing Engines | 84](#)
- [Preparing the Switch for Satellite Device Conversion | 85](#)
- [Converting a Satellite Device to a Standalone Switch | 86](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 86](#)
- [Downgrading Junos OS | 87](#)

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 Software Installation and Upgrade Guide](#).

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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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

## SEE ALSO

---

[What's New | 79](#)

---

[What's Changed | 79](#)

---

[Known Limitations | 80](#)

---

[Open Issues | 80](#)

---

[Resolved Issues | 81](#)

---

[Documentation Updates | 82](#)



# Junos OS Release Notes for Junos Fusion for Provider Edge

## IN THIS SECTION

- [What's New | 88](#)
- [What's Changed | 90](#)
- [Known Limitations | 90](#)
- [Open Issues | 91](#)
- [Resolved Issues | 91](#)
- [Documentation Updates | 92](#)
- [Migration, Upgrade, and Downgrade Instructions | 93](#)

These release notes accompany Junos OS Release 20.2R3 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

### IN THIS SECTION

- [What's New in Release 20.2R3 | 89](#)
- [What's New in Release 20.2R2 | 89](#)
- [What's New in Release 20.2R1 | 89](#)

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

## What's New in Release 20.2R3

There are no new features or enhancements to existing features for Junos fusion for provider edge in Junos OS Release 20.2R3.

## What's New in Release 20.2R2

There are no new features or enhancements to existing features for Junos fusion for provider edge in Junos OS Release 20.2R2.

## What's New in Release 20.2R1

### Hardware

- **Support for QFX5110 as a satellite device in a Junos fusion for provider edge on a GNF(MX480 and MX960)**—With Junos Node Slicing, you can create guest network functions (GNFs), partitions where an aggregation device can be configured. The aggregation device on a GNF supports a maximum of 10 satellite devices. Starting in Junos OS Release 20.2R1, Junos OS supports QFX5110 switches as satellite devices in Junos fusion for provider edge on a GNF.

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

### Junos Fusion

- **MPC10E and MPC11E interoperability with Junos fusion for provider edge (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 20.2R1, Junos OS supports using the MPC10E and MPC11E alongside other MPC line cards in the same MX Series router chassis that has been configured with Junos fusion for provider edge. The line cards can coexist in the same router chassis, and the router passes traffic between the devices connected to the MPC10E/MPC11E and the satellite devices that are connected to other MPC line cards through the switch fabric. You cannot use MPC10E/MPC11E in Junos fusion, which means you cannot connect satellite devices to ports on the MPC10E/MPC11E line cards.

Junos fusion does not support hyper mode. To support Junos fusion in an MX Series router where MPC10E/MPC11E coexists with other MPC line cards, use the **set forwarding-options no-hyper-mode** statement. In addition, you must also use an FPC slot ID in the range of 160–252 for the satellite device interfaces. To configure the FPC slot ID, use the **set chassis satellite-management fpc slot-id** statement.

[See [Junos Fusion Provider Edge Overview](#).]

SEE ALSO

| [What's Changed](#) | 90

---

[Known Limitations | 90](#)

---

[Open Issues | 91](#)

---

[Resolved Issues | 91](#)

---

[Documentation Updates | 92](#)

---

[Migration, Upgrade, and Downgrade Instructions | 93](#)

---

## What's Changed

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

### SEE ALSO

---

[What's New | 88](#)

---

[Known Limitations | 90](#)

---

[Open Issues | 91](#)

---

[Resolved Issues | 91](#)

---

[Documentation Updates | 92](#)

---

[Migration, Upgrade, and Downgrade Instructions | 93](#)

---

## Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.2R3 for Junos fusion for provider edge.

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

### SEE ALSO

---

[What's New | 88](#)

---

[What's Changed | 90](#)

---

[Open Issues | 91](#)

---

[Resolved Issues | 91](#)

---

[Documentation Updates | 92](#)

[Migration, Upgrade, and Downgrade Instructions | 93](#)

## Open Issues

There are no known issues in the Junos OS Release 20.2R3 for Junos fusion for provider edge.

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

### SEE ALSO

[What's New | 88](#)

[What's Changed | 90](#)

[Known Limitations | 90](#)

[Resolved Issues | 91](#)

[Documentation Updates | 92](#)

[Migration, Upgrade, and Downgrade Instructions | 93](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 92](#)
- [Resolved Issues: 20.2R2 | 92](#)
- [Resolved Issues: 20.2R1 | 92](#)

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

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

## Resolved Issues: 20.2R3

There are no fixed issues for Junos OS Release 20.2R3.

## Resolved Issues: 20.2R2

### *Junos Fusion for Provider Edge*

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

## Resolved Issues: 20.2R1

### *Junos Fusion for Provider Edge*

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

SEE ALSO

---

[What's New | 88](#)

---

[What's Changed | 90](#)

---

[Known Limitations | 90](#)

---

[Open Issues | 91](#)

---

[Documentation Updates | 92](#)

---

[Migration, Upgrade, and Downgrade Instructions | 93](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 documentation for Junos fusion for provider edge.

SEE ALSO

---

[What's New | 88](#)

---

[What's Changed | 90](#)

---

[Known Limitations | 90](#)

Open Issues | 91

Resolved Issues | 91

Migration, Upgrade, and Downgrade Instructions | 93

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 93
- Upgrading an Aggregation Device with Redundant Routing Engines | 96
- Preparing the Switch for Satellite Device Conversion | 96
- Converting a Satellite Device to a Standalone Device | 98
- Upgrading an Aggregation Device | 100
- Upgrade and Downgrade Support Policy for Junos OS Releases | 101
- Downgrading from Junos OS Release 20.1 | 101

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 **bundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#).

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

Issue the following command:

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

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

The download and installation process for Junos OS Release 20.2R1 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/jinstall64-20.2R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.2R3.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

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

```
user@host> request system software add validate reboot
source/jinstall64-20.2R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.2R3.SPIN-export-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.



- For software packages that are downloaded and installed from a remote location:
  - `ftp://hostname/pathname`
  - `http://hostname/pathname`
  - `scp://hostname/pathname` (available only for the Canada and U.S. version)

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

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

Rebooting occurs only if the upgrade is successful.

**NOTE:** After you install a Junos OS Release 20.2R1 **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.

- (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 20.2R3, 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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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

---

[What's New | 88](#)

---

[What's Changed | 90](#)

---

[Known Limitations | 90](#)

---

[Open Issues | 91](#)

---

[Resolved Issues | 91](#)

---

[Documentation Updates | 92](#)

# Junos OS Release Notes for MX Series

## IN THIS SECTION

- [What's New | 102](#)
- [What's Changed | 129](#)
- [Known Limitations | 136](#)
- [Open Issues | 139](#)
- [Resolved Issues | 149](#)
- [Documentation Updates | 185](#)
- [Migration, Upgrade, and Downgrade Instructions | 186](#)

These release notes accompany Junos OS Release 20.2R3 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

## IN THIS SECTION

- [What's New in Release 20.2R3 | 103](#)
- [What's New in Release 20.2R2-S3 | 103](#)
- [What's New in Release 20.2R2-S2 | 103](#)
- [What's New in Release 20.2R2 | 103](#)
- [What's New in Release 20.2R1-S1 | 104](#)
- [What's New in Release 20.2R1 | 104](#)

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

## What's New in Release 20.2R3

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 20.2R3.

## What's New in Release 20.2R2-S3

### OAM

- **Inline CCM Support for MPC10E (MX Series)**—Starting in Junos OS Release 20.2R2S3, Junos OS extends support for inline continuity check messages (CCM) on the MPC10E (MPC10E-10C-MRATE and MPC10E-15C-MRATE) line cards. You can configure inline CCM for both UP MEP and Down MEP to monitor services provided by currently deployed topologies such as INET, CCC/VPWS, Bridge, VPLS, EVPN, and others. Junos OS extends MIP support for all current supported topologies.

[See [Inline Transmission Mode](#).]

## What's New in Release 20.2R2-S2

### Services Applications

- **AMS support (MX240, MX480, MX960, MX2010, and MX2020 routers)**—In Release 20.2R2S2, Junos OS supports AMS (Aggregated Multiservices Interfaces) on the MPC10E and MX2K-MPC11E line cards to provide load balancing (LB) and high availability (HA) features for stateful firewall and NAT services. You can configure AMS with next-hop style service-sets and with MS-MPC only.

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

## What's New in Release 20.2R2

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 20.2R2.



## What's New in Release 20.2R1-S1

### *Software Installation and Upgrade*

- **Zero touch provisioning (ZTP) with IPv6 support (EX3400, EX4300, QFX5100 and QFX5200 switches, MX-Series routers)**—Starting in Junos OS Release 20.2R1-S1, ZTP supports the DHCPv6 client. During the bootstrap process, the device first uses the DHCPv4 client to request for information regarding image and configuration file from the DHCP server. The device checks the DHCPv4 bindings sequentially. If there is a failure with one of the DHCPv4 bindings, the device will continue to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device will check for DHCPv6 bindings and follow the same process as for DHCPv4 until the device can be provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

**NOTE:** Only HTTP and HTTPS transport protocols are supported EX3400, EX4300, QFX5100, and QFX5200 devices.

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

## What's New in Release 20.2R1

### *Class of Service (CoS)*

- **Support for rewrite rules on a per-customer basis on MPC10 and MPC11 (MX Series)**—Starting in Junos OS Release 20.2R1, we support creating rewrite rules on a per-customer basis on MPC10 and MPC11 cards. You can create rewrite rules on a per-customer basis through a policy map. You define policy maps at the **[edit class-of-service policy-map]** hierarchy level, and assign the policy map to a customer through a firewall action, an ingress interface, or a routing policy.

[See [Assigning Rewrite Rules on a Per-Customer Basis Using Policy Maps Overview](#).]

### *EVPN*

- **IPv4 unicast VXLAN encapsulation optimization (MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX10016)**—Starting in Junos OS Release 20.2R1, by default, the listed MX Series routers optimize the IPv4 unicast VXLAN encapsulation process for the following tunnel types:
  - PIM-based VXLAN
  - EVPN-VXLAN
  - Static VXLAN

The optimized encapsulation process results in an increased throughput rate for IPv4 unicast packets between 512 to 1500 bytes in size.

The optimization feature does not support the following:

- EVPN Type-5 tunnels, which are already optimized
- Forwarding table filters

[See [Understanding VXLANs.](#)]

- **EVPN on MPLS-over-UDP tunnels (MX Series and vMX)**—Starting in Junos OS Release 20.2R1, Junos OS supports an EVPN network with MPLS-over-UDP tunnels. EVPN uses indirect next hop while MPLS-over-UDP tunnels use tunnel composite next hop (TCNH) in resolving routes in the routing table. In Junos OS releases before Release 20.2R1, indirect next hops for EVPN traffic on MPLS-over-UDP tunnels resolve into unicast next hops. With this release, the indirect next hops for EVPN traffic on MPLS-over-UDP tunnels will resolve into TCNH.

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

- **Support for inline performance monitoring services on EVPN (MX Series)**—Starting in Junos OS Release 20.2R1, you can enable inline performance monitoring services on an EVPN network. With inline performance monitoring, you can configure a greater number of performance monitoring sessions. Inline performance monitoring applies only to delay measurements and synthetic loss measurements. You must also enable both enhanced IP network services and enhanced CFM mode in the device.

To enable inline performance monitoring, include the following statements:

- **hardware-assisted-pm** and **hardware-assisted-keepalives enable** statements at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring]** hierarchy level.
- **enhanced-ip** statement at the **[edit chassis network-services]** hierarchy level.
- **enhanced-cfm-mode** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

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

- **Noncolored SR-TE LSPs with EVPN-MPLS (ACX5448, EX9200, MX Series, and vMX)**—Starting in Junos OS Release 20.2R1, ACX5448, EX9200, MX Series, and vMX routers support noncolored static segment routing-traffic engineered (SR-TE) label-switched paths (LSPs) with an EVPN-MPLS core network and the following Layer 2 services running at the edges of the network:

- E-LAN
- EVPN-ETREE
- EVPN-VPWS with E-Line

Without color, all LSPs resolve using a BGP next hop only.

The Juniper Networks routers support noncolored SR-TE LSPs in an EVPN-MPLS core network with the following configurations:

- EVPN running in a virtual switch routing instance
- Multihoming in active/active and active/standby modes

The Juniper Networks routers also support noncolored SR-TE LSPs when functioning as a Data Center Interconnect (DCI) device that handles EVPN Type 5 routes.

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

- **Layer 3 gateway in an EVPN-MPLS environment (MPC10 and MPC11 line cards with MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, the supported MX Series routers with MPC10 and MPC11 line cards can act as a default Layer 3 gateway for an EVPN instance (EVI), which can span a set of routers. In this role, the MX Series routers can perform inter-subnet forwarding. With inter-subnet forwarding, each subnet represents a distinct broadcast domain.

The Layer 3 gateway supports the following features:

- IRB interfaces through which the default gateway routes IPv4 and IPv6 traffic from one bridge domain to another [See [Example: Configuring EVPN with IRB Solution](#).]
- Dynamic list next hop [See [Configuring Dynamic List Next Hop](#).]
- EVPN proxy ARP and ARP suppression, and proxy NDP and NDP suppression on IRB interfaces [See [EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression](#).]
- The substitution of a source MAC address with a proxy MAC address in an ARP or NDP reply [See [ARP and NDP Request with a Proxy MAC Address](#).]
- Data center interconnectivity using EVPN Type 5 routes [See [EVPN Type-5 Route with MPLS encapsulation for EVPN-MPLS](#).]
- **Multihoming in an EVPN-MPLS environment (MPC10 and MPC11 line cards with MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, you can multihome a customer edge (CE) device to two or more provider edge (PE) devices (the supported MX Series routers with MPC10 and MPC11 line cards) in an EVPN-MPLS network. We support the following multihoming features:
  - Single-active and all-active modes
  - The configuration of an Ethernet segment identifier (ESI) per interface
  - Preference-based designated forwarder election

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

- **EVPN-VXLAN (MPC10 and MPC11 line cards with MX2010, MX2020)**—Starting in Junos OS Release 20.2R1, the MX2010 and MX2020 routers with MPC10 and MPC11 line cards installed support the following EVPN-VXLAN features:
  - Layer 2 VXLAN
    - Multihoming with active/active and active/standby modes, an Ethernet segment identifier (ESI) per interface, and preference-based designated forwarder (DF) election
    - MAC pinning, MAC move, MAC limiting, and MAC aging
    - QoS

- DHCP and DHCP relay
- Prevention of broadcast, unknown unicast, and multicast (BUM) traffic loops when a leaf device is multihomed to more than one spine device
- Layer 3 VXLAN
  - IRB interfaces
  - IPv6 over IRB interfaces
  - Support for OSPF, IS-IS, BGP, and static routing over IRB interfaces
  - Proxy ARP and ARP suppression, and proxy NDP and NDP suppression with and without IRB interfaces
  - IPv6 underlay
  - Virtual machine traffic optimization (VMTO) for ingress traffic
- Data Center Interconnect (DCI)
  - Nonpure and pure EVPN Type-5 routes
- High availability
  - Nonstop active routing (NSR)
  - Graceful Routing Engine switchover (GRES)
  - Graceful restart from a routing process restart or Routing Engine switchover without NSR enabled
- Operations and management
  - Core isolation feature
  - Ping over EVPN Type-5 tunnel
- Static VXLAN
  - Overlay ping and traceroute

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

### ***High Availability (HA) and Resiliency***

- **Support for VRRP on the MPC11 (MX2010 and MX2020)**—Starting in Junos OS Release 20.2R1, VRRP is supported on the MPC11 line card. All VRRP features are supported.

[See [Understanding VRRP](#).]

- **LACP inline support during unified ISSU for multivendor networks (MX104, MX240, MX480, MX960, and MX10003)**—Starting with Junos OS Release 20.2R1, unified ISSU supports LACP interoperability with other vendor devices for fast periodic interval sessions. LACP sessions in full-scale scenarios with interoperability will no longer experience timeouts during unified ISSU.

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

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

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the `commit synchronize` statement at the `[edit system]` hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes roles. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

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

- **Support for VRRP on the MPC10 and MPC11 (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, VRRP is supported on the MPC11 and MPC10 line cards. All VRRP features are supported.

[See [Understanding VRRP.](#)]

- **Unsupported hardware for unified ISSU (MX240, MX480, MX960, MX10003, and PTX3000)**—The following cards do not support unified ISSU upgrading to Junos OS Release 20.2R1:

- MPC7E-MRATE
- MPC8E with MRATE MIC
- MPC9E with MRATE MIC
- MPC10E-10C-MRATE
- MPC10E-15C-MRATE
- PTX5000 with 24-Port 10-Gigabit Ethernet, 40-Gigabit Ethernet PIC with QSFP+ or 15-Port 10-Gigabit, 40-Gigabit Ethernet, 100-Gigabit Ethernet PIC with QSFP28
- MX10003 with QSFP28 Ethernet TIC

### ***Interfaces and Chassis***

- **Transparent forwarding of CFM packets over VPLS (MX Series)**—In Junos OS Release 20.2R1 and later, MX Series router supports VLAN transparency for connectivity fault management (CFM) packets over Virtual private LAN service (VPLS). If the incoming CFM packets have more **vlan-tags** than the configured interface **vlan-tags**, then CFM PDU is treated transparent. In the earlier Junos OS releases, CFM frame filtering was applied on all CFM PDU including on CFM PDU that had more number of tags than the interface configuration.

We do not support the following on MX Series routers:

- Transparency for tagged CFM PDU incoming on untagged interface.

- Transparency for untagged CFM PDU on interface with native VLAN configuration.

[See [Example: Configuring Ethernet CFM over VPLS.](#)]

- **Support for 400-Gbps port speed (MX240, MX480, and MX960)**—In Junos OS Release 20.2R1, you can configure port speed of 400-Gbps for MPC10E (MPC10E-10C-MRATE and MPC10E-15C-MRATE) on MX240, MX480, and MX960 routers. Use the QSFP56-DD optics to configure 400-Gbps port speed on:
  - MPC10E-10C-MRATE: Port 4 of the MPC
  - MPC10E-15C-MRATE: Port 4 of the MPC

[See [Port Speed.](#)]

- **Support for monitoring link degradation (MX Series routers with MPC10E)**—Starting in Junos OS Release 20.2R1, you can monitor link degradation of the 10-Gigabit Ethernet interfaces, 40-Gigabit Ethernet interfaces, and 100-Gigabit Ethernet interfaces on the MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line cards. Link degradation monitoring enables you to monitor the quality of physical links on interfaces and take corrective action when the link quality degrades beyond a certain value.

To enable your device to monitor the links, use the **link-degrade-monitor** statement at the [**edit interfaces interface-name**] hierarchy level.

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

- **Targeted broadcast support (MPC10E and MX2K-MPC11E)**—Starting in Junos OS Release 20.2R1, you can configure targeted broadcast on broadcast interfaces on the MPC10E and MX2K-MPC11E line cards. Targeted broadcast enables a broadcast packet, destined for a remote network, to transit across networks until the destination network is reached. In the destination network, the packet is broadcast as a normal broadcast packet. This feature is useful when the Routing Engine is flooded with packets to process. You can configure targeted broadcast to forward the packets to :
  - Both the egress interface and the Routing Engine.
  - Egress interface only.

To configure targeted broadcast on an interface, include the **targeted-broadcast** statement at the [**edit interfaces interface-name unit logical-unit-number family inet**] hierarchy level.

[See [Understanding Targeted Broadcast.](#)]

### *Juniper Extension Toolkit (JET)*

- **RIB service APIs support dynamic next-hop interface binding (MX Series, PTX Series, and vMX)**—Starting in Junos OS Release 20.2R1, programmed RIB routes react to Up, Down, Add, and Delete events for direct next-hop interfaces. When all direct next-hop interfaces are unusable, the route becomes inactive. This prevents traffic from being dropped and keeps inactive routes from being propagated through the network.

This feature applies to all routes programmed using the `rib_service` JET API where an interface is configured as a direct next hop, including interfaces that are part of a flexible tunnel. It also applies to tunnels configured with the `flexible_tunnel_service` JET API.

To disable this feature, use **`edit routing-options programmable-rpd rib-service dynamic-next-hop-interface disable`**.

[See [rib-service \(programmable-rpd\)](#), [Juniper Extension Toolkit Developer Guide](#), and [Juniper Engineering Network website](#).]

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **`set system scripts language python3`** command.

[See [language \(Scripts\)](#), [Develop Off-Device JET Applications](#), and [Develop On-Device JET Applications](#).]

### *Junos Telemetry Interface*

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models `openconfig-local-routing.yang` and `openconfig-network-instance.yang`.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

- **ON-CHANGE BGP peer information statistics support for JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 provides BGP peer sensor support using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON\_CHANGE statistics are sent to an outside collector.

The following resource paths are supported:

- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/active` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/received` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/sent` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/rejected` (ON\_CHANGE)

- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/admin-state (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/established-transitions (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/last-established (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/received/notification (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/messages/received/update (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/notification (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/update (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/session-state (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/supported-capabilities (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/transport/state/local-address (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-address (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-port (ON\_CHANGE)

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

- **Telemetry support for LDP and MLDP traffic statistics (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, the following LDP and multipoint LDP native sensors are added for the Junos telemetry interface:
  - /junos/services/ldp/label-switched-path/ingress/usage/
  - /junos/services/ldp/label-switched-path/transit/usage/
  - /junos/services/ldp/p2mp/interface/receive/usage/
  - /junos/services/ldp/p2mp/interface/transmit/usage/
  - /junos/services/ldp/p2mp/label-switched-path/usage/

You must enable telemetry streaming with the **sensor-based-stats** option at the **[edit protocols ldp traffic-statistics]** hierarchy level.



The **show ldp traffic-statistics** command is enhanced to display upstream LDP traffic statistics and to display multipoint LDP traffic statistics per interface.

On PTX Series routers, this feature is not supported for the following variants:

- PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine
- PTX10003
- PTX10008 with the PTX10K-LC1201-36CD line card
- FPC2 line cards do not support ingress multipoint LDP statistics.

[See [sensor \(Junos Telemetry Interface\)](#).]

- **gRPC telemetry support for LDP and MLDP traffic statistics (MX Series)**—Starting in Junos OS Release 20.2R1, gRPC support is available to export LDP and multipoint LDP traffic statistics. You can use the following resource paths to export sensor data:

- LDP LSP transit  
traffic—/mpls/signaling-protocols/ldp/lsp-transit-policies/lsp-transit-policy/state/counters
- LDP LSP ingress  
traffic—/mpls/signaling-protocols/ldp/lsp-ingress-policies/lsp-ingress-policy/state/counters
- Multipoint LDP traffic—/mpls/signaling-protocols/ldp/p2mp-lsps/p2mp-lsp/state/counters
- Multipoint LDP egress traffic  
per-interface—/mpls/signalling-protocols/ldp/p2mp-interfaces/p2mp-interface/state/counters
- Multipoint LDP ingress traffic  
per-interface—/mpls/signalling-protocols/ldp/p2mp-interfaces/p2mp-interface/

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

- **JTI sensor support for Packet Forwarding Engine and Routing Engine sensors (MX Series Virtual Chassis and MX Series routers with dual Routing Engines)**—Junos OS Release 20.2R1 extends Junos telemetry interface (JTI) sensor support for all Packet Forwarding Engine and Routing Engine sensors currently supported on MX Series routers to include MX routers with dual Routing Engines or MX Series Virtual Chassis. The level of sensor support currently available for MX Series routers applies, whether through streaming or ON\_CHANGE statistics export, using UDP, remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. Additionally, JTI operational mode commands will provide details for all Routing Engines and MX Series Virtual Chassis, too.

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

- **JTI sensor support for standby Routing Engine statistics (MX480, MX960, MX10003, MX2010, and MX2020)**—Junos OS Release 20.2R1 provides Junos telemetry interface (JTI) sensor support for standby Routing Engine statistics using remote procedure call (gRPC) services. This feature is supported on both single chassis and virtual chassis unless otherwise indicated. Use this feature to better track the state of software components running on a standby Routing Engine. Statistics exported to an outside collector

through the following sensors (primarily under subscriber management) provide a more complete view of the system health and resiliency state:

- Chassis role (backup or master) sensor `/junos/system/subscriber-management/chassis` and `/junos/system/subscriber-management/chassis[chassis-index=chassis-index]` (for specifying an index for an MX Series Virtual Chassis)
- Routing Engine status and GRES notification sensor `/junos/system/subscriber-management/chassis/routing-engines/routing-engine` and `/junos/system/subscriber-management/chassis/routing-engines/routing-engine[re-index=RoutingEngineIndex]` (to specify an index number for a specific Routing Engine)
- Subscriber management process sensor `/junos/system/subscriber-management/chassis/routing-engines/process-status/subscriber-management-processes/subscriber-management-process` and `/junos/system/subscriber-management/chassis/routing-engines/process-status/subscriber-management-processes/subscriber-management-process[pid=ProcessIdentifier]` (to specify a PID for a specific process)
- Per Routing Engine DHCP binding statistics for server or relay sensor `/junos/system/subscriber-management/chassis/routing-engines/routing-engine/dhcp-bindings/dhcp-element[dhcp-type-name=RelayOrServer/v4]` and `/junos/system/subscriber-management/chassis/routing-engines/routing-engine/dhcp-bindings/dhcp-element[dhcp-type-name=RelayOrServer/v6]`
- Virtual Chassis port counter sensor `/junos/system/subscriber-management/chassis/virtual-chassis-ports/virtual-chassis-port` and `/junos/system/subscriber-management/chassis/virtual-chassis-ports/virtual-chassis-port[vcp-interface-name=vcp-interface-port-string]` (to specify the interface name). This resource path is only supported on a virtual chassis.

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

- **CPU statistics support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 supports streaming various CPU statistics and process parameters using remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) services and Junos telemetry interface (JTI). You can stream CPU usage per process (statistics are similar to output from the `show system process detail` operational mode command), as well as CPU usage per Routing Engine core.

This feature supports the private data model `openconfig-procmon.yang`.

To stream statistics to an outside collector, include the following resource paths in a gRPC or gNMI subscription:

- Individual process level information (resource path `/system/processes/process/`)
- Individual Routing Engine core information (resource path `/components/component/cpu/`)

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

- **TARGET\_DEFINED subscription mode support with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX10016)**—Junos OS Release 20.2R1 adds support for TARGET-DEFINED mode for subscriptions made using gRPC Network Management Interface (gNMI) services.

Using a gNMI subscription, an external collector stipulates how sensor data should be delivered:

- STREAMING mode periodically streams sensor data from the DUT at a specified interval.
- ON\_CHANGE mode sends updates for sensor data from the DUT only when data values change.
- Newly supported TARGET\_DEFINED mode (submode 0) instructs the DUT to select the relevant mode (STREAMING or ON\_CHANGE) to deliver each element (leaf) of sensor data to the external collector. When a subscription for a sensor with submode 0 is sent from the external collector to the DUT, the DUT responds, activating the sensor subscription so that periodic streaming does not include any of the ON\_CHANGE updates. However, the DUT will notify the collector whenever qualifying ON\_CHANGE events occur.

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

- **Packet Forwarding Engine sensor support with INITIAL\_SYNC on JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000 line of routers, QFX5100, and QFX5200)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export Packet Forwarding Engine statistics from devices to an outside collector using gNMI submode INITIAL\_SYNC. When an external collector sends a subscription request for a sensor with INITIAL\_SYNC (gnmi-submode 2), the host sends all supported target leaves (fields) under that resource path at least once to the collector with the current value. This is valuable because:
  - The collector has a complete view of the current state of every field on the device for that sensor path.
  - Event-driven data (ON\_CHANGE) is received by the collector at least once before the next event is seen. In this way, the collector is aware of the data state before the next event happens.
  - Packet Forwarding Engine sensors that contain zero counter values (zero-suppressed) that normally do not show up in streamed data are sent, ensuring that all fields from each line card (also referred to as source) are known to the collector.

**NOTE:** ON\_CHANGE data is not available for native (UDP) Packet Forwarding Engine Sensors.

INITIAL\_SYNC submode requires that at least one copy to be sent to the collector; however, sending more than one is acceptable.

INITIAL\_SYNC submode is supported for the following sensors:

- Sensor for CPU (ukernel) memory (resource path `/junos/system/linecard/cpu/memory/`)
- Sensor for firewall filter statistics (resource path `/junos/system/linecard/firewall/`)

- Sensor for physical interface traffic (resource path `/junos/system/linecard/interface/`)
- Sensor for logical interface traffic (resource path `/junos/system/linecard/interface/logical/usage/`)
- Sensor for physical interface queue traffic (resource path `/junos/system/linecard/interface/queue/`)
- Sensor for physical interface traffic except queue statistics (resource path `/junos/system/linecard/interface/traffic/`)
- Sensor for NPU memory (resource path `/junos/system/linecard/npu/memory/`)
- Sensor for NPU utilization (resource path `/junos/system/linecard/npu/utilization/`)
- Sensor for packet statistics (resource path `/junos/system/linecard/packet/usage/`)
- Sensor for software-pollled queue-monitoring statistics (resource path `/junos/system/linecard/qmon-sw/`)

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

- **Export data using JSON encoding format with JTI (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and MX10016)**—Junos OS Release 20.2R1 adds support for JavaScript Object Notation (JSON) encoding to export telemetry data using gRPC network management interface (gNMI) services and Junos telemetry interface (JTI). JSON is an open standard file format and data interchange format that provides a good balance of usability and performance. It uses human-readable text to store and transmit data objects consisting of attribute–value pairs and array data types.

To export telemetry data using JSON encoding, include **format json-gnmi** at the `[edit services analytics export-profile profile-name]` hierarchy level. This is part of the export profile CLI configuration used to configure collector and sensor details in Junos OS.

[See [export-profile \(Junos Telemetry Interface\)](#).]

- **SR-TE statistics for uncolored SR-TE policies streaming on JTI (MX240, MX480, MX960, MX2010, and MX2020 with MPC-10E or MPC-11E)**—Junos OS Release 20.2R1 provides segment routing-traffic engineering (SR-TE) per label-switched path (LSP) route statistics using Junos telemetry interface (JTI) and remote procedure call (gRPC) services. Using JTI and gRPC services, you can stream SR-TE telemetry statistics for uncolored SR-TE policies to an outside collector.

Ingress statistics include statistics for all traffic steered by means of an SR-TE LSP. Transit statistics include statistics for traffic to the binding SID (BSID) of the SR-TE policy.

To enable these statistics, include the **per-source per-segment-list** statement at the `[edit protocols source-packet-routing telemetry statistics]` hierarchy level.

If you issue the **set protocols source-packet-routing telemetry statistics no-ingress** command, ingress sensors are not created.

If you issue the **set protocols source-packet-routing telemetry statistics no-transit** command, transit sensors are not created. Otherwise, if BSID is configured for a tunnel, transit statistics are created.

The following resource paths (sensors) are supported:

- `/junos/services/segment-routing/traffic-engineering/tunnel/lsp/ingress/usage/`
- `/junos/services/segment-routing/traffic-engineering/tunnel/lsp/transit/usage/`

To provision the sensor to export data through gRPC services, use the **telemetrySubscribe** RPC.

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

[See [Guidelines for gRPC and gNMI Sensors \(Junos Telemetry Interface\)](#), [source-packet-routing](#), and [show spring-traffic-engineering lsp detail name \*name\*](#).]

### **Layer 2 VPN**

- **Support for Layer 2 interworking (iw0) interface on the MPC10E and MPC11E line cards (MX Series)**—Starting in Junos OS Release 20.2R1, you can connect Layer 2 networks together by configuring a Layer 2 interworking (iw0) route with iw0 interfaces. This feature supports the following interconnections:
  - Layer 2 circuit to Layer 2 circuit
  - Layer 2 circuit to Layer 2 VPN
  - Layer 2 VPN to Layer 2 circuit
  - Layer 2 VPN to Layer 2 VPN

[See [Using the Layer 2 Interworking Interface to Interconnect a Layer 2 Circuit to a Layer 2 VPN](#) and [Layer 2 VPN to Layer 2 VPN Connections](#).]

### Layer 3 Features

- **MPC10E interoperates with MS-MPC/MS-MICs for Layer 3 Services (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2, the MPC10E interoperates with MS-MPC/MS-MICs for Layer 3 Services such as active flow monitoring, IPSec, NAT, RPM, and stateful firewall. [See [Layer 2 and Layer 3 Features on MX Series Routers](#).]

### Management

- **Error recovery, fault handling, and resiliency support for MX2K-MPC11E (MX2010 and MX2020)**—Starting in Junos OS Release 20.2R1, the MX2010 and MX2020 routers with the MX2K-MPC11E line card support error recovery, fault handling, and software resiliency. The MX2K-MPC11E line cards support detecting errors, reporting them through alarms, and triggering resultant actions. To view application-level errors, use the **show trace node fpc<#> application fabspoked-pfe** command. To check the status of the card, use the **show chassis fpc pic-status** command. Use the **show chassis errors active** command to view the fault details and the **show system alarm** command to view the alarm details.

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

### MPLS

- **Support to change the default re-merge behavior on the P2MP LSP (MX Series)**—Starting with Junos OS Release 20.2R1, you can change the default re-merge behavior on RSVP P2MP LSP. The term re-merge refers to the case of an ingress (headend) or transit node (re-merge node) that creates a re-merge branch intersecting the P2MP LSP at another node in the network. This may occur due to events such as an error in path calculation, an error in manual configuration, or network topology changes during the establishment of the P2MP LSP.

You can configure the no re-merge behavior on P2MP LSPs by enabling the newly introduced **no-re-merge** and **no-p2mp-re-merge** CLI commands at the ingress (headend) and transit devices (re-merge nodes), respectively.

[See [Re-merge Behavior on Point-to-Multipoint LSP Overview](#).]

- **Support for MPLS ping and traceroute for segment routing (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.2R1, we extend the MPLS ping and traceroute support for all types segment routing--traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

We also support the following features:

- FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
- ECMP traceroute support for all types of SR-TE paths.

We do not support the following:

- Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
- OAM for IPv6 prefix.
- BFD

[See [traceroute mpls segment-routing spring-te](#) and [ping mpls segment routing spring-te](#).]

- **MPLS support (MX Series routers with MPC10E and MPC11E)**—Starting in Junos OS Release 20.2R1, some of the MPLS features are supported on MX Series routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MX2K-MPC11E line cards.

[See [Protocols and Applications Supported by the MPC10E](#) and [Protocols and Applications Supported by the MX2K-MPC11E](#).]

### **Multicast**

- **Fast failover according to flow rate (MX Series with MPC10E or MPC11E line cards)**—Starting in Junos OS Release 20.2R1, for routers operating in Enhanced IP Network Services mode, you can configure a threshold that triggers fast failover in next-generation MVPNs with hot-root standby on the basis of aggregate flow rate. For example, fast failover (as defined in *Draft Morin L3VPN Fast Failover 05*) is triggered if the flow rate of monitored multicast traffic from the provider tunnel drops below the set threshold.

[See [min-rate](#).]

### **Network Management and Monitoring**

- **SNMP support for multicast LDP MIB objects (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS SNMP extends support for the following multicast LDP MIB tables and objects:

- mplsMldpInterfaceStatsTable
- mplsMldpFecUpstreamSessPackets
- mplsMldpFecUpstreamSessBytes
- mplsMldpFecUpstreamSessDiscontinuityTime

The multicast LDP standard MIB builds on the objects and tables that are defined in RFC3815, which only supports LDP point-to-point label-switched paths (LSPs). This multicast LDP MIB provides support for managing multicast LDP point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs.

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

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

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

- **Enhanced on-box monitoring support on the control plane (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure traceoptions to track all events related to system-level and process-level memory monitoring. You can also view the history of the actions taken for system-level and process-level memory monitoring by using the **show system monitor memory actions** command.

### **Next Gen Services**

- **Support for Dual Stack Lite (DS-Lite) Softwires**—Starting in Junos OS Release 20.2R1, Dual Stack Lite (DS-Lite) softwires are supported for CGNAT Next Gen Services. DS-Lite allows service providers to migrate to an IPv6 network while continuing to support IPv4 services; even after the exhaustion of the IPv4 address space. You can natively allocate IPv6 addresses to customers while legacy end-user devices accessing the IPv4 Internet remain same. Thus, IPv4 devices continue to access the IPv4 Internet with minimum disruption on their home networks. DS-Lite also de-couples IPv6 deployment in the service provider network from the rest of the Internet, making incremental deployment easier.

[See [DS-Lite Softwires—IPv4 over IPv6 for Next Gen Services](#).]

- **Support for HTTP Content Manager (HCM)**—Starting in Junos OS Release 20.2R1, HTTP Content Manager (HCM) is supported under Next Gen Services. HCM is an application that inspects the HTTP traffic transmitted through port 80 (default) or any other port you use to transmit HTTP traffic. HCM inspects HTTP traffic even if the default port 80 is not used for HTTP traffic and is interoperable with ms, rms, and ams interface types. It supports fragmented HTTP request packets and GET, PUT, and POST requests.

[See [HTTP Content Manager \(HCM\)](#).]

- **Support for Mapping of Address and Port with Encapsulation (MAP-E) Softwires for CGNAT Next Gen Services**—Starting in Junos OS Release 20.2R1, Mapping of Address and Port with Encapsulation (MAP-E) softwires are supported for CGNAT Next Gen Services. MAP-E is an automatic tunneling mechanism tailored for deployment of IPv4 to end users via a service provider's IPv6 network infrastructure. Using MAP-E technology, islands of v4 networks can be connected via v6 tunnels. The IPV4 packets are carried in IPV4-over-IPV6 tunnels from the MAP-E Customer Edge (CE) routers to the MAP-E Border Relay(s) (BR) (through IPV6 routing topology), where they are de-tunneled for further processing. MAP-E can be used by Service Providers to provide IPv4 connectivity to their subscribers over the ISP's IPv6 access network.



[See [Mapping of Address and Port with Encapsulation \(MAP-E\) for Next Gen Services.](#)]

- **Support for Network Address Translation and Protocol Translation for CGNAT Next Gen Services**—Starting in Junos OS Release 20.2R1, Network Address Translation and Protocol Translation (NAT-PT) [RFC2766] are supported for CGNAT Next Gen Services. NAT-PT is a IPv4-to-IPv6 transition mechanism that provides a way for end-nodes in IPv6 realm to communicate with end-nodes in IPv4 realm and vice versa. This is achieved using a combination of Network Address Translation and Protocol Translation.

[See [NAT46 Next Gen Services Configuration Examples.](#)]

- **Support for Port Control Protocol Support (PCP) for DS-Lite for CGNAT Next Gen Services**—Starting in Junos OS Release 20.2R1, Port Control Protocol Support (PCP) for DS-Lite is supported for CGNAT Next Gen Services. DS-Lite is a technology which enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT).

Typically, the home gateway embeds a Basic Bridging BroadBand (B4) capability that encapsulates IPv4 traffic into a IPv6 tunnel to the CGNAT, named the Address Family Transition Router (AFTR). AFTRs are run by service providers.

PCP allows customer applications to create mappings in a NAT for new inbound communications destined to machines located behind a NAT. In a DS-Lite environment, PCP servers control AFTR devices.

[See [Port Control Protocol Overview.](#)]

### ***Operation, Administration, and Maintenance (OAM)***

- **Support for connectivity fault management (CFM) on MPC10E and MX2K-MPC11E**—Starting in Junos OS Release 20.2R1, you can configure the IEEE 802.1ag OAM CFM Down maintenance association end points (MEPs) on MPC10E and MX2K-MPC11E to monitor Ethernet networks for connectivity faults.

Junos OS supports the continuity check messages (CCM) and loopback messages as defined in IEEE 802.1ag.

[See [Configuring Connectivity Fault Management.](#)]

### ***Routing Policy and Firewall Filters***

- **ARP policer support on pseudowire interfaces (MX Series)**—Starting in Junos OS Release 20.2R1, you can create policers for ARP traffic on pseudowire interfaces. Configure rate limiting for the policer by specifying the bandwidth and the burst-size limit of a firewall policer and attaching the policy to a pseudowire interface, just like you would any other interface. Traffic that exceeds the specified rate limits can be dropped or marked as low priority and delivered when congestion permits.

In the case of denial of service (DoS) or ARP broadcast storms, ARP policers protect the Routing Engine against malicious traffic intended to degrade the network.

Apply the ARP policer to a pseudowire interface at the [**edit interfaces *interface-name* unit *unit-number* family inet policer arp *policy-name***] level of the hierarchy.

[See [ARP Policer Overview](#).]

- **Support for P2MP and P2P automatic LSP policers (MX Series)**—Starting in Junos OS Release 20.2R1, support for automatic policers on point-to-multipoint (P2MP) label-switched paths (LSPs) is available on MX240, MX480, MX960, MX2010, and MX2020 routers with MPC10E and MPC11E line cards.

P2MP MPLS LSP is either an LDP-signaled, or RSVP-signaled, LSP with a single source and multiple destinations that can optimize packet replication at the ingress router. With it, packet replication only occurs for packets being forwarded to two or more different destinations requiring different network paths. Automatic LSP policing lets you provide strict service guarantees for network traffic in accordance with the bandwidth configured for the LSPs.

Also supported with this release are the following features:

- Graceful Routing Engine switchover (GRES) at the ingress and egress
- Load balancing over aggregated links
- P2MP statistics
- Multiprotocol BGP-based multicast VPNs (or Layer 3 VPN multicast)

[See [Configuring Automatic Policers](#).]

- **Support for firewall forwarding (MX Series)**—Starting in Junos OS Release 20.2R1, the following traffic policers are supported on MX240, MX480, MX960, MX2010, and MX2020 routers with MPC10E or MPC11E line cards:
  - 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
  - Input and output filter chains
  - Actions, including policy-map filters, do-not-fragment, and prefix
  - Layer 2 policers
  - Policer overhead adjustment
  - Hierarchical policers
  - Shared bandwidth
  - Percentages
  - Logical interfaces

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

## Routing Protocols

- **TI-LFA SRLG protection for IS-IS (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure Shared Risk Link Group (SRLG) protection for segment routing to choose a fast reroute path that does not include SRLG links in the topology-independent loop-free alternate (TI-LFA) backup paths. This is in addition to existing fast reroute options such as **link-protection**, **node protection**, and **fate-sharing protection** for segment routing. IS-IS computes the fast reroute path that is aligned with the post-convergence path and excludes the SRLG of the protected link. All local and remote links that are from the same SRLG as the protected link are excluded from the TI-LFA back up path. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface.

To enable TI-LFA SRLG protection with segment routing for IS-IS, include the **srlg-protection** statement at the **[edit protocols isis interface name level number post-convergence-lfa]** hierarchy level.

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

- **Support for BGP-LU over SR-TE for color-based mapping of VPN Services (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, we are extending support to BGP labeled unicast service for color-based mapping of VPN services over Segment Routing-Traffic Engineering (SR-TE). This enables you to advertise BGP-LU IPv6 and IPv4 prefixes with an IPv6 next-hop address in IPv6-only networks where routers do not have any IPv4 addresses configured. With this feature, BGP-LU can now resolve IPv4 and IPv6 routes over SR-TE core. BGP-LU constructs a colored protocol next hop, which is resolved on a colored SR-TE tunnel in the **inetcolor.0** or **inet6color.0** table. Currently we support BGP IPv6 LU over SR-TE with IS-IS underlay.

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

- **Support for AIGP metric to MED translation (MX2010 and MX2020)**—Starting in Release 20.2R1, Junos OS supports the translation of AIGP metric to MED. You can enable this feature when you want the end to end effective AIGP metric in order to choose the best path. Effective AIGP is the AIGP value advertised with the route plus the IGP cost to reach the nexthop. This is especially useful in Inter-AS MPLS VPNs solution, where customer sites are connected via two different service providers, and customer edge routers want to take IGP metric based decision. You can configure a **minimum-aigp** to prevent unnecessary update of route when effective-aigp changes past the previously known lowest value.

The following configuration statements are introduced at the **[edit protocols bgp group <group-name> metric-out]** hierarchy level:

- **effective-aigp** to track the effective AIGP metric
- **minimum-effective-aigp** to track the minimum effective AIGP metric.

[See [effective-aigp](#) and [minimum-effective-aigp](#).]

- **Support for Layer 2 circuit, Layer 2 VPN, and VPLS services with BGP labeled unicast (MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices)**—Starting with Junos OS Release 20.2R1, MX Series, EX9204, EX9208, EX9214, EX9251, and EX9253 devices support BGP PIC Edge protection for Layer 2 circuit, Layer 2 VPN, and VPLS (BGP VPLS, LDP VPLS and FEC 129 VPLS) services with BGP

labeled unicast as the transport protocol. BGP PIC Edge using the BGP labeled unicast transport protocol helps to protect traffic failures over border nodes (ABR and ASBR) in multi-domain networks. Multi-domain networks are typically used in metro-aggregation and mobile backhaul networks designs.

A prerequisite for BGP PIC Edge protection is to program the Packet Forwarding Engine (PFE) with expanded next-hop hierarchy.

To enable BGP PIC Edge protection, use the following CLI configuration statements:

- Expand next-hop hierarchy for BGP labeled unicast family:

```
[edit protocols]
user@host#set bgp group group-name family inet labeled-unicast nexthop-resolution
preserve-nexthop-hierarchy;
```

- BGP PIC for MPLS load balance nexthops:

```
[edit routing-options]
user@host#set rib routing-table-name protect core;
```

- Fast convergence for Layer 2 circuit and LDP VPLS:

```
[edit protocols]
user@host#set l2circuit resolution preserve-nexthop-heirarchy;
```

- Fast convergence for Layer 2 VPN, BGP VPLS, and FEC129:

```
[edit protocols]
user@host#set l2vpn resolution preserve-nexthop-heirarchy;
```

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

- **Support for dynamic peer AS range for BGP groups (ACX Series, MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, you can configure acceptable autonomous system (AS) ranges for EBGp groups that can be used for bringing up BGP peers while establishing a BGP session. BGP accepts a peer request based on the configured AS range and rejects a peer request if the AS does not fall into the specified range. This allows you to control BGP peering when the neighbor's exact IP address is not known.

To define peer AS range for BGP groups through policy, you can include the **as-list** statement at the **[edit policy-options]** hierarchy level. To include the specified peer AS list, include the **peer-as-list** **peer-as-list** statement at the **[edit protocols bgp group *group-name*]** hierarchy level.

See [[peer-as-list](#) and [as-list](#).]

- **Support for BGP-SR-TE rearchitecture (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS provides support for controller-based BGP segment routing--traffic engineering (SR-TE) routes

to be installed as source packet routing traffic-engineered (SPRING-TE) routes. BGP installs the SR-TE policy in the routing tables `bgp.inetcolor.0` and `bgp.inet6color.0`, and these routes are subsequently installed in the routing tables `inetcolor.0` or `inet6color.0` by SPRING-TE.

In releases before Junos OS Release 20.2R1, controller-based BGP SR-TE routes are installed as BGP routes in the routing table. To maintain consistency and for easy maintenance, all SR-TE based routes appear as SPRING-TE routes irrespective of the source.

You need to enable **source-packet-routing** at the **[edit protocols]** hierarchy level to see the routes installed in `inetcolor.0` or `inet6color.0`. A new option **detail** is introduced under **traceoptions (Protocols Spring-TE)** to trace the detailed information.

See [[Segment Routing Traffic Engineering at BGP Ingress Peer Overview.](#)]

- **Support for egress protection and BGP PIC features (MX Series Routers with MPC10E and MPC11E)**—Starting in Junos OS Release 20.2R1, you can configure the following egress link protection and BGP Prefix Independent Convergence (PIC) features on MX Series devices with MPC10E and MPC11E.
  - **Egress protection for BGP labeled unicast** —Fast protection for egress nodes is available to services in which BGP labeled unicast interconnects IGP areas, levels, or autonomous systems (ASs). If a provider router detects that an egress router (AS or area border router) is down, it immediately forwards the traffic destined to that router to a protector router that forwards the traffic downstream to the destination.
  - **Provider-edge link protection for BGP labeled unicast paths**—You can configure a precomputed protection path in a Layer 3 VPN such that if a BGP labeled-unicast path between an edge router in one AS and an edge router in another AS goes down, you can use the protection path (also known as the backup path) between alternate edge routers in the two ASs. This is useful in a carrier-of-carriers deployments, where a carrier can have multiple labeled-unicast paths to another carrier. In this case, the protection path avoids disruption of service if one of the labeled-unicast paths goes down.
  - **BGP PIC for inet** —We've extended the BGP Prefix Independent Convergence (PIC) support to BGP with multiple routes in the global tables such as `inet` and `inet6` unicast, and `inet` and `inet6` labeled unicast. When you enable the BGP PIC feature on a router, BGP installs to the Packet Forwarding Engine the second best path in addition to the calculated best path to a destination. When an IGP loses reachability to a prefix, the router uses this backup path to reduce traffic loss until the global convergence through BGP is resolved, thereby drastically reducing the outage duration.
  - **BGP (PIC Edge for RSVP** —With BGP PIC Edge in an MPLS VPN network, IGP failure triggers a repair of the failing entries and causes the Packet Forwarding Engine to use the prepopulated protection path until global convergence has re-resolved the VPN routes. The convergence time is no longer dependent on the number of prefixes. When RSVP receives a tunnel down notification at the ingress PE router, it sends a notification to the Packet Forwarding Engine to start making use of the tunnel to the alternate egress PE router.

[See [Egress Protection for BGP Labeled Unicast](#) ,[Understanding Provider Edge Link Protection for BGP Labeled Unicast Paths](#), [Use Case for BGP PIC for Inet](#), and [show rsvp version.](#)]

### Services Applications

- **Interoperability of MPC10E with MS-MPC and MS-MIC for Layer 3 Services ( MX240, MX480, and MX960)**—Starting in Junos OS Release 20.2R1, the MPC10E-15C-MRATE interoperates with MS-MPC and MS-MIC-16G to support the following Layer 3 Services:
  - Stateful firewall
  - NAT
  - IPSec
  - RPM
  - MS-MPC/MS-MIC based Inline flow monitoring services

- **Support for RFC 2544-based benchmarking tests (MX Series routers with MPC10E and MX2K-MPC11E)**—Junos OS Release 20.2 extends support for the reflector function and the corresponding RFC 2544-based benchmarking tests on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MX2010 and MX2020 routers with MX2K-MPC11E. The RFC 2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before activation of the service. The tests measure throughput, latency, frame loss rate, and back-to-back frames.

RFC 2544-based benchmarking tests on MX Series routers support the following reflection functions:

- Ethernet pseudowire reflection (ingress and egress direction) (ELINE service—supported for family `ccc`)
- Layer 2 reflection (egress direction) (ELAN service—supported for family `bridge`, `vpls`)
- Layer 3 IPv4 reflection (limited support)

To run the benchmarking tests on the MX Series routers, you must configure reflection (Layer 2 or pseudowire) on the supported MPC. To configure the reflector function on the MPC, use the `fpc fpc-slot-no slamon-services rfc2544` statement at the `[edit chassis]` hierarchy level.

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

- **Support for random load balancing (MX Series routers with MPC10E and MX2K-MPC11E)**—Starting in Junos OS Release 20.2R1, you can configure per packet random load balancing on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MX2010 and MX2020 routers with MX2K-MPC11E. Per-packet random spray load balancing ensures that the members of ECMP are equally loaded without taking bandwidth into consideration. Random load balancing also eliminates traffic imbalance that occurs as a result of software errors, except for packet hash.

To configure random load balancing on the MPC, include the `load-balance random` statement at the `[edit policy-options policy-statement policy-name term term-name then]` hierarchy level.

[See [Understanding the Algorithm Used to Load Balance Traffic on MX Series Routers](#)].

- **Support for static IP tunnels (MX Series routers with MPC10E and MX2K-MPC11E)**—Starting in Junos OS Release 20.2R1, MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and

MPC10E-10C-MRATE) and MX2010 and MX2020 routers with MX2K-MPC11E support static IP tunnels with:

- Encapsulation support of the following types:
  - IPv4-over-IPv4
  - IPv6-over-IPv4
  - IPv4-over-IPv6
  - IPv6-over-IPv6
- Scaling upto 4000 tunnels per PIC
- Graceful Routing Engine switchover (GRES)

### **Software-Defined Networking (SDN)**

- **Manual (PIM-based) VXLAN support (MPC10 and MPC11 line cards with MX2010 and MX2020)**—Starting in Junos OS Release 20.2R1, the MX2010 and MX2020 routers with MPC10 and MPC11 line cards installed support manual (PIM-based) VXLAN.

[See [Understanding VXLANs.](#)]

- **GNFs with MX-SPC3 support carrier-grade NAT services over abstracted fabric interfaces (MX480 and MX960)**—Starting in Junos OS Release 20.2R1, guest network functions running Next Gen Services with the MX-SPC3 card support carrier-grade NAT services.

The support includes the following:

- NAT translation types—dnat-44, dynamic-nat44, basic-nat44, basic-nat66, twice-basic-nat-44, twice-dynamic-nat44, deterministic NAT. Support for interface and next-hop style service sets, EIM/EIF, PBA, XLAT464, and port forwarding are available. Support for basic-nat44, basic-nat66 over layer 3 VPN is also available.
- SIP and RTSP Application Layer Gateways
- carrier-grade events logging, using the Junos Traffic Vision (J-Flow).
- Class of service (CoS)

**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 [Junos OS Carrier-Grade NAT Implementation Overview.](#)]

- **GNFs with MX-SPC3 support various services over abstracted fabric interfaces (MX480 and MX960)**—Starting in Junos OS Release 20.2R1, guest network functions (GNFs) running Next Gen Services with the MX-SPC3 card support the following services over abstracted fabric interfaces:

- DNS filtering to identify DNS requests for blacklisted website domains.
- URL filtering to determine which Web content is not accessible to users.

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 [DNS Request Filtering for Blacklisted Website Domains](#) and [Configuring URL Filtering](#)]

### ***Subscriber Management and Services***

- **RADIUS-sourced connection status updates to CPE devices (MX Series)**—Starting in Junos OS Release 20.2R1, you can use RADIUS-sourced messages to convey information, such as upstream bandwidth or connection rates, that the BNG transparently forwards to CPE devices. Configure RADIUS to send the router the Juniper Networks Connection-Status-Message VSA (26-4874-218) in Access-Accept or CoA messages. Include the **lcp-connection-update** PPP option in the client dynamic profile to enable PPP to send the VSA contents to the CPE device in the Connection-Status-Message option of an LCP Connection-Update-Request message.

[See [RADIUS-Sourced Connection Status Updates to CPE Devices](#).]

- **Identifying dynamic profile versions with version aliases (MX Series)**—Starting in Junos OS Release 20.2R1, you can use the **versioning-alias** statement to configure a text description that identifies a particular variation of a dynamic client profile. The version alias is conveyed to the RADIUS server in the Access-Accept message in the Juniper Networks Client-Profile-Name VSA (26-4874-174).

[See [Versioning for Dynamic Profiles](#).]

- **IPFIX support for per-subscriber queue statistics (MX Series)**—Starting in Junos OS Release 20.2R1, you can configure the input-jti-ipfix plug-in to collect per-subscriber interface queue statistics. The output ipfix-plugin can then export the statistics as IPFIX template and data records.

[See [Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector](#).]

- **Junos Multi-Access User Plane support (MX204, MX10003)**—Starting with Junos OS Release 20.2R1, you can configure Junos Multi-Access User Plane on MX204 and MX10003 routers. Junos Multi-Access User Plane is a software solution that turns your MX Series router into a high-capacity user plane function called a System Architecture Evolution Gateway-User Plane (SAEGW-U). This MX Series SAEGW-U interoperates with a third-party SAEGW-C (control plane function), according to the 3GPP Release 14 Control User Plane Separation (CUPS) architecture, to provide high-throughput 4G fixed-wireless access service. CUPS enables independent scaling of the user and control planes, network architecture flexibility, operational flexibility, and an easier migration path from 4G to 5G services. The CUPS architecture is optional for 4G but inherent in 5G architecture.

[See [Junos Multi-Access User Plane User Guide](#).]



### System Logging

- **Support to track the maximum number of routing and forwarding (RIB/FIB) routes and VRFs (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can track and display the high-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system (RPD) using the **show route summary** CLI command. High-water mark refers to the maximum number of routing and forwarding (RIB/FIB) table routes and VRFs that was present in the RPD system. The high-water mark data can also be viewed in the syslog at the **LOG\_NOTICE** level.

You can configure the interval of the high-water mark data using the **highwatermark-log-interval** CLI configuration statement at the **[edit routing-options]** hierarchy level. The minimum time gap at which the high-water mark data logged in the syslog is 30 seconds. You can configure the value for **highwatermark-log-interval** CLI configuration statement between 5 to 1200 seconds.

[See [routing-options](#) and [show route summary](#).]

### System Management

- **Support for the G.8275.1 Profile (MX10008 and MX10016 with line card JNP10K-LC2101)**—Starting in Junos OS Release 20.2R1, we support ITU-T G.8275.1 Full path Timing Support (FTS) Profile and G.8273.2 Telecom Boundary Clock. The G.8275.1 Profile is a phased profile that operates with PTP-based packet exchange for Phase and Time recovery, and Synchronous-Ethernet-based based frequency recovery (also called *Synchronous-Ethernet-based assisted PTP mode of operation*). This profile is required in TDD application deployment in both 4G and 5G networks.

The PTP operation must be two-way in this profile in order to transport phase/time synchronization because propagation delay must be measured. Hybrid mode must be enabled for the G.8275.1 profile.

[See [profile-type](#).]

### Virtual Chassis

- **MX Series Virtual Chassis support for the ephemeral database (MX480 and MX960)**—Starting in Junos OS Release 20.2R1, MX Series Virtual Chassis support configuring the ephemeral database. The ephemeral database is an alternate configuration database that provides a fast programmatic interface for performing configuration updates on devices running Junos OS.

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

### SEE ALSO

[What's Changed | 129](#)

[Known Limitations | 136](#)

[Open Issues | 139](#)

[Resolved Issues | 149](#)

[Documentation Updates | 185](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 129](#)
- [What's Changed in Release 20.2R2 | 131](#)
- [What's Changed in Release 20.2R1 | 133](#)

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

### What's Changed in Release 20.2R3

#### *General Routing*

- **Updates to ON-CHANGE and periodic dynamic subscriber interface metadata sensors (MX Series routers and EX9200 line of switches)**—We've made the following updates to the `/junos/system/subscriber-management/dynamic-interfaces/interfaces/meta-data/interfacesid='sid-value'/` sensor:
  - Notifications are sent when subscribers log in on either IP demux or VLAN demux interfaces. In earlier releases, login notifications are sent only for IP demux logins.
  - The **interface-set** end path has been added to the logical interface metadata. The `interface-set` field appears in both ON-CHANGE and periodic notifications. In earlier releases, this field is not included in the sensor metadata or notifications.

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

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

[See [iccp](#).]

### *Junos XML API and Scripting*

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

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

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

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

### *Layer 2 Ethernet Services*

- Active leasequery based bulk leasequery (MX Series)—The `overrides always-write-option-82` and `relay-option-82 circuit-id` configuration at the `edit forwarding-options dhcp-relay` hierarchy level is not mandatory for active leasequery based bulk leasequery. Earlier to this release, the `overrides always-write-option-82` and `circuit-id` configurations are mandatory for active leasequery based bulk leasequery. For regular bulk leasequery between relay and server without any active leasequery, the `overrides always-write-option-82` and `relay-option-82 circuit-id` configurations are mandatory.

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

### *Network Management and Monitoring*

- Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:
  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.
  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and emits only an `<ok>` or `<rpc-error>` element.

[See [Configuring RFC-Compliant NETCONF Sessions.](#)]

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the **edit system services netconf hello-message yang-module-capabilities** hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the **edit system services netconf netconf-monitoring netconf-state-schemas** hierarchy level.

[See [hello-message.](#) and [netconf-monitoring..](#)]

### *User Interface and Configuration*

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

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

## What's Changed in Release 20.2R2

### *EVPN*

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

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

- **Warning message for proxy MAC advertisement (MX Series)**—When **proxy-macip-advertisement** is enabled, the Layer 3 gateway advertises MAC and IP routes (MAC+IP type 2 routes) on behalf of Layer 2 VXLAN gateways in EVPN-VXLAN networks. This behavior is not supported on EVPN-MPLS. Starting in Junos OS Release 20.2R2, the warning message, **WARNING: Only EVPN VXLAN supports proxy-macip-advertisement configuration**, appears when you enable **proxy-macip-advertisement**. The message appears when you change your configuration, save your configuration, or use the **show** command to display your configuration.

[See [proxy-macip-advertisement.](#)]

### General Routing

- **MS-MPC and MS-MIC service package (MX240, MX480, MX960, MX2008, MX2010, and MX2020)**—PICs of MS-MPC and MS-MIC do not support any other service package than extension-provider. These PICs always come up with the extension-provider service-package, regardless 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](#).]

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

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

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

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

### *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. In Junos OS Release 20.2R2, 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 (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. In Junos OS Release 20.2R2, 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.

### *Routing Protocols*

- **Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised as router IDs.

### *Subscriber Management and Services*

- **Improved tunnel session limits display (MX Series)**—Starting in Junos OS Release 20.2R2, 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](#).]

## What's Changed in Release 20.2R1

### 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

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the **persist-groups-inheritance** option at the `[edit system commit]` hierarchy level is enabled by default. To disable this option, use **no-persist-groups-inheritance**.

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

- **Install or activate the RIFT package to include the request rift package activate-as-top-of-fabric option**—Install or activate the RIFT package to include the **request rift package activate-as-top-of-fabric** option. This option is same as the **activate** option but it adds additional configuration to act as a **top-of-fabric** node.
- **Command to view summary information for resource monitor (EX9200 line of switches and MX Series)**—You can use the **show system resource-monitor** command to view statistics about the use of memory resources for all line cards or for a specific line card in the device. The command also displays information about the status of load throttling, which manages how much memory is used before the device acts to reduce consumption.

[See [show system resource-monitor](#) and [Resource Monitoring for Subscriber Management and Services](#).]

### Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the **PASS** keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

[See [Develop Off-Device JET Applications](#) and [Develop On-Device JET Applications](#).]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH\_USAGE\_INVALID**.

[See [Juniper EngNet](#).]

### **Network Management and Monitoring**

- **Support for new SNMP object for the ifJnx MIB**—Starting in Junos OS Release 20.2R1, we introduce a new SNMP object, **ifJnxInputErrors**, that tracks all input errors except the L3 incomplete errors. The **ifJnxInErrors** object continues to track the L3 incomplete errors.
- **Support for Clearing the Event at MEP Level (MX Series)**—In Junos OS 20.2R1, you can define an action profile for connectivity fault management at the local MEP level or at the remote MEP level. You define an action profile to monitor events and thresholds and specify an action that the device performs when the configured event occurs. When you define the action profile at the local MEP level, you can clear the event for the configured action profile at the local MEP level by specifying only the local MEP numeric identifier. When you define the action profile at the remote MEP level, you can clear the event for the configured action profile at the remote MEP level by specifying the local MEP numeric identifier as well as the remote MEP numeric identifier.

See [[clear oam ethernet connectivity-fault-management event](#).]

- **Request support information for IPsec function (MX Series)**—Starting in Release 20.2R1, Junos OS introduces **ipsec-vpn** option to the existing **request support information** command. The **request support information ipsec-vpn** command displays all the configurations, states, and statistics at Routing Engine and Service Card level. This new option helps in debugging IPsec-VPN related issues. The information collection is streamlined and reduces the output file size.

See [[Request support information](#).]

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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



### Services Applications

- **New option for configuring delay in IPsec SA installation**—In Junos OS Releases 20.2R1 and 20.2R2, 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 SA in a NAT-T scenario soon after the IPsec SA negotiation is complete. The default value is 0 seconds.

### Software-Defined Networking (SDN)

- **JDM install and configuration do not impact host SNMP**—Starting in Junos OS Release 20.2R1, JDM does not write any configuration to the host SNMP configuration file (`/etc/snmp/snmpd.conf`). Hence, JDM installation and subsequent configuration do not have any impact on the host SNMP. The SNMP configuration CLI command in JDM is used only to configure JDM's `snmpd.conf` file, which is present within the container.

[See [SNMP Trap Support: Configuring NMS Server \(External Server Model\)](#).]

### SEE ALSO

[What's New | 102](#)

[Known Limitations | 136](#)

[Open Issues | 139](#)

[Resolved Issues | 149](#)

[Documentation Updates | 185](#)

[Migration, Upgrade, and Downgrade Instructions | 186](#)

## Known Limitations

### IN THIS SECTION

- [General Routing | 137](#)
- [Infrastructure | 138](#)
- [Interfaces and Chassis | 138](#)
- [MPLS | 138](#)
- [Network Management and Monitoring | 138](#)
- [Platform and Infrastructure | 138](#)
- [Routing Protocols | 139](#)

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

## General Routing

- On the MPC11E line card, the **number-of-sub-ports** configuration on the 4x10GbE channelized ports might cause the channels to go down. [PR1442439](#)
- On the MPC11E line card, the following error messages are observed when the line card is online: **i2c transaction error (0x00000002)**. [PR1457655](#)
- Traffic stops after reaching the volume limit but the traffic resumes after the Packet Forwarding Engine fails. [PR1463723](#)
- 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](#)
- Dynamic SR-TE tunnels do not get automatically recreated at the new primary Routing Engine after the Routing Engine switchover. [PR1474397](#)
- Packet Forwarding Engine lookup loop occurs when the firewall based redirection under forwarding-options is used to perform route-lookup in a non-default routing instance for destinations reachable over MPLSoUDP tunnels. [PR1478000](#)
- The following message might be observed while configuring MTU: **SNMP\_TRAP\_LINK\_DOWN**. [PR1486542](#)
- The rpd process might generate core files in the absence of an explicit route-distinguisher configuration. [PR1486922](#)
- After executing the **clear interfaces statistics all** command, the value might be different from the values of the output of the **show interfaces** command. [PR1488758](#)
- It takes nearly 20 minutes to display IP-IP tunnel statistics on the backup Routing Engine after GRES at full scale of 4000 tunnels. [PR1489067](#)
- Packets do not get fragmented based on FTI interface MTU in the data path. [PR1489526](#)
- Traffic drop of around 2.5 seconds on switchover from primary physical interface is observed to backup FTI interface with the scaled routes. [PR1490070](#)
- The sequence-numbers (initial-synchronization and regular streaming) might be in the wrong order when multiple collectors are present. [PR1490798](#)
- The basic service set identifier (BSSID) scaling limits for IPv6 policies are 16,000 per ECMP. [PR1495330](#)
- The pppmd restart does not clear the active RFC2544 reflection sessions. [PR1499285](#)

- Active reflection sessions are not aborted when the delete interfaces and the delete services configuration is committed. [PR1499628](#)
- One hundred percent traffic drop at tunnel destination is observed if fragmentation is enabled when the incoming packet size is greater than the egress WAN MTU. [PR1505209](#)
- Changing the scaled firewall profiles on the fly does not release the TCAM resources as expected. [PR1512242](#)

## Infrastructure

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

## Interfaces and Chassis

- Session fails to come up after the outer tag pop when ingress and egress logical interfaces are on the same Packet Forwarding Engine. [PR1487351](#)
- On the MPC10 or MPC11 line card, the convergence goes up to 38 seconds for a highly scaled configuration. [PR1519373](#)

## MPLS

- The P2MP branches stay on bypass even after the link becomes functional after failure. [PR1486813](#)
- The RPD process might crash. [PR1461468](#)
- After enabling the MPLS p2mp-lsp no-re-merge set protocols on ingress, the P2MP branches fail to come up. [PR1487007](#)
- Branches do not select the common ASBR from the available list with the single-asb command enabled after the common ASBR failure. [PR1490637](#)

## Network Management and Monitoring

- On the MPC11E line card, the following trap message is not observed after a line card reboot when the scaled interfaces are present: **SNMP Link up.** [PR1507780](#)

## 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](#)
- The JTI sensor subscription and the related TCP session are still present after the interface is deleted, deactivated, or disabled. [PR1477790](#)

## Routing Protocols

- RPKI validation is broken. [PR1464931](#)

### SEE ALSO

[What's New | 102](#)

[What's Changed | 129](#)

[Open Issues | 139](#)

[Resolved Issues | 149](#)

[Documentation Updates | 185](#)

[Migration, Upgrade, and Downgrade Instructions | 186](#)

## Open Issues

### IN THIS SECTION

- [Class of Service \(CoS\) | 140](#)
- [EVPN | 140](#)
- [Forwarding and Sampling | 140](#)
- [General Routing | 141](#)
- [High Availability \(HA\) and Resiliency | 145](#)
- [Infrastructure | 145](#)
- [Interfaces and Chassis | 145](#)
- [Layer 2 Ethernet Services | 146](#)
- [MPLS | 146](#)
- [Platform and Infrastructure | 146](#)
- [Routing Policy and Firewall Filters | 147](#)
- [Routing Protocols | 147](#)

- Services Applications | 148
- Subscriber Access Management | 148
- User Interface and Configuration | 148
- VPNs | 149

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)

- The following syslog error message is observed: **cosd[10290]: LIBCOS\_COS\_ATTRIBUTE\_RETRIEVE\_FAILED: FAILED to retrieve cos field (cos\_fc\_defaults\_0\_fc\_no\_loss)**. [PR1470252](#)
- The **mpls-inet-both-non-vpn** command does not work as expected. [PR1479575](#)
- When an interface attached to the aggregated Ethernet interface is decoupled and an IP address is assigned to it, ARP resolution issues are observed. [PR1504287](#)

## EVPN

- There might be a few duplicate packets seen in an active/active EVPN scenario when the remote PE device sends packets with IM label due to MAC not being learned on remote PE device but being learned on the active/active local PE device. The non-DF sends the IM-labeled encapsulated packet to the PE-CE interface after MAC looks up instead of dropping the packet, which causes duplicate packets on the CE side. [PR1245316](#)
- The VXLAN OAM host-bound packets are not throttled with DDoS policers. [PR1435228](#)
- The **mustd.core** process generates core file during upgrading or while committing a configuration. [PR1577548](#)

## Forwarding and Sampling

- Packet length for ICMPv6 is shown as **0** in the output of the **show firewall log detail** command. [PR1184624](#)
- The log message of **Prefix-List [] in Filter []** that does not have any relevant prefixes might not be seen when the IPv4 prefix is added on a prefix list referred by the IPv6 firewall filter. [PR1395923](#)

- The following syslog error message might be observed due to SSD hardware failure: **Failed connecting to DFWD, error checking reply - Operation timed out.** [PR1397171](#)
- After restarting the router, the remote mask (indicating from which remote PE devices MAC IP addresses are learned), that the routing daemon sends might be different from the existing remote mask compared to the Layer 2 learning daemon had prior to restart. This causes a mismatch between the Layer 2 learning and routing daemon interpretation as to where the MAC IP address entries are learned (either local or remote) leading to the MAP IP table being out of synchronization. [PR1452990](#)

## General Routing

- The host root file system and the node boot with the previous vmhost software instead of the alternate disk. [PR1281554](#)
- Not using the chained CNH does not bring in a lot of gain because TCNH is based on an ingress rewrite premise. Without this feature, things work just fine. [PR1318984](#)
- With regards to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infra relies on the integrity of the TCP connections. The reactions to failure situations might not be handled gracefully, resulting in TCP connection timeouts because of jlock hog crossing the boundary value (5 seconds), which causes bad consequences in MX Series Virtual Chassis. Currently, there is no other easy solution to reduce this jlock hog besides enabling marker infra in the MX Series Virtual Chassis setup. [PR1332765](#)
- In an MS-MPC or MS-MIC in ALG scenario, the MAC\_STUCK message might be observed and traffic might be dropped. [PR1335956](#)
- The backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- The following error messages are observed with Junos OS Release 17.3 throttle image:  
**localntp\_offload\_tx\_errcheck: failed to send packet 4 times in last one second.** [PR1359149](#)
- On the MX204 and MX10003 routers, the following garbage value on syslog messages from craftd demon is observed: craftd[xxxx]: **fatal error, failed to open smb device: JÎÈ.** [PR1359929](#)
- On the MX2010 and MX2020 routers equipped with SFB2, some error logs might be seen. [PR1363587](#)
- Due to transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events are disabled to prevent alarms and possibly unnecessary hardware replacements. [PR1384435](#)
- The virtio throughput remains the same for the multi-queue and single-queue deployments. [PR1389338](#)
- Revert of RLT to primary might silently discard traffic for around 10 minutes after the primary FPC is online with primary RLT up. [PR1394026](#)
- The FPC generates core files under certain circumstances on addition and deletion of hierarchical CoS from pseudowire devices. [PR1414969](#)
- Traffic statistics are not displayed for the hybrid access gateway session and tunnel traffic. [PR1419529](#)

- 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](#)
- Dynamic tunnel summary displays wrong count of up and total tunnels. [PR1429949](#)
- The ike-esp session are not created after enabling **ike-esp-nat**. [PR1516655](#)
- The ALG timeout value is displayed as default value for the child data sessions even after the configured service set timeout values. [PR1516697](#)
- Need to show which shard a given route is hashed to. [PR1430460](#)
- Layer 2 over GRE is not supported in Junos OS Release 19.3R1. Although, the configuration gets committed, the feature does not work. [PR1435855](#)
- The FPC process might crash when the Packet Forwarding Engine memory is exhausted. [PR1439012](#)
- Interface hold-down timers cannot be achieved for less than 15 seconds on the MPC11E line card. [PR1444516](#)
- The vehostd application fails to generate a minor alarm. [PR1448413](#)
- Physical interface policers are not supported in Junos OS Release 19.3R1 for the MPC11 line card. [PR1452963](#)
- After more than 2 million multicast subscribers are activated without performing GRES or bbe-smgd restart, further multicast subscribers might be unable to log in. [PR1459340](#)
- The following CDA error message is observed: **LkupAsicClient: Index Dmem block read failed, PFE:0.0**. [PR1459665](#)
- The CFM REMOTE MEP does not come up after configuration or if the MEP remains in the **Start** state. [PR1460555](#)
- Need to add the Backport jemalloc profiling CLI support to all Junos OS releases where jemalloc is present. [PR1463368](#)
- In DNS filtering when DNS requests are sent from the server and implicit filters as well as routes to the service PIC are configured, it causes the DNS packets to loop. [PR1468398](#)
- 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](#)
- For the MPC10E line card, the IS-IS and micro-BFD sessions do not come up during baseline. [PR1474146](#)
- Expected number of 512,000 MAC entries are not relearned 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 and MPLS configurations: **[Error] L2alm : l2alm\_mac\_process\_hal\_delete\_msg:667 Ignoring MAC delete with ifl index 355, fwd\_entry has 7888**. [PR1475785](#)
- A 64-bit cMGD should be used if cMGD is running on a 64-bit OS to avoid random issues. [PR1481335](#)

- Invalid packets are dropped by dut with TCC encapsulation configuration as intended, but the statistics counters are incremented. [PR1481698](#)
- The following 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:0xffffffff indirect:333988 hwInstall:1#012.** [PR1486158](#)
- Next-hop learning command is enabled by default in the MPC10 and MPC11 line cards irrespective of the command configuration. [PR1489121](#)
- Login or logout of high scale (around 1 million bearers) causes some sessions not to re-login. [PR1489665](#)
- Need to support upgrading of the PSM firmware on the MX2000 line of devices. [PR1489939](#)
- On the MPC10 line card, AFT crash is seen 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 following error message is observed: **unable to set line-side lane config (err 30).** [PR1492162](#)
- The delta PSM firmware upgrade status is incorrectly displayed. [PR1493045](#)
- On the MX2020 router, the AER image for non-correctable or correctable PCI error is needed. [PR1493065](#)
- Component sensor does not export data under components CB0 or CB1 in the expected time. [PR1493579](#)
- Backup Routing Engine reboots because of power cycle or failure when the offline and online operations are performed on CB1. [PR1497592](#)
- The MPC11 line card is not supported in Junos OS Release 19.4R1. [PR1503605](#)
- The WAN-PHY interface continuously flaps with the default hold-time down of value **0**. [PR1508794](#)
- For EVPN-VXLAN feature verification, the **set chassis loopback-dynamic-tunnel** command is used. [PR1509690](#)
- On the MPC11 line card, dfw crash is seen after removing and restoring configurations on the backup Routing Engine. [PR1512770](#)
- Sometimes external 1 pps cTE is slightly above Class B requirement of the ITU-T G.8273.2 specification. [PR1514066](#)
- On the MX960 router, expected traffic is not received with multicast and PIM scaling configurations. [PR1514646](#)
- The NGMPC2 process generates the core file at **bv\_entry\_active\_here::bv\_vector\_op::gmph\_reevaluate\_group::gmph\_destroy\_client\_group.** [PR1537846](#)
- On the MX480 routers, in an EVPN-VLAN scenario, the **set routing-instances protocols evpn mac-table-aging-time 30** statement does not work. [PR1543238](#)
- Even though enhanced-ip is active, the following alarm is observed during ISSU: **RE0 network-service mode mismatch between configuration and kernel setting.** [PR1546002](#)



- The LACP state is in the **Down** state after enabling and disabling the exclude protocol LACP under **Set security**. [PR1331412](#)
- Disabled interfaces might still transmit power after the device reboots. [PR1487554](#)
- In the output of the **show interface** command, the **smart-sfp-present** leaf is missed. [PR1492551](#)
- Traffic loss might be seen if the routing-instance is deactivated and then re-activated quickly. [PR1498087](#)
- Set of Info level cron logs is displayed from FPC every 1 minute. [PR1527266](#)
- CFM do not consider the 8021AD configuration for the rewrite and classification tables. [PR1527303](#)
- MACSEC PIC stays offline in new primary after ISSU in GNF alone. [PR1534225](#)
- On the MX2020 router, the next hops are less than a total of nhdb 4MPOST GRES. [PR1539305](#)
- On the MX480 routers, COS shaping is not adjusted as per the ANCP actual down stream rate. [PR1544713](#)
- Commit error is introduced during **deactivate chassis synchronization source** and **smc-transmit** are all configured. [PR1549051](#)
- IGMP joins are more than the expected value while verifying the IGMP snooping membership in the CE router. [PR1560588](#)
- Some BFD sessions get stuck in the **Down** or **Init** state after an iterative operations triggers on DUT. [PR1560772](#)
- On the MX2010 or MX2020 routers, the following error message might be observed after switchover with GRES/NSR: **CHASSISD\_IPC\_FLUSH\_ERROR**. [PR1565223](#)
- On the MX480 routers, traffic loss is observed with a scale of 4000 tunnels 800 vrf test. [PR1568414](#)
- The mspmand process might crash if the packet flow-control issue occurs on MS-MPC/MS-MIC. [PR1569894](#)
- CFP unplugged message is not logged in Junos OS Release 17.3 and later. [PR1573209](#)
- The rpd process on the transit node might crash when MPLS traceroute on the ingress node is performed. [PR1573517](#)
- From the regress user shell prompt, vhlclient access does not display the following error message: **rcmd: socket: Operation not permitted**. [PR1574240](#)
- PIM rib-group fails to add in VR. [PR1574497](#)
- On the MX150 routers, the interface might take a long time to power down while rebooting, powering-off, halting, or upgrading. [PR1575328](#)
- FPC CPU utilization gets stuck at 100 percent during the longevity case. [PR1575355](#)
- The **show services service-sets statistics syslog** command returns an error when the service-set does not have a syslog configuration: **usp\_ipc\_client\_recv\_ 1237: ipc\_pipe\_read fails! error:No error: 0(0), tries:.** [PR1576044](#)

- On the MX10016 routers, when Fan Tray 1 fan fails the alarm is cleared, the **Fan/Blower OK** SNMP traps are generated for the Fan Tray 0 [Fan 31 - 41] and Fan Tray 1 [Fan 11 - 41]. [PR1576521](#)
- In the NAT64 scenario during session creation, the IPv6 atomic fragments are not processed correctly. [PR1581348](#)
- MS-MIC or MS-MPC based jflow (flow-sampling) on the logical systems is not supported. [PR1585824](#)

## High Availability (HA) and Resiliency

- Unexpected XML structure change with the **show system switchover** command is observed. [PR1158986](#)
- Performing GRES with the interface em0 (or fxp0) disabled on the primary Routing Engine; when you enable the interface on the new backup Routing Engine, you might not be able to access the network. [PR1372087](#)
- During ZPL ISSU, traffic loss is observed with the IGP or BGP protocol session. [PR1487144](#)

## Infrastructure

- The HSRPv2 IPv6 packets might get dropped if IGMP-snooping is enabled. [PR1232403](#)
- The following error message is seen during FTP: **ftpd[14105]: bl\_init: connect failed for /var/run/blacklistd.sock(No such file or directory)**. [PR1315605](#)
- 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](#)
- Memory corruption of any binary in **/usr/bin/** or **/usr/sbin/** can be triggered by the execution of the binary when a recovery snapshot is being copied to the OAM volume. [PR1563647](#)

## Interfaces and Chassis

- The **cfmd** process might continuously crash after the upgrade. [PR1281073](#)
- 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](#)
- 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](#)
- MPLS VPN label can point to the discarded next hop after a Routing Engine switchover without NSR if the egress interface is pp0. [PR1488302](#)
- The **show interface x extensive** command might not be accurate. [PR1505100](#)
- LB fails to MIP on VT with a default md. [PR1516583](#)

- After DUT with MPC10 or MPC11 line card takes over as vrrp primary role, the logical interface undergoes 100 seconds of traffic loss. [PR1519374](#)
- 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 **input errors counter** command on the **monitor interface** command does not work. [PR1561065](#)

## Layer 2 Ethernet Services

- The DHCP decline packets are not forwarded to the DHCP server when forward-only is set within dhcp-reply. [PR1429456](#)
- the OSPF and OSPF3 adjacency uptime is more than expected after the NSSU upgrade and outage is higher than the expected. [PR1551925](#)

## MPLS

- Aggressive switchovers due to MBB or CSPF computations causes traffic loss on all branches of the tree even if a single branch fails to come up due to remerge detection on the transit router. [PR1487916](#)
- The GRES or NSR Routing Engine switchovers followed by restart routing on the primary Routing Engine does not honor the remerge behavior. [PR1489168](#)
- Extended-admin-groups on links are shown as SRLG attribute in TED. [PR1575060](#)

## Platform and Infrastructure

- The Packet Forwarding Engine might produce error messages during interface deletions in configurations with IRB interfaces. [PR1054798](#)
- The following error message is observed during ISSU from 19.1-20190325.0 to 19.3I-20190324\_dev\_common.0.1957: **Async XTXN Error PPE/Context 9/13 @ PC 0x6f77: sampling\_li\_launch\_nh** [PR1426438](#)
- For the bridge-domains configured under an EVPN instance, the ARP suppression is enabled by default. This enables the EVPN to proxy the ARP and reduces the flooding of ARP in the EVPN networks. As a result, storm-control does not effect the ARP packets on the ports under such bridge-domain. [PR1438326](#)
- The npc process generates the core file at **trinity\_rt\_iff\_attach,pfe\_ifl\_family\_attach,ifrt\_ifl\_family\_adder,ifrt\_ifl\_family\_add\_vector,ifrt\_command\_handler.** [PR1461892](#)
- The cosmetic error messages of NTP time synchronization might be observed during device booting. [PR1463622](#)

- 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](#)
- Traffic loss might be observed after ISSU. [PR1493723](#)
- Upgrading satellite devices might lead to some SDs in the **SyncWait** state. [PR1556850](#)
- On the MX480 router, during the verification of GRES and NSR functionality with VXLAN feature, the convergence is not as expected L2-DOMAIN-TO-L3VXLAN. [PR1520626](#)
- The vmxt\_lnx process generates core file at **KtreeSpace::FourWayLeftAttachedNode::getNextDirty Trinity\_Ktree::walkSubTree Trinity\_Ktree::walkSubTree**. [PR1525594](#)
- IPv6 VRRP sessions are not established when Duplicate Address Detection (DAD) is enabled. [PR1534835](#)
- Upgrading satellite devices might lead to some SDs in the **SyncWait** state. [PR1556850](#)
- Monitor traffic interface fxp0 resets the last flapped time for the interface. [PR1564323](#)
- The FPC process might crash when the next-hop memory of ASIC is exhausted in the EVPN-VXLAN scenario. [PR1571439](#)

## Routing Policy and Firewall Filters

- The routing policy actions fail to configure neighbor-sets and tag-sets. [PR1491795](#)

## Routing Protocols

- While interoperating with other vendors in a draft-rosen multicast VPN, by default Junos OS attaches a route target to the multicast distribution tree (MDT), subsequent address family identifier (SAFI), and network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching the route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities are prevented from propagating of the route-target fil. [PR993870](#)
- Certain BGP traceoption flags (for example, open, update, and keepalive) might result in trace logging of debugging messages that do not fall within the specified traceoption category, which results in some unwanted BGP debug messages being logged to the BGP traceoption file. [PR1252294](#)
- LDP OSPFs are in the **Synchronization** state because the IGP interface is down with ldp-synchronization enabled for OSPF. [PR1256434](#)
- In rare cases, RIP replication might fail as a result of performing NSR Routing Engine switchovers when the system is not NSR ready. [PR1310149](#)

- The **show version detail** command triggers the following severity error logs: **mcsnoopd: INFO: krt mode is 1 "JUNOS SYNC private vectors set"**. [PR1315429](#)
- SCP command with routing option (-JU) is not supported. [PR1364825](#)
- On the MX2010 Series routers, the BFD session on the IS-IS step up flaps during the ISSU - FRU upgrade stage. [PR1453705](#)
- Even when protocols mpls traffic-engineering bgp-igp command is configured, the UDP tunnel routes are not added to inet.0. The UDP tunnel routes are added only to inet.3 table whether the command is configured or not. [PR1457426](#)
- BGP graceful restart might have some traffic loss when sharding is enabled. [PR1475773](#)
- Some PIM join or prune packets might not be processed in the first attempt in the scaling scenario where the PIM routers establish neighborhood and immediately join the multicast group. [PR1500125](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- BFD with authentication for BGP flaps after GRES or NSR switchover on the NG-RE and SCBE2 setup. [PR1522261](#)
- The virtual-router option is not supported under a routing-instance in a lean RPD image. [PR1494029](#)
- Dynamic tunnels are still up after deactivating the BGP nexthop type UDP policy. [PR1579225](#)

## Services Applications

- All the unreachable destinations are not kept in the **Locked out** state post GRES. [PR1541271](#)
- The **Tunnel-Assignment-Id** string is not present while checking the packets from coming in for the attributes. [PR1543628](#)

## Subscriber Access Management

- BBE-SMGD configures in-correct vbf\_accurate\_accounting\_bits to the Packet Forwarding Engine. [PR1515899](#)
- Subscriber might get stuck in the **Terminating** state if the **Access-Challenge** packet is received from the RADIUS server during the subscriber authentication. [PR1583090](#)

## User Interface and Configuration

- A 64-bit cMGD must be used if cMGD runs on a 64-bit OS to avoid random issues. [PR1481335](#)
- The **port\_speed** configuration details not present in the **picd** configuration for ports et-0/0/128 and et-0/0/129. [PR1510486](#)

## VPNs

- In an MVPN environment with SPT-only option, if the source or receiver is connected directly to c-rp PE and the MVPN data packets arrive at the c-rp PE before its transition to SPT, the MVPN data packets might be dropped. [PR1223434](#)
- The output value of the `show mvpn c-multicast inet source-pe | display xml` command is not proper. [PR1509948](#)
- Interface statistics do not match for the Mroute VPN-B. [PR1517039](#)
- The PIM (S,G) join state might stay forever when there are no MC receivers and source is inactive. [PR1536903](#)

### SEE ALSO

[What's New | 102](#)

[What's Changed | 129](#)

[Known Limitations | 136](#)

[Resolved Issues | 149](#)

[Documentation Updates | 185](#)

[Migration, Upgrade, and Downgrade Instructions | 186](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 150](#)
- [Resolved Issues: 20.2R2 | 160](#)
- [Resolved Issues: 20.2R1 | 169](#)

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

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

## Resolved Issues: 20.2R3

### *Class of Service (CoS)*

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

### *EVPN*

- With dynamic list next hop configured, a forwarding problem occurs after graceful switchover. [PR1513759](#)
- **no-arp-suppression** is required for MAC learning across the EVPN domain on the static VTEP. [PR1517591](#)
- The BUM traffic might get dropped in the EVPN-VXLAN setup. [PR1525888](#)
- The route table shows additional paths for the same EVPN or VXLAN Type 5 destination after upgrading from Junos OS Release 18.4R2-S3 to Junos OS Release 19.4R1-S2. [PR1534021](#)
- All the ARP reply packets toward some address are flooded across the entire fabric. [PR1535515](#)
- Rpd memory leak might occur when the EVPN configuration is changed. [PR1540788](#)
- The l2ald process might generate the core file after changing the EVPN or VXLAN configuration. [PR1541904](#)
- The rpd process might crash after adding route-target on a dual-Routing Engine system under the EVPN multihoming scenario. [PR1546992](#)
- VLAN ID information is missed while installing the EVPN route from the BGP Type 2 Route after modifying a routing-instance from instance-type EVPN to instance-type virtual-switch. [PR1547275](#)
- The ARP replies from the EVPN CE device might get dropped incorrectly if the EVPN routes are resolved through the MPLS-over-UDP tunnels. [PR1563802](#)

### *Forwarding and Sampling*

- The srrd process might crash in a high route churns scenario or if the process flaps. [PR1517646](#)
- The commit might fail if a filter enabled with enhanced-mode to et- interface is configured. [PR1524836](#)
- The l2ald process might crash when a device configuration flaps frequently. [PR1529706](#)
- MAC learning issue might occur when EVPN-VXLAN is enabled. [PR1546631](#)
- All traffic are dropped on the aggregated Ethernet interface bundle without the VLAN configuration if the bandwidth-percent policer is configured. [PR1547184](#)
- The l2ald process might crash due to next-hop issue in the EVPN-MPLS. [PR1548124](#)
- In the VXLAN scenario, the locally originated packets have UDP source port 0. [PR1571970](#)

## General Routing

- The **max-drop-flows** statement is not available. [PR1375466](#)
- The MPC2E-NG or MPC3E-NG line card with specific MIC might crash after a high rate of interface flaps. [PR1463859](#)
- The following error message is observed after GRES: **[user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, agglfifName:ps1.0 memberIfifName:lt-3/0/0.32767]**. [PR1466531](#)
- The following line card errors are seen: **HALP-trinity\_nh\_dynamic\_mcast\_add\_irb\_topo:3520 snooping-error: invalid IRB topo/ IRB ifl zero in l2 nh 40495 add IRB**. [PR1472222](#)
- Dynamic SR-TE tunnels do not get automatically recreated at the new primary Routing Engine after the Routing Engine switchover. [PR1474397](#)
- 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 vmcore process crashes sometimes along with the mspmand process on MS-MPC or MS-MIC if large-scale traffic flows are processed. [PR1482400](#)
- SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. [PR1484322](#)
- False positive TSensor errors are reported on vjunos0. [PR1508580](#)
- Not able to forward traffic to VCP FPC after the MX Virtual Chassis reboots, FPC reboots, or adding VCP link. [PR1514583](#)
- On the MX960 routers, the **show interfaces redundancy RLTO** statement shows current status as primary down as FPC is still in the **Ready** state after RLT failover (restart FPC). [PR1518543](#)
- During an upgrade, vSRX3.0 displays the following incorrect license warnings when utilizing licensable features even if the license is present on the device: **requires 'idp-sig' license**. [PR1519672](#)
- The BFD session status remains down at the non-anchor FPC even though the BFD session is up after anchor the FPC reboots. [PR1523537](#)
- Problem with static VLAN deletion with active subscribers and the FPC might be stuck at the **Ready** state during restart. [PR1525036](#)
- The following error message is observed during GRES if an IRB interface is configured without a profile: **RPD\_DYN\_CFG\_GET\_PROF\_NAME\_FAILED**. [PR1526481](#)
- The transit PTP packet might be modified unexpectedly while passing through MPC2E-NG, MPC3E-NG, and MPC5E line cards. [PR1527612](#)
- The **speed** command cannot be configured under the interface hierarchy on an extended port when the MX204 or MX10003 router works as an aggregation device. [PR1529028](#)
- The SFP-LX or SFP-SX optics on MIC-3D-20GE-SFP-E/EH might show as unsupported after ISSU. [PR1529844](#)



- On the MX204 and MX10003 routers, PEM0 always shows as **Absent** or **Empty** even if PEM0 is present. [PR1531190](#)
- Commit might fail after Routing Engine switchovers. [PR1531415](#)
- On the MX150 routers, configuring the **no-flow-control** command under `gether-options` does not work. [PR1531983](#)
- Wavelength unlocked alarm is set as **On** while using the SFP+-10G-T-DWDM-ZR optics. [PR1532593](#)
- The interface with the **pic-mode 10GE** configuration might not come up if upgraded to Junos OS Release 18.4R3-S4 or later. [PR1534281](#)
- Some routes might get incorrectly programmed in the forwarding table in the kernel, which is no longer present in rpd. [PR1534455](#)
- Snmp mib walk for `jnxSubscriber` OIDs returns a general error. [PR1535754](#)
- All SFBs might go offline due to fabric failure and fabric self-ping probes performing the **disable-pfe** action. [PR1535787](#)
- Enhancements are needed for debugging `I2ald`. [PR1536530](#)
- The chassisd memory leak might cause traffic loss. [PR1537194](#)
- The following error message might be observed when the JAM packages for the MX204, MX10003, and MX10008 are installed: **JAM: Plugin installed for summit\_xxx PIC**. [PR1537389](#)
- Version-alias gets missed for the subscribers that are configured with the dynamic profiles after ISSU. [PR1537512](#)
- Deactivating or activating PTP or synchronized Ethernet in the upstream router causes the 100GbE links on the LC2103 to flap. [PR1538122](#)
- On the AFT based FPCs (MPC10 and MPC11 line cards), the **show jnh exceptions inst** command of the Packet Forwarding Engine might cause the FPC process to crash. [PR1538138](#)
- Traffic drop might be seen while executing the **request system reboot** command. [PR1538252](#)
- After configuring the **global system name-server** configuration, commit should fail but instead the commit is successful. [PR1538514](#)
- Upon receiving of a specific BGP FlowSpec message, network traffic might be disrupted. [PR1539109](#)
- The accounting interim-updates for subscriber does not work after GRES and subsequent reboot of FPCs in the node-slicing setup. [PR1539474](#)
- The rpd memory leak might be observed on the backup Routing Engine due to the flapping of the link. [PR1539601](#)
- The `mspan` process leaks memory in relation to the MX Series telemetry reporting the following error message: **RLIMIT\_DATA exceed**. [PR1540538](#)
- With hold time configuration, the `ge` interfaces remain down on reboot. [PR1541382](#)

- Subscriber might not come up on some dynamic VLAN ranges in a subscriber management environment. [PR1541796](#)
- The KRT queue might get stuck after the Routing Engine switchovers. [PR1542280](#)
- Port mirroring with the **maximum-packet-length** configuration does not work over the GRE interface. [PR1542500](#)
- The license errors might get returned on the backup Routing Engine while trying to commit the configuration. [PR1543037](#)
- The mspmand process might generate the core file on activating or deactivating the interface. [PR1544794](#)
- Traffic loss might be observed when the Switch Fabric Board 3 and MPC8E 3D combination is used in the MX2010 and MX2020 routers. [PR1544953](#)
- Continuous rpd errors might be seen and new routes fails to be programmed by the rpd process. [PR1545463](#)
- Backup Routing Engine vmcore might be seen due to the absence of the next-hop acknowledgement infra. [PR1547164](#)
- In the syslog output, the **sylog-local-tag** name is truncated as **SYSLOG\_SF** when the **sylog-local-tag** name is configured as **SYSLOG\_SFW**. [PR1547505](#)
- The **verbose** command unexpectedly becomes hidden after Junos OS Release 16.1 for **set system export-format json**. [PR1547693](#)
- The **SENSOR APP DWORD** leak is observed during the period of churn for routes bound to the sensor group. [PR1547698](#)
- Multicast traffic drop might be seen after ISSU. [PR1548196](#)
- The adapted sample rate might get reset to the configured sample rate without changing the sampling rate information in sFlow datagrams after enabling sFlow technology on a new interface. [PR1550603](#)
- The rpd crash might be seen when the BGP service route is resolved over the color-only SR-TE policy. [PR1550736](#)
- The PPPoE subscribers might fail to login. [PR1551207](#)
- The **LCM Peer Absent** message might be seen. [PR1551760](#)
- The fabric errors are observed and the FPC processes might get offline with the SCBE3, MPC3E-NG, or MPC3E and MPC7 or MPC10 line card in the **increased-bandwidth fabric** mode. [PR1553641](#)
- Configuring HFRR (link-protection) on an interface might cause rpd to crash. [PR1555866](#)
- The following message is not generated on the MPC11E line card due to no power: **Chassisd SNMP trap Fru Offline**. [PR1556090](#)
- On the MX150 routers, the following continuous license error is observed:  
[licinfra\_set\_usage\_nextgen\_async:1733] Invalid input parameters. [PR1559361](#)

- The **request system software validate** command might corrupt installation of the **junos-openconfig** package. [PR1560234](#)
- The rpd crash might be observed during processing a huge amount of PIM prune messages. [PR1561984](#)
- MX platforms with MX-SCBE3 might reboot continuously. [PR1564539](#)
- PPPoE service-name-tables does not correctly count active sessions that matches agent-specifier aci/ari used for delay. [PR1565258](#)
- On the MX150 routers, the **request system software add** command is disabled in Junos OS Release 19.4R3-S1, 20.1R2, and 20.4R1. [PR1568273](#)
- Family IPv6 does not come up for Layer 2 TP subscriber when additional attributes are not passed in the Framed-IPv6-Route VSA. [PR1526934](#)
- DHCP discover packet might be dropped if the DHCP inform packet is received first. [PR1542400](#)
- The **show dynamic-profile session client-id** command displays only one IPv6 framed-route information. [PR1555476](#)
- Slow response might be observed when the **show | compare** or **commit check** action in a large-scale configuration environment is committed. [PR1500988](#)
- Transit IPv4 traffic forwarding over BGP SR-TE might not work. [PR1505592](#)
- The **No response from the other routing engine for the last 2 seconds** error triggers the SNMP trap generated **Fru Offline** messages. [PR1524390](#)
- Multiple FRUs disconnection alarms might be displayed post the firmware upgrade. [PR1529710](#)
- The following error message for port might be observed: **FAILED(-1) read of SFP eeprom**. [PR1529939](#)
- The unilists are incorrectly formed and the list of forwarded next hops are not resolved properly if the ECMP is set to 128. [PR1530803](#)
- BGP SR-TE IPv6 routes might get hidden after the chassisd restarts. [PR1534511](#)
- Multiple vmxt processes might generate core files. [PR1534641](#)
- Snmp mib walk for jnxSubscriber OIDs returns a general error. [PR1535754](#)
- The kmd process might crash when the interface flaps. [PR1544800](#)
- The l2ald process might crash due to next-hop issue in the EVPN-MPLS. [PR1548124](#)
- The Broadcom chip FPC might crash during the system booting. [PR1545455](#)
- The performance of the Packet Forwarding Engine process on the MX204 routers might be degraded after Junos OS Release 19.3R1. [PR1545989](#)
- Unexpected log messages appears related to the Neighbor Solicitation (NS) messages with multicast as source address. [PR1546501](#)
- The nsd daemon might crash after configuring the inline NAT in the USF mode. [PR1547647](#)
- SR-TE might stay in the **Up** state when the routes are deleted through policy. [PR1547933](#)

- Validation of the OCSP certificate might not go through in case of certain CA servers. [PR1548268](#)
- The l2alm processes high CPU utilization might be observed in the EVPN-VXLAN environment. [PR1551025](#)
- The following error messages are observed: **Disable-pfe with intermittent ipc\_pipe\_get\_packet(): packet\_get() failed error message** and **CM\_CMERROR\_FABRIC\_SELFPING failure**. [PR1554209](#)
- During ISSU, BNG losses subscriber sessions without sending the Session Stop message but stay in authd. [PR1554539](#)
- The framed route installed for a demux Interface has no MAC address. [PR1556980](#)
- ISSU are aborted and the chassisd process generates core file on the backup Routing Engine during the Junos OS upgrade to version Junos OS Release 20.2R2-S1. [PR1557413](#)
- Packets corruption on 100G or 40G interface are configured with protocol PTP. [PR1557758](#)
- Need to allow the tunnel interface as the peer-address for ALQ. [PR1567735](#)
- On the MX204 routers, FPC might display high CPU utilization because of the JGCI background thread that runs for a long period. [PR1567797](#)
- Core files are generated at **export\_svc\_set\_nat\_idl@nsd\_malloc** while verifying the no-translation with destination-nat. [PR1568997](#)
- The RPD process might crash while using BFD API to bring up the BFD sessions. [PR1569040](#)
- The agent sensor **\_\_default\_fabric\_sensor\_\_** are partly applied to some FPCs, which causes zero payload issue **AGENTD received empty payload for pfe sensor \_\_default\_fabric\_sensor\_\_**. [PR1569167](#)
- The MPLS traffic passed through the back-to-back PE topology might match the wrong CoS queue. [PR1569715](#)
- OAM might not work as expected after FPC reboots or flaps. [PR1569790](#)
- The following log message might be observed: **/tmp//mpci\_info: No such file or directory :error[1]**. [PR1570135](#)
- On the MX960 routers, the **Require a Fan Tray upgrade** alarm is raised when the top Fan Tray 0 is removed, even though the enhanced Fan Tray is already used. [PR1572778](#)
- Fabric errors are observed and FPC processes might get offline when the MPC3-NG/MPC3E/SRX5K-IOC2 line cards are installed along with the MPC7/MPC10/SRX5K-IOC04 and SCBE3/SCB4 line cards operating in an increased-bandwidth fabric mode. [PR1573360](#)
- Slow FPC heap memory leak might be triggered by flapping the subscribers terminated over multiple pseudowires. [PR1574383](#)
- On the EA-based cards igmp group membership is displayed incorrectly. [PR1575031](#)
- The LLDP neighbor information displays hex string instead of chassis ID when subtype 1 is used. [PR1576721](#)

### **Infrastructure**

- The output of the **show interfaces extensive** command might display 0 temporarily during a race condition when SNMP query for JnxCos is issued. [PR1533314](#)

### **Interfaces and Chassis**

- The configuration might not be applied after deleting all existing logical interfaces and adding a new logical interface for an IFD in a single commit. [PR1534787](#)
- Inline Y.1731 SLM or DM does not work in **enhanced-cfm-mode** for the EVPN UP MEP scenario. [PR1537381](#)
- The following error message might occur after commit for configuration under interface hierarchy: should have at least one member link on a different FPC. [PR1539719](#)
- After VRRP failover, the VRRP backup router keeps receiving traffic for about 2 minutes. [PR1546635](#)
- The following commit error is observed while trying to delete unit 1 logical system interfaces: **ae2.1: Only unit 0 is valid for this encapsulation.** [PR1547853](#)
- An IRB interface that has large unit value over 32767 cannot be an active group for the inheriting VRRP. [PR1550993](#)
- The VCP port is marked as administratively down on the wrong MX-VC member. [PR1552588](#)
- The dcd process might leak memory on pushing the configuration to the ephemeral database. [PR1553148](#)
- Junos device might send VRRP advertisement packets in the VRRP **Init** or **Idle** state before **startup-silent-period timer expiry** on the VRRP primary device with NSR disabled after GRES. [PR1558560](#)
- MAC address entry issue might be observed after the MC-LAG interface. failover/failback [PR1562535](#)

### **Layer 2 Ethernet Services**

- The `jnxJdhcpLocalServerMacAddress (.1.3.6.1.4.1.2636.3.61.61.1.4.3)` returns incorrect format of the MAC address. [PR1565540](#)
- DHCP packet drop might be seen when the DHCP relay is configured on a leaf device. [PR1554992](#)
- The Option 82 information is incorrectly cleared by the DHCP Relay agent. [PR1568344](#)

### **MPLS**

- The rpd scheduler might slip after the link flaps. [PR1516657](#)
- The rpd process might crash when the LDP route with indirect next hop is deleted on the aggregated Ethernet interface. [PR1538124](#)
- If link-protection is enabled for an externally provisioned LSP, any commit for the first time after provisioning causes a break (MBB) even if the configuration is not related to the LSP. [PR1546824](#)
- A new LSP might not be up even if bypass LSP is up and **setup-protection** is configured. [PR1555774](#)

### **Network Management and Monitoring**

- Commit error occurs while deleting the routing instance when the SNMP trap-group also have the same routing instance referred. [PR1555563](#)

### **Platform and Infrastructure**

- The state of the flow detection configuration might not be displayed properly if DDoS-SCFD is configured globally. [PR1519887](#)
- An internal timer on the backup Routing Engine might cause an ARP storm upon GRES switchover on the new primary (old backup) Routing Engine. [PR1547583](#)
- The following major error message might cause the Packet Forwarding Engine(s) to disable: **XQ\_CMERROR\_SCHED\_L3\_PERR\_ERR**. [PR1538960](#)
- The VXLAN encapsulation over IPv6 underlay might not work. [PR1532144](#)
- PE-CE OAM CFM might have issues in the aggregated Ethernet interface. [PR1501656](#)
- Flow programming issue for It- interface in the Packet Forwarding Engine level is observed. [PR1525188](#)
- The following error message is observed when alarms after interface reset: **7836 ifl 567 chan\_index 8 NOENT & jnh\_ifl\_topo\_handler\_pfe(13015): ifl=567 err=1 updating channel table nexthop**. [PR1525824](#)
- PPE errors or traps might be observed in the Layer 2 flooding scenarios. [PR1533767](#)
- The FPC process might crash when the next-hop memory of ASIC is exhausted in the EVPN-MPLS scenario. [PR1533857](#)
- The npc process generates the core file in **igmp\_process\_wakeup\_events,igmp\_pfe\_thread,thread\_detach\_tty**. [PR1534542](#)
- Subscribers do not come up on VPLS in the PS interface. [PR1536043](#)
- Packet loss might be observed when the RFC2544 egress reflector session is configured on the non-zero Packet Forwarding Ethernet interface. [PR1538417](#)
- The rmopd process memory leak might be seen if the TWAMP client is configured. [PR1541808](#)
- FPC might crash when the underlying Layer 2 interface for ARP over IRB interface is changed from the physical interface to the LSI interface. [PR1542211](#)
- The RP expired timer on the backup Routing Engine is not the same as the primary Routing Engine if the aging-timer is configured. [PR1544398](#)
- The kernel might crash if GRES is performed on either new iteration or after swapping the Routing Engine and restoring the HA configuration. [PR1549656](#)
- The BGP session replication might fail to start after the session crashes on a backup Routing Engine. [PR1552603](#)
- Traffic is not forwarded over IRB to a Layer 2 circuit on the It interfaces. [PR1554908](#)
- The IPv4 EXP rewrite might not work properly when inet6-vpn is enabled. [PR1559018](#)

- The BUM frame might be duplicated on an aggregate device if the extended-port on the satellite device is an aggregated Ethernet interface. [PR1560788](#)
- The DHCPv4 request packets might be wrongly dropped when DDoS attack occurs. [PR1562474](#)
- The **enforce-strict-scale-limit-license** configuration enforces subscriber license incorrectly in the ESSM subscriber scenario. [PR1563975](#)

### ***Routing Policy and Firewall Filters***

- The policy configuration might be mismatched between the rpd and mgd processes when deactivating the policy-options prefix-list in the configuration sequence. [PR1523891](#)
- Generated route goes to the **Hidden** state when the protect core command is enabled. [PR1562867](#)
- Global variable policy\_db\_type is not set to the correct value on failure. [PR1561931](#)

### ***Routing Protocols***

- The BFD session might get stuck in the **Init** or **Down** state after the BFD session flaps. [PR1474521](#)
- With BGP rib-sharding enabled, the RPD memory exhaustion might be observed. [PR1546347](#)
- Traffic loss might be seen in the next-hop-based dynamic tunnels of the Layer 3 VPN scenario after changing the dynamic-tunnel preference. [PR1542123](#)
- Traffic loss might occur during VRF route resolution over indirect next hop. [PR1525363](#)
- Traffic might be silently discarded when the BGP route gets deleted, which is part of multipath. [PR1514966](#)
- 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 rpd might crash with BGP RPKI enabled in a race condition. [PR1487486](#)
- The ppmd process generates the core file after MS-MPC restarts. [PR1490918](#)
- The BGP session with VRRP virtual address might not come up after the session flaps. [PR1523075](#)
- The VRF label is not assigned at ASBR when the inter AS is implemented. [PR1523896](#)
- The IS-IS LSP database synchronization issue might be seen while using the flood-group feature. [PR1526447](#)
- Transit labels for Layer 3 VPN routes are pushed momentarily to the MPLS.0 table. [PR1532414](#)
- Configuring the next hop and then rejecting it on a route policy for the same route might cause the rpd process to crash. [PR1538491](#)
- After the peer is moved out of the protection group, the path protection is not removed from the PE device. Multipath route is still present. [PR1538956](#)
- The rpd process generates the core file at **gp\_rtargert\_tsi\_update,bgp\_rtargert\_flash\_rt,bgp\_rtargert\_flash**. [PR1541768](#)

- Continuous rpd crash might be observed if a static group is added to protocol PIM. [PR1542573](#)
- The metric of prefixes in intra-area-prefix LSA might be changed to 65535 when the metric of one of the OSPFv3 P2P interfaces is set to 65535. [PR1543147](#)
- The neighbor shutdown configuration of the BGP session does not effect the non-established peer. [PR1554569](#)
- The changes do not get effective when the values are set under the **static default** hierarchy. [PR1555187](#)
- Sending multicast traffic to downstream receiver on the Trio based Virtual Chassis platforms might fail. [PR1555518](#)
- Multipath information is displayed for BGP route even after disabling the interface for one path. [PR1557604](#)
- All the Layer 3 VPN route resets when a VRF is added or removed. [PR1560827](#)
- Duplicate LSP next hop is shown on inet.0, inet.3, and mpls.0 route table when OSPF Traffic-Engineering shortcuts and mpls bgp-igp-both-ribs are enabled. [PR1561207](#)
- SNMP MIB ospfv3NbrState returns a drifted value. [PR1571473](#)
- Six PE device prefixes might not be removed from RIB upon the reception of withdrawal from a BGP neighbor when RIB sharding is enabled. [PR1556271](#)
- Wrong SPF calculation might be observed for OSPF with ldp-synchronization hold-time configured after the interface flaps. [PR1561414](#)
- BGP routes might be stuck in routing table in the **Accepted DeletePending** state when the BGP peering session goes down. [PR1562090](#)
- VRF table does not get refreshed after a change made to **maximum-prefixes** in the VRF. [PR1564964](#)
- Traffic might be lost during mirror data transmit from primary ppm/bfdd. [PR1570228](#)
- SNMP MIB ospfv3NbrState returns drifted value. [PR1571473](#)
- BGP session flap might be observed after the Routing Engine switchovers when the VRRP virtual address is used as the local address for the BGP session. [PR1576959](#)

### ***Services Applications***

- Layer 2 TP subscribers might fail to establish a session on MX if the CPE is a virtual host. [PR1527343](#)
- The following error message is observed: **SPD\_CONN\_OPEN\_FAILURE: spd\_pre\_fetch\_query: unable to open connection to si-1/0/0**. [PR1550035](#)

### ***User Interface and Configuration***

- The configuration under groups stanza is not inherited properly. [PR1529989](#)
- Commit might fail after the Routing Engine switchovers. [PR1531415](#)



- The license errors might be returned on the backup Routing Engine when you try to commit the configuration. [PR1543037](#)
- The **verbose** command unexpectedly becomes hidden after Junos OS Release 16.1 for **set system export-format json**. [PR1547693](#)

#### VPNs

- MVPN multicast route entry might not be properly updated with the actual downstream interfaces list. [PR1546739](#)

### Resolved Issues: 20.2R2

#### *Application Layer Gateways (ALGs)*

- The srxpfe or mspmand process might crash if FTPS is enabled in a specific scenario. [PR1510678](#)

#### *EVPN*

- EVPN-VXLAN core isolation does not work when the system is rebooted or the routing is restarted. [PR1461795](#)
- 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 next-hop 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 the kernel. [PR1477140](#)
- 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](#)
- The rpd process might crash when auto-service-id is configured in the EVPN-VPWS scenario. [PR1530991](#)
- All the ARP reply packets towards to some address are flooded across the entire fabric. [PR1535515](#)

#### *Forwarding and Sampling*

- The DHCP subscribers might get stuck in the **Terminated** state for around 5 minutes after disabling cascade ports. [PR1505409](#)
- UTC timestamp is used in the flat-file-accounting files when a profile is configured. [PR1509467](#)
- Traffic might be dropped for not exceeding the configured bandwidth under policer. [PR1511041](#)

- The pfe process might crash while running the **show pfe FPC x** command. [PR1509114](#)
- The l2ald process generates core file at **libl2\_trigger\_flush libl2\_enqueue\_pkt libl2\_send\_keepalive**. [PR1529706](#)

### General Routing

- In some MX Series deployments running Junos OS, the following random syslog messages are observed for FPCs: **FPCx ppe\_img\_ucose\_redistribute Failed to evict needed instr to GUMEM - xxx left**. These messages might not have a service impact. These messages are addressed as INFO level messages. On a Packet Forwarding Engine, there are dedicated UMEM and shared GUMEM memory blocks. This informational message indicates some evicting events between UMEM and GUMEM and can be safely ignored. [PR1298161](#)
- The **show security group-vpn member IPsec security-associations detail | display xml** command is not in the expected format. [PR1349963](#)
- On the MX2000 router, the following error message might be observed if the MPC7 line card is offline when Routing Engine switchover occurs: **Failed to get xfchip**. [PR1388076](#)
- The rpd scheduler might slip 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 million). [PR1425515](#)
- The MPC9E line card does not get offline due to unreachable destinations in the phase 3 stage. [PR1443803](#)
- The FPC process or 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 an MVPN instance, the traffic drops on multicast receivers within the range of 0.1 to 0.9 percent. [PR1460471](#)
- On the MX960 router, the following error message might be observed: **SCHED L4NP[0] Parity errors**. [PR1464297](#)
- On the MX150 routers, the **request system halt** and **request system power-off** commands do not work as expected. [PR1468921](#)
- The syslog message reports simultaneous zone change reporting for all green, yellow, orange, red zones for one or more service PICs. [PR1475948](#)
- All PPPoE subscribers might not log in after the FPC restarts. [PR1479099](#)
- Fabric healing logic incorrectly makes all MPC line cards to go offline in the MX2000 router while the hardware fault is located on one specific MPC line card slot. [PR1482124](#)
- Traffic decreases during throughput testing. [PR1483100](#)
- Any change in the nested groups might not be detected on commit and does not take effect. [PR1484801](#)
- XML is not properly formatted. [PR1488036](#)

- Prolonged flow control might occur with MS-MPC or MS-MIC. [PR1489942](#)
- The following error message is observed on the MPC line card in the manual mode:  
`clksync_as_evaluate_sync_ref: 362 - Failed to configure clk.` [PR1490138](#)
- The MX10003 RCB always detects the fire temperature and shuts down in a short time after downgrade. [PR1492121](#)
- The MPC10 or MPC11 line card might crash if the interface is configured with the firewall filter referencing shared-bandwidth policer. [PR1493084](#)
- VPLS flood next-hop might not get programmed correctly. [PR1495925](#)
- B4 might not be able to establish the softwire with AFTR. [PR1496211](#)
- Heap memory leak might be seen on the MPC10 and MPC11 line cards. [PR1499631](#)
- Some of the virtual services might not come up after GRES or rpd restart. [PR1499655](#)
- After disabling and enabling the ams0 interfaces, the NAT sessions do not get synchronized back to the current standby SDG. [PR1500147](#)
- Unexpected behavior during the `show | display inheritance` command is observed when the foreground is deactivated. [PR1500569](#)
- The `show services alg conversations` and `show services alg sip-globals` commands are not supported in the USF mode. [PR1501051](#)
- VPN traffic gets silently discarded in a cornered Layer 3 VPN scenario. [PR1501935](#)
- The chassisd process might become nonresponsive. [PR1502118](#)
- The packets from a non-existing 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 is added as part of `jnxLicenseInstallTable: jnxLicenseStartDate jnxLicenseEndDate`. [PR1503790](#)
- The gNMI stream does not follow the frequency on the subscription from the collector. [PR1504733](#)
- The rpd process might crash in case of a network churn when the telemetry streaming is in progress. [PR1505425](#)
- 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](#)
- `GnmiJuniperTelemetryHeader` incompatibility is introduced in Junos OS Release 19.3. [PR1507999](#)
- The heap memory utilization might increase after extensive subscriber login or logout. [PR1508291](#)

- Outbound SSH connection flap or memory leak issues is observed during push configuration to the ephemeral database with a high rate. [PR1508324](#)
- 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 switch 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 messages 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](#)
- 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](#)
- The MACsec session might fail to establish if the 256-bit cipher suite is configured for MACsec connectivity association assigned to a logical interface. [PR1514680](#)
- On the MX2010 and MX2020 routers, the SPMB CPU is elevated when an SFB3 is installed. [PR1516287](#)
- Active sensor check fails while checking the **show agent sensors|display xml** command. [PR1516290](#)
- Used-Service-Unit of the CCR-U has Output-Bytes counter zero. [PR1516728](#)
- 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 or 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 the accepted protocol configured as PPPoE. [PR1523902](#)
- 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](#)
- Commit is successful while deactivating CB0 and CB1 interfaces with a running GNF. [PR1524766](#)
- According to the OC data model, the **openconfig-alarms.yang** subscription path must be used as a system, alarms, or alarm. [PR1525180](#)

- Addition and removal of an aggregated Ethernet interface member link might cause the PPPoE subscriber session and traffic to drop. [PR1525585](#)
- WAG control route prefix length is observed. [PR1526666](#)
- Commit error messages comes twice while validating the **physical-cores** statement. [PR1527322](#)
- The cpcdd process might generate the core file after upgrading to Junos OS Release 19.4 and later. [PR1527602](#)
- The transit PTP packet might be modified unexpectedly when the packet is passed through MPC2E-NG, MPC3E-NG, and MPC5E. [PR1527612](#)
- The **commit confirm** command might not roll back the previous configuration when the commit operation fails. [PR1527848](#)
- 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](#)
- In the subscriber management environment, the RADIUS interim accounting records does not get populated with the subscriber statistics. [PR1529602](#)
- Deletion of the address of the jmgmt0 interface might fail if the shortened version of the CLI command is used. [PR1532642](#)
- The clear ike statistics with remote gateway does not work. [PR1535321](#)
- Multicast traffic might be sent out through unexpected interfaces with distributed IGMP enabled. [PR1536149](#)
- Version-alias is missed for subscribers configured with dynamic profiles after ISSU. [PR1537512](#)
- With hold time configuration, the ge interfaces remain down on reboot. [PR1541382](#)
- Port mirroring with the **maximum-packet-length** configuration does not work over GRE interface. [PR1542500](#)
- MPC10 or MPC11 line card might crash in case of Composite Chain Nextthop creation failures. [PR1538559](#)
- During an upgrade, vSRX3.0 would display the following incorrect license warnings when utilizing licensable features even if the license is present on the device: **warning: requires 'idp-sig' license.** [PR1519672](#)
- On the MX150 router, the logical interfaces stay up during vmhost halt or power-off. [PR1526855](#)
- ERO update by the controller for branch LSP might cause issues. [PR1508412](#)
- PEM 0 always shows as absent or empty even if PEM 0 is present on the MX10003 router. [PR1531190](#)

### **Infrastructure**

- If the serial number of the PEM starts with 1F1, the following alarm might be generated: **Minor FPC PEM Temp Sensor Failed.** [PR1398128](#)

- 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. [PR1508281](#)
- SNMP polling might return an unexpected high value for the ifHCOutOctets counter for a physical interface when any jnxDom OID is processed at the same time. [PR1508442](#)

### ***Interfaces and Chassis***

- The **sonet-options configuration** statement is disabled for the xe interface that works in the wan-phy mode. [PR1472439](#)
- Failure to configure proactive ARP detection. [PR1476199](#)
- Control logical interface 32767 is not created on the VLAN-tagged IFD even after removing the VLAN 0 configuration. [PR1483395](#)
- 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 the track priority-hold-time configured. [PR1506747](#)
- The vrrpd process might crash when the dual VLAN on VRRP interfaces is configured. [PR1512658](#)
- Commit failure is observed while deleting all the units under the ps0 interface. [PR1514319](#)
- When multiple CFM sessions are configured on IFD, the SNMP walk of **ieee8021CFMStack** table fails. [PR1517046](#)
- Inline Y.1731 SLM or DM does not work in enhanced-cfm-mode for the EVPN UP MEP scenario. [PR1537381](#)
- Buffer overflow vulnerability in a device control daemon is observed. [PR1519334](#)
- FPC crash might be observed with an inline mode with CFM configured. [PR1500048](#)

### ***Intrusion Detection and Prevention (IDP)***

- When creating the custom IDP signatures that match the raw bytes (hexadecimal), the commit check fails if the administrator configures the depth parameter. [PR1506706](#)

### ***Junos Fusion for 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***

- The aggregated Ethernet interface sometimes might not come up after the switch is rebooted. [PR1505523](#)
- The DHCPv6 lease query is not as expected while verifying the DHCPv6 server statistics. [PR1506418](#)
- The **show dhcp relay statistics** command displays **DHCPLEASEUNASSIGNED** instead of **DHCPLEASEUNASSINGED**, which is spelling error. [PR1512239](#)

- The `show dhcpv6 relay statistics` command 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](#)
- Memory leak in `jdhcpd` might be seen if access-profile is configured under the `dhcp-relay` or `dhcp-local-server` statement. [PR1525052](#)
- Receipt of the malformed DHCPv6 packets causes the `jdhcpd` process to crash. [PR1511782](#)
- The `jdhcpd` process crashes when a specific DHCPDv6 packet is processed in the DHCPv6 relay configuration. [PR1512765](#)

### **MPLS**

- The RSVP interface bandwidth calculation rounds up. [PR1458527](#)
- The same device responds twice for traceroute if it goes through the MPLS network under specific conditions. [PR1494665](#)
- 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 auto-bandwidth feature might not work correctly in an MPLS scenario. [PR1504916](#)
- Activating or deactivating the LDP-sync under OSPF might cause the LDP neighborship to go down and stay down. [PR1509578](#)
- The `rpd` process might crash after upgrading Junos OS Release 18.1 to a later release. [PR1517018](#)
- The SNMP trap is sent with the incorrect OID `jnxSpSvcSetZoneEntered`. [PR1517667](#)
- The LDP session-group might throw a commit error and flap. [PR1521698](#)
- `ping mpls rsvp` does not take into account for the lower MTU in the path. [PR1530382](#)
- The `rpd` process might crash when the LDP route with the indirect next-hop is deleted on the aggregated Ethernet interface. [PR1538124](#)
- The inter-domain LSP with loose next-hops path might get stuck in the **Down** state. [PR1524736](#)
- The RPD scheduler might slip after the link flaps. [PR1516657](#)

### **Network Address Translation (NAT)**

- Need to improve the maximum eNode connections for one persistent NAT binding from 8 to 32. [PR1532249](#)

### **Network Management and Monitoring**

- The SNMPv3 informs might not work properly after rebooting. [PR1497841](#)

### **Platform and Infrastructure**

- Packets are dropped when next-hop is IRB over an It interface. [PR1494594](#)
- Traffic to VRRP virtual IP or MAC addresses might be dropped when ingress queuing is enabled. [PR1501014](#)
- Traffic that originates from another subnet is sent out with 0x8100 instead of 0x88a8. [PR1502867](#)
- MPCs might crash when there is a change on routes learnt on the IRB interface configured in the VPLS or EVPN instances. [PR1503947](#)
- 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 performing virtual IP related change. [PR1511833](#)
- During the route table object fetch failure, the FPC process might crash. [PR1513509](#)
- The output value of the `show jnh qmon queues-sensor stats 0` command has no content. [PR1514881](#)
- VPLS connection might be stuck in the primary fail status when a dynamic profile is used on the VPLS pseudowire logical interface. [PR1516418](#)
- Configured scheduler-map is not applied on the ms- interface if the service PIC is in the **Offline** state during commit. [PR1523881](#)
- TWAMP interoperability issue between Junos OS releases is observed. [PR1533025](#)
- Packet loss might be observed when the RFC2544 egress reflector session is configured on the non-zero Packet Forwarding Ethernet interface. [PR1538417](#)
- Trio-based FPC might crash when the underlying layer 2 interface for ARP over IRB interface is changed from the physical interface to LSI interface. [PR1542211](#)

### **Routing Protocols**

- Multicast traffic loss might be seen in certain conditions while enabling IGMP snooping under the EVPN-VXLAN ERB scenario. [PR1481987](#)
- The output value 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. [PR1482983](#)
- BGP RPKI ROA withdrawal might lead to an unexpected BGP route flap. [PR1483097](#)
- There might be rpd memory leak in a certain looped MSDP scenario. [PR1485206](#)



- The rpd process might crash in a multicast scenario with the configured BGP. [PR1501722](#)
- On all Junos OS dual-Routing Engine GRES or NSR enabled routers, the rpd process might crash on a new primary Routing Engine if the Routing Engine switchover occurs right after massive routing-instance deletion. [PR1507638](#)
- The rpd process might crash due to RIP updates being sent on an interface in the **Down** state. [PR1508814](#)
- The rpd process might crash on the backup Routing Engine if the BGP (standby) receives a route from the peer, which is rejected due to an invalid target community. [PR1508888](#)
- The rpd process might report 100 percent CPU usage with the BGP route damping enabled. [PR1514635](#)
- ISIS-SR routes might not be updated to reflect the change in the SRMS advertisements. [PR1514867](#)
- The rpd process might crash after deleting and re-adding a BGP neighbor. [PR1517498](#)
- The rpd process might crash if there is a huge number of SA messages in the MSDP scenario. [PR1517910](#)
- Tag matching in the VRF policy does not work properly when the **independent-domain** option is configured. [PR1518056](#)
- The BGP-LS NLRI handling improvements are needed for BGP-LS ID TLV. [PR1521258](#)
- The IS-IS LSP database synchronization issue might be seen while using the flood-group feature. [PR1526447](#)
- Configuring **then next-hop** and **then reject** on a route policy for the same route might cause rpd to crash. [PR1538491](#)
- After moving the peer out of protection group, the path protection not removed from the PE router. [PR1538956](#)

### ***Services Applications***

- The FPC process might crash with the npc core file if the service interface is configured under service-set in the USF mode. [PR1502527](#)
- The output value 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***

- Subscriber accounting message retransmissions exist even after configuring accounting retry 0. [PR1405855](#)
- The 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 the user name of the subscriber contains space. [PR1518016](#)

### *User Interface and Configuration*

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

### *VPNs*

- MPLS label manager might allow configuration of a duplicated VPLS static label. [PR1503282](#)
- The rpd process might crash after removing the last interface configured under the Layer 2 circuit neighbor. [PR1511783](#)
- The rpd process might crash when deleting the Layer 2 circuit configuration in a specific sequence. [PR1512834](#)

### **Resolved Issues: 20.2R1**

#### *Application Layer Gateways (ALGs)*

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

#### *Class of Service (CoS)*

- The MX Series generated OAM/CFM LTR messages are sent with a different priority than the incoming OAM/CFM LTM messages. [PR1466473](#)
- The MX10008 and MX100016 routers might generate cosd core files after executing the **commit/commit check** command if the **policy-map** configuration is set. [PR1475508](#)
- Error message **GENCFG write failed (op, minor\_type) = (delete, Scheduler map definition) for tbl id 2 ifl 0 TABLE Reason: No such file or directory** is observed. [PR1476531](#)
- MX Series platforms with MPC1-Q and MPC2-Q line cards might report memory errors. [PR1500250](#)

#### *EVPN*

- Remote MAC address present in EVPN database might be unreachable. [PR1477140](#)
- Deleting a Layer 2 logical interface generates an error if the interface is not deleted first from EVPN. [PR1482774](#)
- The ESI of IRB interface does not update after autonomous-system number change if the interface is down. [PR1482790](#)
- Dead next-hops might flood in a rare scenario after remote PE devices are bounced. [PR1484296](#)

- The ARP entry gets deleted from the kernel after adding and deleting the virtual-gateway-address. [PR1485377](#)
- The rpd core file might be generated when doing Routing Engine switchover after disabling BGP protocol globally. [PR1490953](#)
- VXLAN bridge domain might lose VTEP logical interface after restarting chassisd. [PR1495098](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)
- The MAC address of the LT interface might not be installed in the EVPN database. [PR1503657](#)

### ***Forwarding and Sampling***

- IP-IP de-encapsulation fails if de-encapsulation filter is applied on loopback interface. [PR1469219](#)
- Traffic might be forwarded into the default queue instead of the correct queue when the VPLS traffic has three or more VLAN tags with VLAN priority 5. [PR1473093](#)
- The filter might not be installed if the **policy-map xx** is present under the filter. [PR1478964](#)

### ***General Routing***

- Syslog error message **PFEIFD: Could not decode media address with length 0** might be generated by the Packet Forwarding Engine. [PR1341610](#)
- The nondefault routing instance is not supported correctly for NTP packets in a subscriber scenario. [PR1363034](#)
- Egress monitored traffic is not mirrored to destination for analyzers on MX Series routers. [PR1411871](#)
- **FPC x Voltage Tolerance Exceeded** alarm raised and cleared upon bootup of JNP10K-LC2101. [PR1415671](#)
- The pccd starts running from the system start. [PR1417052](#)
- Resetting the Playback Engine logs are seen on the MPC5E line cards. [PR1420335](#)
- PF core voltage is not set according to the required e-fuse value and remains as default value of 0.9V on the JNP10008-SF and JNP10016-SF Switch Interface Boards (SIBs). [PR1420864](#)
- FPC might crash after GRES when you commit the changes in firewall filter with the **next term** statement in the subscriber scenario. [PR1421541](#)
- PTP might not work on the MX104 platform if phy-timestamping is enabled. [PR1421811](#)
- When you run the **show route label X | display json** command, two **nh** keys are present in the output. [PR1424930](#)
- PTP and show warning are disabled when hyper mode is configured. [PR1429527](#)
- Interfaces on the MPC-3D-16XGE-SFPP might go down due to CB0 clock failure. [PR1433948](#)
- ZF interrupts for out-of-range destination Packet Forwarding Engine INTR for Gnt are observed when the MPC6 or MPC9 line card is brought up. [PR1436148](#)

- System reboot is required when GRES is enabled or disabled with the **mobile-edge** configuration. [PR1444406](#)
- On the MPC10E-15C-MRATE with 25-Gigabit Ethernet ports, FEC statistics are not getting reset after changing FEC mode. [PR1449088](#)
- RE-MX2008-X8-128G secure BIOS version mismatch alarms. [PR1450424](#)
- Need to add support for drop flows when the packet drops. [PR1451921](#)
- When MVLAN interface (OIF map) is changed, the existing multicast subscribers with membership reports in place experience loss of multicast traffic until traffic is forwarded to a new OIF map. [PR1452644](#)
- Interfaces shutdown by the **disable-pfe** action might not be up using MIC offline or online command. [PR1453433](#)
- When scale configurations are applied from approximately 10 minutes, chassisd CLI will either have a delay in response or will time out. [PR1454638](#)
- On 4-port 1-Gigabit Ethernet using QSFP28 optics, continuous logging in chassisd process occurs when speed 1-Gigabit Ethernet is configured with **pic\_get\_nports\_inst** and **ch\_fru\_db\_key**. [PR1456253](#)
- On the MPC11E line card, need to add the support of optics-options low light. [PR1456894](#)
- LSP statistics are not getting reset after restart routing. [PR1458107](#)
- Inline S-BFD packets are dropped on MPC6E MIC1/PIC1 ports: 0-11. [PR1459529](#)
- Occasional warning message such as **TCP Connect error** can be seen during FPC reboot. [PR1460153](#)
- Multiple leaf devices and prefixes are missing when LLDP neighbor is added after streaming is started at the global level. [PR1460347](#)
- Support of **del\_path** for the LLDP neighbor change at various levels. [PR1460621](#)
- When you receive IPv6 over IPv4 IBGP session, the IPv6 prefix is hidden. [PR1460786](#)
- Explicit deletion notification (**del\_path**) is not received when LLDP neighbor is lost as a result of disabling local interface on the DUT through CLI (gNMI). [PR1461236](#)
- On the MPC10E line cards, more output packets than expected are seen when ping function is performed. [PR1461593](#)
- The **show dynamic-tunnel database** CLI command output does not filter IP-IP tunnels based on destination. [PR1461659](#)
- The **CHASSISD\_SNMP\_TRAP6: SNMP trap generated: Power Supply failed** message appears when both DIP switches and power switch are turned off. [PR1462065](#)
- Inline BFD session might flap on renegotiation of timers from slow to aggressive interval. [PR1462775](#)
- The MVPN traffic might be dropped after performing switchover. [PR1463302](#)
- The **native-vlan-id** functionality does not work and untagged traffic does not pass with the **native-vlan-id** configuration. [PR1463544](#)

- The `jdhcpd` process might consume high CPU use, and no further subscribers can be brought up if there are more than 4000 dhcp-relay clients in the MAC-MOVE scenario. [PR1465277](#)
- On the MPC10E and MPC11E line cards, the bandwidth-percent with shaping-rate might not work as expected on aggregated Ethernet interfaces after shaping-rate change. [PR1465766](#)
- The `bbe-smgd` process generates core files on the backup Routing Engine. [PR1466118](#)
- ICMP error messages are still unreceived after enabling the **enable-asymmetric-traffic-processing** configuration statement. [PR1466135](#)
- A few DHCP INFORM packets specific to a particular VLAN might be taking the incorrect resolve queue. [PR1467182](#)
- On the MPC11E line card, the DOM MIB alarm for the channelized 10-Gigabit Ethernet interface is not showing any alarm for LF/RF. [PR1467446](#)
- Daemons might not be started if **commit** is executed after **commit check**. [PR1468119](#)
- PPP IPv6 NCP fails to negotiate during the PPP login. [PR1468414](#)
- The `rpd` process might crash if BGP sharding is enabled. [PR1468676](#)
- The `tcp-log` connections fail to reconnect and get stuck in the **Reconnect-In-Progress** state. [PR1469575](#)
- Unable to set up 26M sessions (NAPT44) at 900,000 pps. [PR1470833](#)
- In rare occasions, the router might send out one extra URR quota value for a bearer. [PR1470890](#)
- Syslog message **FPCX user.notice logrotate: ALERT exited abnormally with [1]** pops at 04:02:01. [PR1471006](#)
- DHCP relay with forward-only might fail to send OFFER messages when DHCP client is terminated on logical tunnel interface. [PR1471161](#)
- Sudden FPC shutdown due to hardware failure or ungraceful removal of line card might cause major alarms on other FPCs in the system. [PR1471372](#)
- The `clksyncd` crash might be seen when PTP over aggregated Ethernet is configured on the MX104 platform. [PR1471466](#)
- On the MPC11E line card, locating a specific 100-Gigabit Ethernet, 40-Gigabit Ethernet, and 10-Gigabit Ethernet port in the card by blinking the corresponding port LED does not work. [PR1471894](#)
- Chassis alarm on BSYS might be observed: **RE0 to one or many FPCs is via em1: Backup RE**. [PR1472313](#)
- Performing back-to-back `rpd` restarts might cause `rpd` to crash. [PR1472643](#)
- Manually configured ERO on NS controller might be lost when PCEP session bounces. [PR1472825](#)
- SDB goes down very frequently if the **reauthenticate lease-renewal** statement is enabled for DHCP. [PR1473063](#)
- Some routes might not be installed into the FPC after it gets restarted. [PR1473079](#)

- On the MPC11E line card, **show dynamic-tunnels database** command does not show traffic statistics. [PR1473096](#)
- On MPC11, oversubscription drops are not accounted in Routing Engine CLI under resource drops when Flow control is disabled. [PR1473191](#)
- Dynamic-profile for VPLS-PW pseudowire incorrectly reports Dynamic Static Subscriber Base Feature license alarm. [PR1473412](#)
- On the MPC11E line card, after doing Routing Engine switchover on BSYS, the AF interface on peer router shows status as down with the reason being that the Packet Forwarding Engine is down on the GNF. [PR1473555](#)
- When both MSTP and ERP are enabled on the same interface, then ERP does not come up properly. [PR1473610](#)
- Drops counter does not increment for the aggregated Ethernet even after the member link shows the drops. [PR1473665](#)
- Ingress multicast replication does not work with GRES configuration. [PR1474094](#)
- DHCP-server RADIUS-given mask is being reversed. [PR1474097](#)
- On the MX150 platform, core files are not seen under **show system core-dumps**. [PR1474118](#)
- A newly added LAG member interface might forward traffic even though its micro BFD session is down. [PR1474300](#)
- Upon external X86 node slicing server reboot, the host SNMP configuration gets overwritten by the JDM SNMP configuration settings. [PR1474349](#)
- When traffic loss is observed on a 100-Gigabit Ethernet logical interface, the MACsec sessions are up and live. [PR1474714](#)
- On the MPC11E line card, basic circuit cross-connect traffic flow does not occur with the logical systems. [PR1474983](#)
- The clksyncd process generates core file after the GRES. [PR1474987](#)
- Memory leak leads to restart of the MPC10E line card. [PR1475036](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- The full list should be returned. A leaf should be considered atomic, regardless of whether it is a single value or a list for on-change event. [PR1475293](#)
- The RADIUS accounting updates of the service session have incorrect statistic data. [PR1475729](#)
- When xSTP protocols are enabled on interface all, it might run on **vlan-tagging/flexible-vlan-tagging** Layer 3 interfaces and lead to blocking of SXE interface. [PR1475854](#)
- Traffic loss might be seen as backup Routing Engine takes around 20 seconds to acquire the primary role. [PR1475871](#)

- Traffic drop might be observed while performing a unified ISSU on the MX2020, MX2010, and MX960 platforms. [PR1476505](#)
- The bbe-mibd might crash on an MX Series platform in subscriber environment. [PR1476596](#)
- On the MPC10 or MPC11 line cards, Routing Engine might not be able to send packets with traffic-manager enhanced-priority-mode configuration enabled. [PR1476683](#)
- The host-generated packets which might get dropped at the other end. [PR1476764](#)
- Traffic loss might occur to the LNS subscribers in case the **routing-service** statement is enabled under the dynamic profile. [PR1476786](#)
- Traffic loss might be seen in SAEGW scenario after the daemon restarts or after the GRES operation. [PR1477461](#)
- In NAT-T scenario, IKE version 2 IPsec tunnel flaps if the tunnel initiator is not behind NAT. [PR1477483](#)
- The rpd process might crash when the JET RIB API is used to set the "bandwidth" attribute. [PR1477745](#)
- On the MX2010 platform, syslog message **spmb0 cmt\_y\_sfb\_temp\_check: sfb[0] is powered OFF** & **"spmb0 cmt\_y\_sfb\_voltage\_check\_one: sfb[0] is powered OFF** is flooding even though SFBs are online. [PR1477924](#)
- Error log message **chassisd[7836]: %DAEMON-3-CHASSISD\_IOCTL\_FAILURE: acb\_get\_fpga\_rev: unable to get FPGA revision for Control Board (Inappropriate ioctl for device)** is observed after every commit. [PR1477941](#)
- The Packet Forwarding Engine might be disabled because of the major error on MPC2E-NG, MPC3E-NG, MPC5, MPC6, MPC7, MPC8, and MPC9. [PR1478028](#)
- The **show evpn statistics instance** command gets stuck in a multihomed scenario. [PR1478157](#)
- At-scale logins of both default and dedicated bearers might require retries from the control plane. [PR1478191](#)
- The ukern-platformd process might crash on MX2000 platforms with MPC11 line card. [PR1478243](#)
- Output chain filter counters are not proper. [PR1478358](#)
- MX Series-based MPC line card might crash when there is bulk route update failure in a corner case. [PR1478392](#)
- The FPC with **vpn-localization vpn-core-facing-only** configuration might be stuck in ready state. [PR1478523](#)
- On MX240, MX480, MX960, MX2000, MX10003, MX10008, and MX10016 with the MPC7E, MPC8E, and MPC9E line cards, hardware sensor information is logged every 30 minutes. [PR1478816](#)
- The protocol MTU might not be changed on lt- interface from the default value. [PR1478822](#)
- The TCP-log sessions might be in Established state but no logs are sent out to the syslog server. [PR1478972](#)

- Mobile-edge sessions might be lost if GRES is being performed while sessions are logged in with URR enabled. [PR1478985](#)
- The SCBE3 fabric plane gets into check state in MX Series Virtual Chassis. [PR1479363](#)
- Interface states are not showing correctly between main and shards on one of the interfaces. [PR1479801](#)
- After kmd restarts, IPsec SA comes up but the traffic fails for some time in certain scenarios. [PR1480692](#)
- 100-Gigabit interface might randomly fail to come up after maintenance operations. [PR1481054](#)
- Issue with binding non-default routing instance to existing soft-gre group. [PR1481278](#)
- After unified ISSU on the primary and the backup Routing Engine, **ISSU enhanced-mode: Performing action get-state for error /FPC/5/pfe/0/cm/0/PCle\_Error/0/PCIE\_CMERROR\_UNCORRECTABLE (0x190001)** error message is generated. [PR1481859](#)
- The rpd might crash when you execute the **show route protocol l2-learned-host-routing** or **show route protocol rift** CLI command on a router. [PR1481953](#)
- Log in to some PPPoE subscribers through aggregate Ethernet interface might cause the device to reboot. [PR1482431](#)
- Fragmentation limit and reassembly timeout configuration under services option is missing for SPC3. [PR1482968](#)
- When checking the BFD functionality over Layer 2 VPN client, BFD session is not coming up. [PR1483014](#)
- Link errors might be seen after restarting the FPC or fabric plane. [PR1483124](#)
- Traffic impact might be seen when the **policy-multipath** is configured without LDP on the SPRING-TE scenario. [PR1483585](#)
- The downstream IPv4 packet greater than BR MTU are getting dropped in MAP-E. [PR1483984](#)
- Traffic rate is not as expected on aggregated Ethernet interface when child links are from MPC11 and MPC9 line card after applying a policer. [PR1484193](#)
- ARP entry might not be created in the EVPN-MPLS environment. [PR1484721](#)
- The logical tunnel interface might not work on the MPC10 line card. [PR1484751](#)
- Fix and enhancement has been done for **request rift package activate** for the junos-rift package. [PR1485098](#)
- Attribute sending zero value should be compressed because it uses too much bandwidth in periodic streaming. [PR1485257](#)
- Interface input error counters are not increasing on the MX150 platforms. [PR1485706](#)
- The **krt-nextthop-ack-timeout** command might not automatically be picked up on restarting the rpd process. [PR1485800](#)
- MPC10E line card installed in the FPC slot 4 might drop host outbound traffic. [PR1485942](#)



- Command completion help text for LLDP-MED coordinate configuration statement contains spelling errors. [PR1486327](#)
- The afd process might crash when MPC10 line card is installed. [PR1487416](#)
- Incorrect frame length of 132 bytes might be captured in packet header. [PR1487876](#)
- XML is not properly formatted. [PR1488036](#)
- Add support for PSM firmware upgrade on the MX2000 platform. [PR1488575](#)
- During multiple login and logout of 250,000 sessions, there can be daemon restart due to mishandling of data. [PR1489512](#)
- NAT rule-sets processing order is not getting processed based on the order configured under **service-set**. It is getting processed based on the NAT rules defined under [**services nat source**] hierarchy level configuration. [PR1489581](#)
- With 4-member AMS used in the service-set, commit check fails when /30 subnet address is used as NAT pool IP. [PR1489885](#)
- Error syslog message **Failed to connect to the agentx primary agent (/var/agentx/primary): Unknown host (/var/agentx/primary) (No such file or directory)** is continuously being generated with dns-sinkholing. [PR1490487](#)
- When NAT/SFW rule is configured with application-set with multiple applications having different TCP inactivity-timeout, sessions are not getting TCP inactivity-timeout as per the configured application order. [PR1491036](#)
- The DAC cable is not detected after reboot or plug out or plug in. [PR1491116](#)
- The unified ISSU is not supported on next-generation MPC cards. [PR1491337](#)
- Multiple deactivating and activating of security traceoptions along with clear single NAPT44 session could result in generation of flowd core file. [PR1491540](#)
- MS-MIC is down after loading some releases in the MX Virtual Chassis scenario. [PR1491628](#)
- FPCs might stay down or restart when you swap the MPC7, MPC8, and MPC9 line cards with the MPC10 and MPC11 line cards or vice versa in the same slot. [PR1491968](#)
- User-configured MTU might be ignored after the unified ISSU upgrade uses **request vmhost software in-service-upgrade**. [PR1491970](#)
- Behavior change in clients with multiple gRPC channels to same target. [PR1492088](#)
- The delay of LT interfaces coming up is seen on MPC11E line card after you configure scaled PS interfaces anchoring to RLT. [PR1492330](#)
- On the MX10008 platform, SNMP table entPhysicalTable does not match the PICs shown for the **show chassis hardware** command. [PR1492996](#)
- DHCP subscribers do not come up as expected after deactivating the Virtual Chassis port. [PR1493699](#)

- The **ptp-clock-global-freq-tracable** leaf value becomes false and does not change to true when the internal lock is in the **Acquiring** state. [PR1493743](#)
- The LSP might not come up in LSP externally-provisioned scenario. [PR1494210](#)
- Error message **PFE\_ERROR\_FAIL\_OPERATION: Unable to unbind cos scheduler from physical interface 147** is observed on the MPC9E line card after restarting the MPC11E line card. [PR1494452](#)
- Missing firmware image file in **usr/share/pfe/firmware**. [PR1494557](#)
- In node slicing setup after GRES, RADIUS interim updates might not carry actual statistics. [PR1494637](#)
- Group address is not programmed back after deactivating and activating the bridge domain. [PR1495480](#)
- Flood next-hop ID is not same in both the primary and backup Routing Engines. [PR1495925](#)
- Error message **PFEIFD: Could not decode media address with length 0** is generated by the Packet Forwarding Engine when subscribers come up over a pseudowire interface. [PR1496265](#)
- Port numbers logged in ALG syslog are incorrect. [PR1497713](#)
- Subscribers might be disconnected after one of the aggregated Ethernet participating FPCs comes online in a Junos OS node slicing scenario. [PR1498024](#)
- SNMP polling does not show correct PSM jnxOperatingState when one of the PSM inputs failed. [PR1498538](#)
- The rpd might crash when multiple VRFs with 'IFLs link-protection' are deleted at a single time. [PR1498992](#)
- The commit check might fail when adding IFL into a routing instance with the **no-normalization** statement enabled under the **[routing-instances]** hierarchy. [PR1499265](#)
- The heap memory leak might be seen on the MPC10 and MPC11 line cards. [PR1499631](#)
- The SPC3 card might crash if SIP ALG is enabled. [PR1500355](#)
- On the MX2010 and MX2020 routers, the **pem\_tiny\_power\_remaining** message will be continuously logged in chassisd log. [PR1501108](#)
- Application ID does not display under NAT/SFW rule configured with application 'any' rule. [PR1501109](#)
- Support license start and end date in MIBs. [PR1503790](#)
- The **show bridge statistics** command does not display the statistics information for pseudowire subscriber interfaces. [PR1504409](#)
- The l2cpd crash might be seen if you add or delete ERP configuration and then restart l2cpd. [PR1505710](#)
- GnmiJuniperTelemetryHeader incompatibility is introduced in Junos OS Release 19.3. [PR1507999](#)
- The host generated packets might get dropped if the **force-control-packets-on-transit-path** statement is configured. [PR1509790](#)
- The multicast traffic might be dropped if ALB is enabled on the aggregated Ethernet interface. [PR1512157](#)

### **High Availability (HA) and Resiliency**

- Unified ISSU might fail on MX204 and MX10003 Virtual Chassis with an error message. [PR1480561](#)

### **Infrastructure**

- Slow response from SNMP might be observed after an upgrade to Junos OS Release 19.2R1 and later. [PR1462986](#)
- F-label veto code checks for per-pfe f-label pools. [PR1466071](#)

### **Interfaces and Chassis**

- Syslog error `scchassisd[ ]: CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC x` is observed after MX Virtual Chassis local or global switchover. [PR1428254](#)
- Decoupling of Layer 2 logical interfaces from bridge and EVPN configurations. [PR1438172](#)
- The MC-LAG configuration-consistency ICL configuration might fail after committing some changes. [PR1459201](#)
- On the MPC11E line card, the IPv6 local stats are counted against the IPv6 transit traffic statistics as well. [PR1467236](#)
- When you configure ESI on a physical interface, the traffic drops when you disable the logical interface under the physical interface. [PR1467855](#)
- Executing commit might hang because of stuck dcd process. [PR1470622](#)
- Traffic is not forwarded properly when traffic-control-profiles with logical interface queues are configured. [PR1475350](#)
- Commit error is not thrown when member link is added to multiple aggregation group with different interface specific options. [PR1475634](#)
- The interface on MIC3-100G-DWDM might go down after performing an interface flap. [PR1475777](#)
- When you delete and add a logical interface (both the logical interfaces with the same VLAN ID) in a single commit, the configuration check fails with the error **duplicate VLAN-ID**. [PR1477060](#)
- A stale IP address might be seen after a specific order of configuration changes in logical systems scenario. [PR1477084](#)
- Traffic is seen for 248 seconds when an aggregated Ethernet member link is brought down with minimum link configuration. [PR1477821](#)
- MC-AE interface might be shown as unknown status if you add the subinterface as part of the VLAN on the peer MC-AE node. [PR1479012](#)
- For ATM interfaces configuration, if any logical interface has the **allow-any-vci** configuration, then the commit operation might fail. [PR1479153](#)
- PPPoE subscribers are not up while verifying static IPv4 subscriber in passive mode. [PR1483395](#)

- CFM over BD along with negative events lead to restart and CFM DM two-way verification fails. [PR1489196](#)
- The **vrrp-inherit-from** change operation leads to packet loss when traffic is forwarded to the VIP gateway. [PR1489425](#)

### ***Intrusion Detection and Prevention (IDP)***

- The CLI now provides helpful remarks about IDP's tunable detector parameters. [PR1490436](#)
- When creating custom IDP signatures that match on raw bytes (hexadecimal), the commit check fails if the administrator has configured the depth parameter. [PR1506706](#)

### ***J-Web***

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

### ***Junos Fusion for Enterprise***

- SDPD core file is found at **vFPC\_all\_eports\_deletion\_complete vFPC\_dampen\_FPC\_timer\_expiry**. [PR1454335](#)
- Loop detection might not work on extended ports in Junos fusion scenarios. [PR1460209](#)

### ***Junos Fusion Satellite Software***

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

### ***Layer 2 Ethernet Services***

- On MX2010 and MX2020 platforms, no alarm is generated when FPC is connected to primary Routing Engine through backup Routing Engine/CB. [PR1461387](#)
- Member links state might be unsynchronized on a connection between a PE device and a CE device in an EVPN active/active scenario. [PR1463791](#)
- Telemetry data for relay/bindings/binding-state-v4relay-binding and relay/bindings/binding-state-v4relay-bound is not correct. [PR1475248](#)
- On the MX204 platform, the Vendor-ID is set as MX10001 in factory-default configuration and DHCP client messages. [PR1488771](#)
- With ALQ and VRRP configurations, DHCP subscribers are not coming up. [PR1490907](#)
- 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](#)

## Layer 2 Features

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

## MPLS

- Traffic loss might be seen if P2MP with NSR is enabled. [PR1434522](#)
- P2MP LSP might flap after VT interface in MVPN routing instance is reconfigured. [PR1454987](#)
- The RSVP interface bandwidth calculation rounds up. [PR1458527](#)
- The rpd might crash in PCEP for the RSVP-TE scenario. [PR1467278](#)
- The fast reroute detour next-hop down event might cause the primary LSP go in the **Down** state in a particular scenario. [PR1469567](#)
- The rpd process might crash during shutdown. [PR1471191](#)
- The LDP and BFD sessions are not coming up in a scaled setup. [PR1474204](#)
- The RSVP LSPs might not come up in a scaled network with a very high number of LSPs if NSR is used on the transit router. [PR1476773](#)
- PCC might flood with event logs to controller. [PR1476822](#)
- Kernel crashes and device might restart. [PR1478806](#)
- The rpd process crashes on the backup Routing Engine when LDP tries to create LDP P2MP tunnel upon receiving corrupted data from the primary Routing Engine. [PR1479249](#)
- On MX Series with MPC10E line card, rpd core files in `rsvp_copy_route (rt=< optimized out>, rtparms_p=< optimized out>)` at `../../../../../../../../src/junos/usr.sbin/rpd/mpls_te/proto/rsvp/proto/rsvp_route.c:3033` are seen after GRES. [PR1485985](#)
- The rpd might crash on restart of primary Routing Engine or backup Routing Engine when chain-NH has inner and outer labels in the SR-TE scenario. [PR1486077](#)
- High CPU utilization for rpd might be seen if RSVP is implemented. [PR1490163](#)
- The rpd might crash when BGP with FEC 129 VPWS enabled flaps. [PR1490952](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. [PR1493431](#)
- The rpd might crash in a rare condition under SR-TE scenario. [PR1493721](#)
- The rpd core files are generated during unified ISSU. [PR1493969](#)
- The rpd process might crash when SNMP polling is done using OID `jnxMplsTeP2MPTunnelDestTable`. [PR1497641](#)
- The rpd process might crash with RSVP configured in a rare timing case. [PR1505834](#)

### *Platform and Infrastructure*

- Core.vmxt.mpc0 is seen at 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](#)
- With chained composite next-hop enabled, the MPLS CoS rewrite does not work for IPv6 PE device traffic. [PR1436872](#)
- Traffic loss might be seen in case of Ethernet frame padding with VLAN. [PR1452261](#)
- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- On the MX204 platform, Packet Forwarding Engine errors might occur when incoming GRE tunnel fragments get sampled and undergo inline reassembly. [PR1463718](#)
- The CoS might not work on MPC10E and MPC11E line cards. [PR1465870](#)
- VXLAN packet might be discarded with flow caching enabled on MX150 and vMX. [PR1466470](#)
- All the subscriber services might be unavailable on vBNG running on MX150 and vMX running in payg mode. [PR1467368](#)
- The JNH memory leaks after CFM session flap for LSI and VT interfaces. [PR1468663](#)
- The switch might not be able to learn MAC address with `dot1x` and `interface-mac-limit` configured. [PR1470424](#)
- SSH login might hang and the TACACS+ server closes the connection without sending any authentication failure response. [PR1478959](#)
- Remote MEPs are not coming up as expected while verifying MIP functionality with bridge domains. [PR1484303](#)
- The `show system buffer` command displays all zeros in the MX104 chassis. [PR1484689](#)
- MAC learning under bridge domain stops after MC-LAG interface flaps. [PR1488251](#)
- MAC malformation might happen in a rare scenario under MX Series Virtual Chassis setup. [PR1491091](#)
- In node slicing setup, MPLS TTL might be set to zero when the packet goes through af interface configured with CCC family. [PR1492639](#)
- A specific IPv4 packet might lead to FPC restart. [PR1493176](#)
- Python or SLAX script might not be executed. [PR1501746](#)
- MPCs might crash when there is a change on routes learned on IRB interface configured in VPLS and EVPN instances. [PR1503947](#)
- Traffic convergence failed with ICL failure case. [PR1505465](#)

### ***Routing Policy and Firewall Filters***

- The router-id from martian address range cannot be committed even if the range is allowed by configuration. [PR1480393](#)

### ***Routing Protocols***

- The BGP session might be stuck with high BGP OutQ value after GRES on both sides. [PR1323306](#)
- PIM RPF selection for the specific multicast group might get incorrectly applied to other multicast groups. [PR1443056](#)
- TI-LFA might be unable to install backup path in the routing table in a specific case. [PR1458791](#)
- BGP NSR with more than 40,000 IPv6 peers is not qualified or supported. [PR1461436](#)
- IS-IS IPv6 routes might flap when there is an unrelated commit under protocol stanza. [PR1463650](#)
- The rpd might crash if IPv4 routes are programmed with IPv6 next-hop through JET APIs. [PR1465190](#)
- BGP peers might flap if the parameter of hold-time is set small. [PR1466709](#)
- The configured BGP damping policy might not take effect after BGP is disabled and then enabled followed by **commit**. [PR1466734](#)
- The rpd might stop when both instance-import and instance-export policies contain the as-path-prepend action. [PR1471968](#)
- Removing cluster from BGP group might cause prolonged convergence time. [PR1473351](#)
- Adjacency SID might be missed and not be advertised to peer/controller/BMP monitor in BGP-LS NLRI. [PR1473362](#)
- SFTP does not connect properly and the following error is displayed: **Received message too long**. [PR1475255](#)
- BGP TCP MD5 authentication support is not available. [PR1476669](#)
- The rpd process might crash with BGP multipath and route withdraw occasionally. [PR1481589](#)
- The rpd process crashes due to specific BGP UPDATE packets. [PR1481641](#)
- The rpd process might crash when deactivating logical systems. [PR1482112](#)
- BGP multipath traffic might not fully load-balance for a while after adding a new path for load sharing. [PR1482209](#)
- The rpd might be crashed after BGP peer flapping. [PR1482551](#)
- RIPV2 packets stop transmitting when changing interface-type configuration from P2MP to broadcast. [PR1483181](#)
- The rpd process crashes if the same neighbor is set in different RIP groups. [PR1485009](#)
- On MX Series, MSDP memory leak is observed. [PR1485206](#)
- The BGP-LU routes do not have the label when BGP sharding is used. [PR1485422](#)

- Removal of the BGP and rib-sharding configuration might cause routing protocols to become unresponsive. [PR1485720](#)
- Layer 3 VPN RR with **family route-target** and **no-client-reflect** statements does not work as expected. [PR1485977](#)
- Traffic loss is seen on a scaled MPLS setup after unified ISSU in enhanced mode. [PR1486657](#)
- The rpd process crashes if the BGP LLGR with RIB sharding and traceoptions for graceful-restart are configured. [PR1486703](#)
- The rpd might crash when you perform GRES with MSDP configured. [PR1487636](#)
- High CPU utilization might be observed when the outgoing BGP updates are sent slowly. [PR1487691](#)
- The rpd process might generate core file after **always-compare-med** is configured for BGP path-selection. [PR1487893](#)
- BGP RIB sharding feature cannot be run on a system with a single CPU. [PR1488357](#)
- The rpd crashes when reset OSPF neighbors. [PR1489637](#)
- The BGP route target family might prevent route reflector from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- The rpd might crash because of rpd resolver problem of INH. [PR1494005](#)
- The static route in inet6.0 or inet6.3 RIB might fail to delete. [PR1495477](#)
- For SPRING support SRv6, continuous rpd core files are generated at **isis\_set\_rt\_pfx\_sid\_tsi,isis\_route\_change\_rt** after configuring **[set protocols isis topologies ipv6-unicast]**. [PR1495994](#)
- Receipt of certain genuine BGP packets from any BGP speaker causes rpd to crash. [PR1497721](#)
- The rpd might crash if the import policy is changed to accept more routes that exceed the teardown function threshold. [PR1499977](#)
- The rpd process crashes when processing a specific BGP packet. [PR1502327](#)
- The **show bgp neighbors** command shows change in x-path output for **input-updates** value. [PR1504399](#)
- BGP might not advertise routes to peers after a peer flap. [PR1507195](#)

### **Services Applications**

- **flow-tap** add function might not work after the dynamic flow capture services process is restarted. [PR1472109](#)
- On an MX Series router, L2TP LTS fails to forward the **agentCircuitId** and **agentRemotId** AVP toward the LNS. [PR1472775](#)



- The kmd might crash due to the incorrect IKE SA establishment after the remote peer's NAT mapping address has been changed. [PR1477181](#)
- NPC core files are found at `services_inline_handle_svc_set_add services_inline_gencfg_handler gencfg_specific_handler`. [PR1502527](#)

### ***Subscriber Access Management***

- The authd process might crash after the unified ISSU from Junos OS Release 18.3 and earlier to Junos OS Release 18.4 and later. [PR1473159](#)
- Syslog messages `pfe_tcp_listener_open_timeout: Peer info msg not received from addr: 0x6000080. Socket 0xffff804ad23c2e0 closed` is observed. [PR1474687](#)
- The delete request of a specified service session through CoA could fail. [PR1479486](#)
- The CoA request might not be processed if it includes the `proxy-state` attribute. [PR1479697](#)
- The `mac-address` CLI option is hidden under the `access profile profile-name radius options calling-station-id-format` statement. [PR1480119](#)
- The authd log events might not be sent to syslog host when `destination-override` is used. [PR1489339](#)

### ***VPNs***

- Traffic loss might be observed when the inter-AS next-generation MVPN VRF is disabled on one of the ASBRs. [PR1460480](#)
- The rpd might crash when "link-protection" is added or deleted from LSP for MVPN ingress replication selective provider tunnel. [PR1469028](#)
- On MVPN scenario, the LSP might stay down on removing all VT interfaces from a single hop egress. [PR1474830](#)
- The MPC10E-15C-MRATE next-generation MPVN ingress replication flushing out is not proper when in egress the ingress replication configuration is deactivated. [PR1475834](#)
- The Layer 2 circuit neighbor might be stuck in RD state at one end of MG-LAG peer. [PR1498040](#)
- The rpd core files are generated while disabling Layer 2 circuit with connection protection, backup neighbor configuration, and Layer 2 circuit trace logs enabled. [PR1502003](#)
- The rpd might crash when you delete l2circuit configuration in a specific sequence. [PR1512834](#)

### SEE ALSO

[What's New | 102](#)

[What's Changed | 129](#)

[Known Limitations | 136](#)

---

[Open Issues | 139](#)

---

[Documentation Updates | 185](#)

---

[Migration, Upgrade, and Downgrade Instructions | 186](#)

## Documentation Updates

### IN THIS SECTION

- [Advanced Subscriber Management Provider | 185](#)

This section lists the errata and changes in Junos OS Release 20.2R3 documentation for MX Series.

### Advanced Subscriber Management Provider

- The Broadband Subscriber Services User Guide incorrectly stated that for Routing Engine-based, converged HTTP redirect services, a CPCD service rule can include both a redirect term and a rewrite term. It also incorrectly stated that you can include separate rewrite and redirect rules in the same service profile.

### SEE ALSO

---

[What's New | 102](#)

---

[What's Changed | 129](#)

---

[Known Limitations | 136](#)

---

[Open Issues | 139](#)

---

[Resolved Issues | 149](#)

---

[Migration, Upgrade, and Downgrade Instructions | 186](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.2R3 | 187](#)
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS | 187](#)
- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | 190](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 191](#)
- [Upgrading a Router with Redundant Routing Engines | 192](#)
- [Downgrading from Release 20.2R3 | 192](#)

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 20.2R3

**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-x86-32-20.2R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.2R3.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-20.2R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.2R3.9-limited.tgz
```

Replace source with one of the following values:

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

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

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

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

Rebooting occurs only if the upgrade is successful.

**NOTE:**

- You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 20.2R3, 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 20.2R3 `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-20.2R3.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot  
source/jinstall-ppc-20.2R3.9-limited-signed.tgz
```

Replace source with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

**NOTE:** After you install a Junos OS Release 20.2R3 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

## 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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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 EOL release and then upgrade or downgrade from that EOL release to your target release.

For more information about EOL releases and to review a list of EOL 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 20.2R3

To downgrade from Release 20.2R3 to another supported release, follow the procedure for upgrading, but replace the 20.2R3 jinstall package with one that corresponds to the appropriate release.

**NOTE:** You cannot downgrade more than three releases.

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

### SEE ALSO

---

[What's New | 102](#)

---

[What's Changed | 129](#)

---

[Known Limitations | 136](#)

---

[Open Issues | 139](#)

---

[Resolved Issues | 149](#)

---

# Junos OS Release Notes for NFX Series

## IN THIS SECTION

- What's New | 193
- What's Changed | 195
- Known Limitations | 196
- Open Issues | 197
- Resolved Issues | 199
- Documentation Updates | 201
- Migration, Upgrade, and Downgrade Instructions | 202

These release notes accompany Junos OS Release 20.2R3 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

### IN THIS SECTION

- What's New in Release 20.2R3 | 194
- What's New in Release 20.2R2 | 194
- What's New in Release 20.2R1 | 194

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

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

### What's New in Release 20.2R3

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 20.2R2.

### What's New in Release 20.2R1

#### *Application Security*

- **AppQoE multihoming with active-active deployment (NFX150, NFX250, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and vSRX)**—Starting In Junos OS Release 20.2R1, AppQoE is enhanced to support multihoming with active/active deployment. In previous releases, AppQoE supports multihoming with active/standby deployment.

In active/active deployment, the spoke device connects to multiple hub devices. Application traffic can transit through any of the hub devices if the link to the hub device meets SLA requirements. Application traffic can switch seamlessly between the hub devices in case of SLA violation or if the active hub device is not responding.

To support active/active mode, you must enable the BGP multipath to allow the device to select multiple equal-cost BGP paths to reach a given destination.

[[Application Quality of Experience \(AppQoE\)](#).]

- **Packet capture for unknown application traffic (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.2R1, you can generate packet capture information for unknown application traffic on your security device. You can use this information to get more insight on unknown applications.

After you configure packet capture for the application traffic on your device, the packet capture function captures the packet details and stores the information in a packet capture (.pcap) file. You can use the packet capture details of an unknown application to define a new custom application signature and create a security policy rule to manage the application traffic more efficiently.

You can submit the packet capture information to Juniper Networks to debug why an application is not detected, and if required, request to create an application signature.

[See [Application Identification](#).]

### High Availability

- **High availability on NFX250 NextGen devices**—Starting in Junos OS Release 20.2R1, NFX250 NextGen devices support the high availability feature. You can configure a cluster of two NFX250 NextGen devices to act as primary and secondary devices for protection against device failures. The high availability feature supports Layer 2 and Layer 3 features in dual CPE deployments.

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

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

### Interfaces

- **ADSL and VDSL interfaces on NFX350 devices**—Starting in Junos OS Release 20.2R1, NFX350 devices support ADSL and VDSL interfaces.

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

### SEE ALSO

[What's Changed | 195](#)

[Known Limitations | 196](#)

[Open Issues | 197](#)

[Resolved Issues | 199](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 196](#)
- [What's Changed in Release 20.2R2 | 196](#)
- [What's Changed in Release 20.2R1 | 196](#)

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

### What's Changed in Release 20.2R3

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 20.2R3 for NFX Series devices.

### What's Changed in Release 20.2R2

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in Junos OS Release 20.2R2 for NFX Series devices.

### What's Changed in Release 20.2R1

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 20.2R1 for NFX Series devices.

#### SEE ALSO

---

[What's New | 193](#)

---

[Known Limitations | 196](#)

---

[Open Issues | 197](#)

---

[Resolved Issues | 199](#)

---

[Documentation Updates | 201](#)

---

[Migration, Upgrade, and Downgrade Instructions | 202](#)

## Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.2R3 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.

#### SEE ALSO

---

[What's New | 193](#)

[What's Changed | 195](#)

---

[Open Issues | 197](#)

---

[Resolved Issues | 199](#)

---

[Documentation Updates | 201](#)

---

[Migration, Upgrade, and Downgrade Instructions | 202](#)

## Open Issues

### IN THIS SECTION

- [High Availability | 198](#)
- [Interfaces | 198](#)
- [Platform and Infrastructure | 198](#)
- [Virtual Network Functions \(VNFs\) | 198](#)

Learn about open issues in Junos OS Release 20.2R3 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

- On an NFX150 chassis cluster device, the first packet is getting dropped while validating VLAN support on reth interfaces and child link. [PR1488462](#)

## Interfaces

- When you configure analyzers on VNF interfaces with output port as other VNF interfaces, all the incoming and outgoing packets can be mirrored on to the designated analyzer port. However, it is noticed that after a system reboot, this functionality stops working and no packets are mirrored on the output analyzer port. [PR1480290](#)

## Platform and Infrastructure

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

## 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 | 193](#)

---

[What's Changed | 195](#)

---

[Known Limitations | 196](#)

---

[Resolved Issues | 199](#)

---

[Documentation Updates | 201](#)

---

[Migration, Upgrade, and Downgrade Instructions | 202](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 199](#)
- [Resolved Issues: 20.2R2 | 200](#)
- [Resolved Issues: 20.2R1 | 200](#)

Learn which issues were resolved in the Junos OS Release 20.2R3 for NFX Series devices.

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

### Resolved Issues: 20.2R3

#### *General Routing*

- False positive TSensor errors are reported on vjunos0. [PR1508580](#)

#### *High Availability*

- On NFX150 devices, upgrade from Junos OS Release 19.4 to Junos OS Release 20.2 fails and the `/usr/sbin/boot_mgmt_fsm: line 40: echo: write error: No space left on device` message is displayed. [PR1532334](#)

#### *Interfaces*

- On NFX Series and MX150 devices, the following error messages are seen in the messages log file for the interfaces that have SFP installed in them: `fpc0 FAILED(-1) read of SFP eeprom for port: 13`. [PR1529939](#)

#### *Platform and Infrastructure*

- On NFX150, NFX250 NextGen, and NFX350 devices, the following command is not supported: `request load merge filename`. [PR1533284](#)
- On NFX250 devices, the l2cpd core files might be seen on reboot. This is a one-time core and does not impact the functionality of the device. [PR1561235](#)



## Resolved Issues: 20.2R2

### *High Availability (HA)*

- On NFX150 devices, upgrade from Junos OS Release 19.4 to Junos OS Release 20.2 fails and the `/usr/sbin/boot_mgmt_fsm: line 40: echo: write error: No space left on device` issue message is displayed. [PR1532334](#)

### *Interfaces*

- On NFX350 devices, the `show interfaces | no-more` command output stops appearing for around 20 seconds after displaying the d10 interface. [PR1502626](#)

### *Platform and Infrastructure*

- 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](#)
- SDWAN NFX150 HA - while upgrade from 19.4 -> 20.2 observed `"/usr/sbin/boot_mgmt_fsm: line 40: echo: write error: No space left on device issue"` that is not allowing to upgrade.

## Resolved Issues: 20.2R1

### *Application Security*

- AppQoE is sending active prob packets for the deleted active-probe-params. [PR1492208](#)

### *High Availability*

- On NFX250 chassis cluster, L3 interfaces are not getting created after secondary automatic reboot when control port recovery is enabled. [PR1502449](#)

### *Interfaces*

- On NFX150 devices, no error is displayed when the commit fails after you configure `native-vlan-id` on an access VNF interface. [PR1438854](#)
- On NFX250 NextGen devices, the `monitor interface traffic` command might not display the pps output for SXE and physical interfaces. [PR1464376](#)
- On NFX350 devices, the `clear interface statistics all` command takes a longer time to execute. [PR1475804](#)
- On NFX350 devices, if you delete and add an SXE interface, the SXE interface moves to the Spanning Tree Protocol blocking (STP BLK) state, and the traffic drops on that interface. [PR1475854](#)

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

- On NFX Series devices, IP identification (IP ID) is not changed after MAP-E NAT44 is performed on fragment packets when the packets reach the customer edge (CE) device. [PR1478037](#)

### **Platform and Infrastructure**

- On NFX150 devices, MAC aging does not work. You must remove aged MAC entries from the CLI. [PR1502700](#)
- On NFX350 devices, if you execute the **show vmhost mode** command multiple times, JDM might crash and cause the **show vmhost mode** commands to stop working. [PR1474220](#)
- Core files on NFX250 while adding the second LAN subnet. [PR1490077](#)
- 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](#)

### **Virtualized Network Functions (VNFs)**

- On NFX150 and NFX250 NextGen devices, when two flowd interfaces are mapped to the same physical interface and if you delete the interface mapping to VFO, the traffic flow is disrupted. Even though the mapping is moved to VFO, the MAC address is not cleared in VF1, which disrupts the traffic. [PR1448595](#)
- On NFX350 devices, VNF instantiation is not working properly. [PR1478456](#)

#### SEE ALSO

[What's New | 193](#)

[What's Changed | 195](#)

[Known Limitations | 196](#)

[Open Issues | 197](#)

[Documentation Updates | 201](#)

[Migration, Upgrade, and Downgrade Instructions | 202](#)

## **Documentation Updates**

There are no errata or changes in Junos OS Release 20.2R3 documentation for NFX Series devices.

#### SEE ALSO

---

[What's New | 193](#)

---

[What's Changed | 195](#)

---

[Known Limitations | 196](#)

---

[Open Issues | 197](#)

---

[Resolved Issues | 199](#)

---

[Migration, Upgrade, and Downgrade Instructions | 202](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 202](#)
- [Basic Procedure for Upgrading to Release 20.2 | 203](#)

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 20.2

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

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

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

To download and install Junos OS Release 20.2R3:

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

9. Copy the software to the device or to your internal software distribution site.
10. Install the new package on the device.

#### SEE ALSO

---

[What's New | 193](#)

---

[What's Changed | 195](#)

---

[Known Limitations | 196](#)

---

[Open Issues | 197](#)

---

[Resolved Issues | 199](#)

---

[Documentation Updates | 201](#)

## Junos OS Release Notes for PTX Series

#### IN THIS SECTION

- [What's New | 205](#)
- [What's Changed | 213](#)
- [Known Limitations | 216](#)
- [Open Issues | 218](#)
- [Resolved Issues | 221](#)
- [Documentation Updates | 226](#)
- [Migration, Upgrade, and Downgrade Instructions | 226](#)

These release notes accompany Junos OS Release 20.2R3 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

### IN THIS SECTION

- [What's New in Release 20.2R3 | 205](#)
- [What's New in Release 20.2R2 | 205](#)
- [What's New in Release 20.2R1 | 205](#)

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

### What's New in Release 20.2R3

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features or enhancements to existing features for PTX Series routers in Junos OS Release 20.2R2.

### What's New in Release 20.2R1

#### *High Availability (HA) and Resiliency*

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes roles. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

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

- **Unsupported hardware for unified ISSU (MX240, MX480, MX960, MX10003, and PTX3000)**—The following cards do not support unified ISSU upgrading to Junos OS Release 20.2R1:

- MPC7E-MRATE
- MPC8E with MRATE MIC
- MPC9E with MRATE MIC
- MPC10E-10C-MRATE
- MPC10E-15C-MRATE
- PTX5000 with 24-Port 10-Gigabit Ethernet, 40-Gigabit Ethernet PIC with QSFP+ or 15-Port 10-Gigabit, 40-Gigabit Ethernet, 100-Gigabit Ethernet PIC with QSFP28
- MX10003 with QSFP28 Ethernet TIC

### *Interfaces and Chassis*

- **Support for 1-Gbps speed on QFX10000-60S-6Q line card (PTX10008 and PTX10016)**—In Junos OS Release 20.2R1 and later, the QFX10000-60S-6Q line card supports 1-Gbps speed on its ports (0 to 59). The QFX10000-60S-6Q line card contains 60 SFP+ ports that support 10 Gbps, two dual-speed QSFP28 ports that support either 40 Gbps or 100 Gbps, and four QSFP+ ports that support 40 Gbps. You can individually configure ports 0 to 59 for 10-Gbps or 1-Gbps port speed. Use the **set chassis fpc fpc-slot-number pic pic-number port port-number speed 1G** command to change the mode of a port from 10 Gbps to 1 Gbps. The transceivers supported for 1 Gbps are QFX-SFP-1GE-LX, QFX-SFP-1GE-SX, and QFX-SFP-1GE-T.

By default, the QFX1000-60S-6Q line card (ports 0 to 59) operates at 10-Gbps speed.

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

### *Juniper Extension Toolkit (JET)*

- **RIB service APIs support dynamic next-hop interface binding (MX Series, PTX Series, and vMX)**—Starting in Junos OS Release 20.2R1, programmed RIB routes react to Up, Down, Add, and Delete events for direct next-hop interfaces. When all direct next-hop interfaces are unusable, the route becomes inactive. This prevents traffic from being dropped and keeps inactive routes from being propagated through the network.

This feature applies to all routes programmed using the `rib_service` JET API where an interface is configured as a direct next hop, including interfaces that are part of a flexible tunnel. It also applies to tunnels configured with the `flexible_tunnel_service` JET API.

To disable this feature, use **edit routing-options programmable-rpd rib-service dynamic-next-hop-interface disable**.

[See [rib-service \(programmable-rpd\)](#), [Juniper Extension Toolkit Developer Guide](#), and [Juniper Engineering Network website](#).]

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

[See [language \(Scripts\)](#), [Develop Off-Device JET Applications](#), and [Develop On-Device JET Applications](#).]

### *Junos Telemetry Interface*

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models `openconfig-local-routing.yang` and `openconfig-network-instance.yang`.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

- **ON-CHANGE BGP peer information statistics support for JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 provides BGP peer sensor support using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON\_CHANGE statistics are sent to an outside collector.

The following resource paths are supported:

- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/active` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/received` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/sent` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/rejected` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/admin-state` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/established-transitions` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/last-established` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/received/notification` (stream)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/messages/received/update` (stream)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/notification` (stream)



- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/update` (stream)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/session-state` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/supported-capabilities` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/transport/state/local-address` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-address` (ON\_CHANGE)
- `/network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-port` (ON\_CHANGE)

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

- **Telemetry support for LDP and MLDP traffic statistics (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, the following LDP and multipoint LDP native sensors are added for the Junos telemetry interface:
  - `/junos/services/ldp/label-switched-path/ingress/usage/`
  - `/junos/services/ldp/label-switched-path/transit/usage/`
  - `/junos/services/ldp/p2mp/interface/receive/usage/`
  - `/junos/services/ldp/p2mp/interface/transmit/usage/`
  - `/junos/services/ldp/p2mp/label-switched-path/usage/`

You must enable telemetry streaming with the **sensor-based-stats** option at the **[edit protocols ldp traffic-statistics]** hierarchy level.

The **show ldp traffic-statistics** command is enhanced to display upstream LDP traffic statistics and to display multipoint LDP traffic statistics per interface.

On PTX Series routers, this feature is not supported for the following variants:

- PTX3000 and PTX5000 with the RE-DUO-C2600-16G Routing Engine
- PTX10003
- PTX10008 with the PTX10K-LC1201-36CD line card
- FPC2 line cards do not support ingress multipoint LDP statistics.

[See [sensor \(Junos Telemetry Interface\)](#).]

- **CPU statistics support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 supports streaming various CPU statistics and process parameters using remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) services

and Junos telemetry interface (JTI). You can stream CPU usage per process (statistics are similar to output from the **show system process detail** operational mode command), as well as CPU usage per Routing Engine core.

This feature supports the private data model `openconfig-procmon.yang`.

To stream statistics to an outside collector, include the following resource paths in a gRPC or gNMI subscription:

- Individual process level information (resource path `/system/processes/process`)
- Individual Routing Engine core information (resource path `/components/component/cpu/`)

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

- **Packet Forwarding Engine sensor support with INITIAL\_SYNC on JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000 line of routers, QFX5100, and QFX5200)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export Packet Forwarding Engine statistics from devices to an outside collector using gNMI submode `INITIAL_SYNC`. When an external collector sends a subscription request for a sensor with `INITIAL_SYNC` (`gnmi-submode 2`), the host sends all supported target leaves (fields) under that resource path at least once to the collector with the current value. This is valuable because:
  - The collector has a complete view of the current state of every field on the device for that sensor path.
  - Event-driven data (`ON_CHANGE`) is received by the collector at least once before the next event is seen. In this way, the collector is aware of the data state before the next event happens.
  - Packet Forwarding Engine sensors that contain zero counter values (zero-suppressed) that normally do not show up in streamed data are sent, ensuring that all fields from each line card (also referred to as source) are known to the collector.

**NOTE:** `ON_CHANGE` data is not available for native (UDP) Packet Forwarding Engine sensors.

`INITIAL_SYNC` submode requires that at least one copy to be sent to the collector; however, sending more than one is acceptable.

`INITIAL_SYNC` submode is supported for the following sensors:

- Sensor for CPU (ukernel) memory (resource path `/junos/system/linecard/cpu/memory/`)
- Sensor for firewall filter statistics (resource path `/junos/system/linecard/firewall/`)
- Sensor for physical interface traffic (resource path `/junos/system/linecard/interface/`)
- Sensor for logical interface traffic (resource path `/junos/system/linecard/interface/logical/usage/`)
- Sensor for physical interface queue traffic (resource path `/junos/system/linecard/interface/queue/`)

- Sensor for physical interface traffic except queue statistics (resource path `/junos/system/linecard/interface/traffic/`)
- Sensor for NPU memory (resource path `/junos/system/linecard/npu/memory/`)
- Sensor for NPU utilization (resource path `/junos/system/linecard/npu/utilization/`)
- Sensor for packet statistics (resource path `/junos/system/linecard/packet/usage/`)
- Sensor for software-pollled queue-monitoring statistics (resource path `/junos/system/linecard/qmon-sw/`)

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

## **MPLS**

- **Support for MPLS ping and traceroute for segment routing (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 20.2R1, we extend the MPLS ping and traceroute support for all types segment routing--traffic engineering (SR-TE) tunnels, including static segment routing tunnels, BGP-SR-TE tunnels, and PCEP tunnels.

We also support the following features:

- FEC validation support, as defined in RFC 8287, for paths consisting of IGP segments. Target FEC stack contains single or multiple segment ID sub-TLVs. This involves validating IPv4 IGP-Prefix Segment and IGP-Adjacency Segment ID FEC-stack TLVs.
- ECMP traceroute support for all types of SR-TE paths.

We do not support the following:

- Ping and traceroute for SR-TE tunnel for non-enhanced-ip mode.
- OAM for IPv6 prefix.
- BFD

[See [traceroute mpls segment-routing spring-te](#) and [ping mpls segment routing spring-te](#).]

## **Network Management and Monitoring**

- **SNMP support for multicast LDP MIB objects (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS SNMP extends support for the following multicast LDP MIB tables and objects:
  - `mplsMldpInterfaceStatsTable`
  - `mplsMldpFecUpstreamSessPackets`
  - `mplsMldpFecUpstreamSessBytes`
  - `mplsMldpFecUpstreamSessDiscontinuityTime`

The multicast LDP standard MIB builds on the objects and tables that are defined in RFC3815, which only supports LDP point-to-point label-switched paths (LSPs). This multicast LDP MIB provides support for managing multicast LDP point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs.

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

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

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

- **Enhanced on-box monitoring support on the control plane (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure traceoptions to track all events related to system-level and process-level memory monitoring. You can also view the history of the actions taken for system-level and process-level memory monitoring by using the **show system monitor memory actions** command.

### *Routing Policy and Firewall Filters*

- **Support for additional route filter qualifiers in a policy statement (PTX1000 and PTX10000)**—Starting in Junos OS Release 20.2R1, the following list-level qualifiers are supported: **exact**, **longer**, **orlonger**, **prefix-length-range**, and **upto**.

You can use route filter lists to group individual route filters created at the **[edit policy-options]** hierarchy level. Each item in a list consists of a complete route filter statement, including a destination prefix, a match type, and an optional action. Reuse the list in different policies, adding whatever qualifiers you need, instead of re-creating a different one for every use case.

[See [Understanding Route Filters for Use in Routing Policy Match Conditions](#).]

### *Routing Protocols*

- **TI-LFA SRLG protection for IS-IS (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can configure Shared Risk Link Group (SRLG) protection in topology-independent loop free alternate (TI-LFA) networks. IS-IS computes the fast reroute path that is aligned with the post-convergence path and excludes the SRLG of the protected link. All local and remote links that share any SRLG with the

protecting link are excluded. The point of local repair (PLR) sets up the label stack for the fast reroute path with a different outgoing interface.

To enable TI-LFA SRLG protection with segment routing for IS-IS, include the **srlg-protection** statement at the `[edit protocols isis interface name level number post-convergence-lfa]` hierarchy level.

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

- **Support for BGP-LU over SR-TE for color-based mapping of VPN Services (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, we are extending support to BGP labeled unicast service for color-based mapping of VPN services over Segment Routing-Traffic Engineering (SR-TE). This enables you to advertise BGP-LU IPv6 and IPv4 prefixes with an IPv6 next-hop address in IPv6-only networks where routers do not have any IPv4 addresses configured. With this feature, BGP-LU can now resolve IPv4 and IPv6 routes over the SR-TE core. BGP-LU constructs a colored protocol next hop, which is resolved on a colored SR-TE tunnel in the **inetcolor.0** or **inet6color.0** table. Currently, we support BGP IPv6 LU over SR-TE with IS-IS underlay.

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

- **Support for BGP-SR-TE rearchitecture (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, Junos OS provides support for controller-based BGP segment routing--traffic engineering (SR-TE) routes to be installed as source packet routing traffic-engineered (SPRING-TE) routes. BGP installs the SR-TE policy in the routing tables **bgp.inetcolor.0** and **bgp.inet6color.0**, and these routes are subsequently installed in the routing tables **inetcolor.0** or **inet6color.0** by SPRING-TE.

In releases before Junos OS Release 20.2R1, controller-based BGP SR-TE routes are installed as BGP routes in the routing table. To maintain consistency and for easy maintenance, all SR-TE based routes appear as SPRING-TE routes irrespective of the source.

You need to enable **source-packet-routing** at the `[edit protocols]` hierarchy level to see the routes installed in **inetcolor.0** or **inet6color.0**. A new option **detail** is introduced under **traceoptions (Protocols Spring-TE)** to trace the detailed information.

[See [Segment Routing Traffic Engineering at BGP Ingress Peer Overview.](#)]

## System Logging

- **Support to track the maximum number of routing and forwarding (RIB/FIB) routes and VRFs (MX Series and PTX Series)**—Starting in Junos OS Release 20.2R1, you can track and display the high-water mark data of routing and forwarding (RIB/FIB) table routes and VRFs in a system (RPD) using the **show route summary** CLI command. High-water mark refers to the maximum number of routing and forwarding (RIB/FIB) table routes and VRFs that were present in the RPD system. The high-water mark data can also be viewed in the syslog at the **LOG\_NOTICE** level.

You can configure the interval of the high-water mark data using the **highwatermark-log-interval** CLI configuration statement at the **[edit routing-options]** hierarchy level. The minimum time gap at which the high-water mark data logged in the syslog is 30 seconds. You can configure the value for **highwatermark-log-interval** CLI configuration statement between 5 and 1200 seconds.

[See [routing-options](#) and [show route summary](#).]

## SEE ALSO

[What's Changed | 213](#)

[Known Limitations | 216](#)

[Open Issues | 218](#)

[Resolved Issues | 221](#)

[Documentation Updates | 226](#)

[Migration, Upgrade, and Downgrade Instructions | 226](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 214](#)
- [What's Changed in Release 20.2R2 | 215](#)

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

## What's Changed in Release 20.2R3

### *Junos OS XML API and Scripting*

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

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

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

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

### *Network Management and Monitoring*

- Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:
  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.
  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and emits only an `<ok>` or `<rpc-error>` element.

[See [Configuring RFC-Compliant NETCONF Sessions.](#)]

### *User Interface and Configuration*

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

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

## What's Changed in Release 20.2R2

### *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 (PSUs) 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.](#)]

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**—Starting in this release, the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group is renamed simply **arp**. This packet type option enables you to change default control plane DDoS protection policer parameters for ARP traffic. After this change, the **edit system ddos-protection protocols arp** protocol group includes **aggregate**, **arp**, and **unclassified** packet type options.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\).](#)]

- **PTX10001-36MR, PTX10008, and PTX10016 routers support a maximum of two drop profile pairs (PTX Series)**—Pair one drop probability must be less than or equal to 25%. Pair two drop probability value must be greater than point one drop probability value. Pair two fill level must be greater than or equal to 1.2 times the pair one fill level.

[See [CoS Features and Limitations on PTX Series Routers.](#)]

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



## 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](#).]

## Routing Protocols

- **Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.

## SEE ALSO

[What's New | 205](#)

[Known Limitations | 216](#)

[Open Issues | 218](#)

[Resolved Issues | 221](#)

[Documentation Updates | 226](#)

[Migration, Upgrade, and Downgrade Instructions | 226](#)

## Known Limitations

### IN THIS SECTION

- [General Routing | 217](#)
- [MPLS | 217](#)
- [Routing Protocols | 217](#)

Learn about known limitations in Junos OS Release 20.2R3 for PTX Series routers.

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

## General Routing

- On the PTX10008 or PTX10016 routers, the GRES takes more than 3 minutes to complete when shutdown is initiated by the internal `vmhost init 0` command. [PR1312065](#)
- The filter-based GRE encapsulation does not work in the egress direction when the filter attachment interface and the interface to reach the next hop are the same. [PR1465837](#)
- During reconfigurations and link events at the physical interface level, the `pe.ipw.misc_int.status:iq_disabled` error message can be seen. This does not impact traffic. [PR1476553](#)
- The `sflow record` command shows incorrect output interface for the egress sampling during the incoming MPLS|IPv4 and outgoing IPv4 with ECMP. [PR1478012](#)
- The PTX10000 routers include the incoming MPLS label stack length also in the jvision counters when acting as the PE device egress counter. [PR1482408](#)
- 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`. [PR1485770](#)
- The following error messages are seen after configuring `set chassis maximum-ecmp 64`:  
`JPRDS_NH:jprds_nh_alloc(),990: JNH[3] failed to grab new region for EGRESS`. [PR1490813](#)
- The `show dynamic-tunnels database statistics <dest>` command must be structured so that the statistics are fetched deterministically for the IPv4 and IPv6 based tunnels. [PR1488715](#)

## MPLS

- Traffic outage during FRR is observed with ingress node logs data errors. [PR1430361](#)

## Routing Protocols

- Router receives and discards traffic for three-and-a-half minutes after bootup when IGP overload is configured. [PR1495435](#)

## SEE ALSO

---

[What's New | 205](#)

---

[What's Changed | 213](#)

---

[Open Issues | 218](#)

---

[Resolved Issues | 221](#)

[Documentation Updates | 226](#)

[Migration, Upgrade, and Downgrade Instructions | 226](#)

## Open Issues

### IN THIS SECTION

- [General Routing | 218](#)
- [Infrastructure | 220](#)
- [Interfaces and Chassis | 220](#)
- [Layer 2 Ethernet Services | 220](#)
- [MPLS | 220](#)
- [Routing Protocols | 220](#)

Learn about open issues in the Junos OS Release 20.2R3 for PTX Series routers.

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

### General Routing

- When CFP2-DCO-T-WDM-1 is plugged in to a PTX Series PIC, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)
- On 30-port MACsec-enabled line card (LC1101-M-30C, LC1101-M-30Q, and LC1101-M-96X) of the PTX10008 chassis, when the exclude-protocol lacp statement configured at the [edit security macsec connectivity-association connectivity-association-name] hierarchy level is deleted or deactivated, the LACP protocol's Mux State shown under the output of CLI command show lacp interface, might remain as attached or detached and might not change to distributing state. [PR1331412](#)
- The PTX Series platform drops the wireless access point (WAP) heartbeat packets; as a result, the WAP cannot work. [PR1352805](#)
- Due to transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. This change applies to all platforms using Hybrid Memory Controller (HMC). [PR1384435](#)

- On the PTX10000 Series platform, the CPU overuse on priority-based flow control might be observed if the adaptive feature is enabled to load-balance for an aggregated Ethernet interface. [PR1399369](#)
- On the PTX3000 routers, the firewall counter for lo0 does not increment. [PR1420560](#)
- The em2 interface configuration causes FPC to crash during initialization and FPC does not come online. After deleting the em2 configuration and restarting the router, FPC comes online. [PR1429212](#)
- When the firewall filter has Port-Mirror as an action along with discard action, the mirrored packet will have two L2 headers. The first L2 header will be the original L2 header and the second L2 header will be egress interface L2 header. This causes packet corruption and discard. [PR1437546](#)
- Memory leaks are expected in this release. [PR1438358](#)
- On Junos OS platforms with next generation Routing Engine installed, the process vhostd may crash without the core file and automatic restart of vhostd may fail. Vhostd is a mib2d MIB II process for managing the lifecycle of system-critical Junos OS VMs in the system. If the process vhostd gets in a crash state, it will impact the management of Junos OS VMs. [PR1448413](#)
- The Layer 2 VPN with asynchronous-notification might flap when the link goes up between the PE device and CE device. After Layer 2 VPN flaps, the interfaces with asynchronous-notification might show - Inf dBm laser output power even if the Layer 2 VPN is in the up status. [PR1486181](#)
- Traceroute on IPoIP tunnel might not work if decap and encap routes are present in two different routing instances. [PR1488379](#)
- On PTX1000 and PTX10001 routers, the port mirror will not work when the port-mirroring is configured with the firewall filter. [PR1491789](#)
- Dynamic-tunnels traceoptions might cause scheduler slips with single underlay route bounce for large scale. [PR1493236](#)
- MPLS sensor does not receive Junos Telemetry Interface data on the server. [PR1514959](#)
- When you continuously run the sync (using the show interfaces aex extensive command) and the async (using SNMP polling) queries in parallel on aggregated Ethernet interfaces, you might notice spikes in aggregated Ethernet interface framing errors counter in between correct values. [PR1539537](#)
- Problem: A. Affected Platform and PIC Type "15x100GE/15x40GE/60x10GE QSFP28 PIC" in PTX5000 Chassis (No other PIC type was affected.) B. Symptom: When the concerned port is configured in 4x10G mode (using QSFP+), if one of the 10G channel detected a clear of Rx LOS (Loss Of Signal), all other 3 10G channels would experience a short link flapping. The interface links would become DOWN, and UP after several seconds. Affected Software Releases: Junos Release 20.2R1 and later releases. Root Cause: The concerned symptom was due to a hidden bug of BRCM re-timer (BCM82381) new firmware that was upgraded on the concerned PIC SW. On the event of Rx LOS CLEAR, the SW for the concerned PIC type would perform an enable/disable line-side digital loopback on the BCM82381 device's 10G channel. This caused a Tx signal glitch of other 3 sibling channels on the same re-timer device. Workaround: For the affected software releases, avoid configuring the ports on this PIC type in 4x10G mode. [PR1578511](#)

## Infrastructure

- Memory corruption of a binary from /usr/bin/ or /usr/sbin/ directory can occur if such binary is invoked when a recovery snapshot creation is in progress. The exact symptoms will be different depending on the exact binary and JUNOS version - some programs will show an error, and some programs will crash every time it is executed. Such memory corruption will be persistent until the affected routing engine is restarted. [PR1563647](#)

## Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with a CFM configuration might cause the cfmd process to crash after the upgrade. This is because of the presence of an old version of /var/db/cfm.db. [PR1281073](#)

## Layer 2 Ethernet Services

- It is observed rarely that issuing **request system zeroize** did not trigger zero-touch provisioning. A workaround is to re-initiate the ZTP. [PR1529246](#)

## MPLS

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

## Routing Protocols

- With an aggregated Ethernet interface with BFD configured, the aggregated Ethernet interface and BFD session remain down after the interface is disabled or enabled. [PR1354409](#)
- The show dynamic-tunnels database command does not show the current value of traffic statistics. It shows the cached value of traffic statistics, which might not be equal to the current value. [PR1445705](#)

## SEE ALSO

---

[What's New | 205](#)

---

[What's Changed | 213](#)

---

[Known Limitations | 216](#)

---

[Resolved Issues | 221](#)

---

[Documentation Updates | 226](#)

---

## Resolved Issues

### IN THIS SECTION

- Resolved Issues: 20.2R3 | [221](#)
- Resolved Issues: 20.2R2 | [223](#)
- Resolved Issues: 20.2R1 | [224](#)

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

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

### Resolved Issues: 20.2R3

#### General Routing

- Flexible PIC concentrator reboot might be observed in the events of J-Lock hog for more than 5 seconds. [PR1439929](#)
- On PTX10016 routers, if aggregated Ethernet member or interface flow control is in disabled state, then it does not enable its own. [PR1478715](#)
- SNMP index in the Packet Forwarding Engine reports as 0. This causes the sFlow records to have either IIF (Input interface value) or OIF (Output interface value) as 0 value in sFlow record data at collector. [PR1484322](#)
- PTX10008: FPC UKERN core dump is not transferred to Routing Engine in scaled setup. [PR1500418](#)
- Error messages `t6e_dfe_tuning_state:et-6/0/0 - Failed to dfe tuning count 10` might be seen after links flap [PR1512919](#)
- The chassisd memory leak might cause traffic loss. [PR1537194](#)
- The error message `expr_dfw_action_topo_connect_anh:1434 expr_dfw_action_topo_connect_anh:eda_anh_discard is FALSE for nh-id 568 - return` is observed in PTX1000 routers. [PR1540064](#)
- The Packet Forwarding Engine might crash in an MPLS IPv6-tunneling scenario when the next hop changes. [PR1540793](#)

- Traffic might drop silently after swapping an FPC Type 3 card with an FPC Type 1 card in the same slot on a PTX3000 router. [PR1547790](#)
- The rpd crash might be seen when BGP service route is resolved over color-only SR-TE. [PR1550736](#)
- Interface filter with source-port 0 is matching everything instead of just port 0. [PR1551305](#)
- An enhancement to enable watchdog petting log on PTX10000 Line Cards [PR1561980](#)
- The system might become unresponsive due to the out of memory issue when the BMP rib-out monitoring is enabled for large-scale BGP peers. [PR1477212](#)
- The following SNMP traps are not generated: **Power Supply failed** and **Power Supply OK**.[PR1479133](#)
- The BGP BMP session with connection passive is not established with the mgmt\_junos routing instance. [PR1507974](#)

### ***Forwarding and Sampling***

- l2ald might crash due to next-hop issue in the EVPN-MPLS. [PR1548124](#)

### ***Infrastructure***

- Output drops in **show interfaces extensive** might display 0 temporarily during a race condition when SNMP query for JxCos is also issued [PR1533314](#)
- Interface drop counters might display 0 during a race condition when VOQ statistics are also polled simultaneously. [PR1537960](#)
- The kernel crash with core file might be seen if churn happens for a flood composite next hop [PR1548545](#)

### ***Interfaces and Chassis***

- EOAM IEEE802.3ah link discovery state is **Down** instead of **Active Send Local** after deactivating interfaces on routers. [PR1532979](#)
- Logs are not being written in /var/log/messages on certain PTX platforms.[PR1551374](#)

### ***Network Management and Monitoring***

- The syslog messages might not be sent with the correct port number info. [PR1545829](#)

### ***Platform and Infrastructure***

- The BGP session replication might fail to start after the session crashes on a backup Routing Engine. [PR1552603](#)

### ***Routing Policy and Firewall Filters***

- Generate route goes to hidden state when protect core statement is enabled. [PR1562867](#)

### ***Routing Protocols***

- Traffic might be silently discarded when the BGP route gets deleted which is a part of multipath [PR1514966](#)

- The rpd process generates the core file at `gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash`. [PR1541768](#)
- BGP LU session flap might be seen with AIGP used scenario [PR1558102](#)

## Resolved Issues: 20.2R2

### General Routing

- On PTX5000 and PTX10008 routers, the output of the `show filter index number counter` command shows value as zero at `28-02-HOSTBOUND_NDP_DISCARD_TERM`. [PR1420057](#)
- The `show snmp mib walk jnxContentsDescr` command output does not show the fan controllers. [PR1455640](#)
- On PTX10016 routers, after device reboot, the FPC takes a long time to come up and hence MKA session 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](#)
- Any change in nested groups might not be detected on commit and does not take effect. [PR1484801](#)
- Outbound SSH connection flaps or a memory leak issue is observed during the push configuration to the ephemeral database with a high rate. [PR1497575](#)
- The error message `mpls_extra NULL` might be seen when you add, change, or delete MPLS route. [PR1502385](#)
- 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](#)
- ERO update by the controller for branch LSP might cause issues. [PR1508412](#)
- On PTX3000 and PTX5000 routers, unable to bring the ports up when plugging in the optic QSFP-100G-LR4-T2 (740-061409). [PR1511492](#)
- The route update might fail because of an HMC memory issue and traffic impact might be seen. [PR1515092](#)
- On PTX1000 and PTX10002-60C routers, sFlow adaptive-sampling, with rate limiter statement enabled, crosses the sampling rate 65535. [PR1525589](#)

### Interfaces and Chassis

- When multiple CFM sessions are configured on a physical interface, SNMP walk of `ieee8021CFMStack` table fails. [PR1517046](#)
- EOAM IEEE802.3ah link discovery state is **Down** instead of **Active Send Local** after deactivating interfaces on routers. [PR1532979](#)



## MPLS

- SNMP trap is observed with incorrect OID jnxSpSvcSetZoneEntered. [PR1517667](#)

## Routing Protocols

- On PTX3000 and PTX5000 routers, the ppm process generates a core file after configuring the S-BFD responder on the RE-DUO-2600. [PR1477525](#)
- The rpd process might report 100 percent CPU usage with BGP route damping enabled. [PR1514635](#)

## Resolved Issues: 20.2R1

### General Routing

- PTX interface stays down after the maintenance. [PR1412126](#)
- With Junos OS Release 19.4R1 on PTX10008 device along with 4x1GE feature, continuous logging in the chassisd file is observed. [PR1456253](#)
- Upgrading fails due to communication failure between the Junos VM and host OS. [PR1438219](#)
- The local-loopback test fails with the together options. [PR1458814](#)
- The PTX1000 or PTX10002 router might discard traffic silently after the transient SIB or FPC voltage alarms. [PR1460406](#)
- On the PTX5000 for FPC3, optics-options syslog and link-down do not work as expected. [PR1461404](#)
- The sample, syslog, or log action in the output firewall filter with packet size less than 128 might cause ASIC wedge (all packet loss). [PR1462634](#)
- On modifying TNL DST NETWORK (more specific TNL DST NETWORK), the IP-IP tunnel gets flushed but fails to get created even though a less specific matching TNL DST NETWORK exists. [PR1462805](#)
- On the PTX10000 line of routers, FPC might restart during runtime. [PR1464119](#)
- The PTX5000 SIB3 might fail to come up in the slot 0 with or without slot 8 when the Routing Engine 1 is the master. [PR1471178](#)
- The input-vlan-map or output-vlan-map might not work properly in the Layer 2 circuit local-switching scenario. [PR1474876](#)
- Sampling process might crash when the MPLS or MPLS over the UDP traffic is sampled. [PR1477445](#)
- Multicast routes add or delete events might cause adjacency and LSPs to go down. [PR1479789](#)
- FPC might crash when dealing with the invalid next hops. [PR1484255](#)
- In the StrictPriority mode, the MedH and MedL should be of separate priorities; StrcH and High become one priority. [PR1490505](#)
- The BFD sessions flap when the firewall filter in the loopback0 is changed. [PR1491575](#)

- Traffic impact might be seen when policy-multipath is configured without LDP on the Spring-TE scenario. [PR1483585](#)
- On a dual Routing Engine GRES or NSR enabled PTX10008 or PTX10016 router, a few TCP-based application sessions like BGP or LDP might flap upon Routing Engine primary-role switch. [PR1503169](#)
- The router might become nonresponsive and bring traffic down when the disk space becomes full. [PR1470217](#)
- Unable to bring the ports up when plugging the optic QSFP-100G-LR4-T2(740-061409) to PTX3000 or PTX5000. [PR1511492](#)
- PHP device has NH mis-programming for members of ECMP for SR label route used for reaching the IPV6 destinations. [PR1457230](#)
- Kernel Routing Table (KRT) queue gets stuck after the J-Flow samples a malformed packet. [PR1495788](#)

### **Infrastructure**

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

### **Layer 2 Ethernet Services**

- Member links state might be asynchronized on a connection between the PE device and the CE devices in the EVPN A/A scenario. [PR1463791](#)

### **MPLS**

- Kernel crash and device restart might occur. [PR1478806](#)
- The BGP session might keep flapping between two directly connected BGP peers because of the wrong usage of the TCP-MSS. [PR1493431](#)
- The rpd process might crash in a rare condition under the SR-TE scenario. [PR1493721](#)

### **Routing Protocols**

- The BGP NSR must be able to synchronize 4000 or more IPv6 sessions. [PR1461436](#)
- On the PTX3000 or PTX5000 line of routers, the ppmd process generates a core file after configuring the sbfd responder on the RE-DUO-2600. [PR1477525](#)
- The rpd process might crash with the BGP multipath and route withdraw occasionally. [PR1481589](#)
- The BGP route-target family might prevent RR from reflecting Layer 2 VPN and Layer 3 VPN routes. [PR1492743](#)
- BGP multi-path traffic might not fully load-balance for a while after adding a new path for the load sharing. [PR1482209](#)
- LSP auto-bandwidth adjust-interval change does not get detected on commit in some cases. [PR1484801](#)

## SEE ALSO

[What's New | 205](#)[What's Changed | 213](#)[Known Limitations | 216](#)[Open Issues | 218](#)[Documentation Updates | 226](#)[Migration, Upgrade, and Downgrade Instructions | 226](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 documentation for PTX Series routers.

## SEE ALSO

[What's New | 205](#)[What's Changed | 213](#)[Known Limitations | 216](#)[Open Issues | 218](#)[Resolved Issues | 221](#)[Migration, Upgrade, and Downgrade Instructions | 226](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.2 | 227](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 229](#)
- [Upgrading a Router with Redundant Routing Engines | 230](#)

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 20.2

When upgrading or downgrading Junos OS, use the **jinstall** package. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#). Use other packages, such as the **bundle** 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 20.2R3:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage:

<https://support.juniper.net/support/downloads/>

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.
4. 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-x86-64-20.2R3.9.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-x86-64-20.2R3.9-limited.tgz
```

Replace the source with one of the following values:

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

- `scp://hostname/pathname`

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

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

Rebooting occurs only if the upgrade is successful.

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

**NOTE:** After you install a Junos OS Release 20.2 `jinstall` package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the `jinstall` package that corresponds to the previously installed software.

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

## Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from

Junos OS Release 19.3 to Release 20.1. 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 | 205](#)

---

[What's Changed | 213](#)

---

[Known Limitations | 216](#)

---

[Open Issues | 218](#)

---

[Resolved Issues | 221](#)

---

[Documentation Updates | 226](#)

# Junos OS Release Notes for the QFX Series

## IN THIS SECTION

- [What's New | 231](#)
- [What's Changed | 256](#)
- [Known Limitations | 260](#)
- [Open Issues | 262](#)
- [Resolved Issues | 267](#)
- [Documentation Updates | 280](#)
- [Migration, Upgrade, and Downgrade Instructions | 281](#)

These release notes accompany Junos OS Release 20.2R3 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

### IN THIS SECTION

- [What's New in Release 20.2R3 | 232](#)
- [What's New in Release 20.2R2 | 232](#)
- [What's New in Release 20.2R1-S1 | 232](#)
- [What's New in Release 20.2R1 | 234](#)

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 20.2R3: QFX5100, QFX5110 (32Q and 48S), QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.

Junos on White Box runs on Accton Edgecore AS7816-64X switches in this release. The software is based on Junos OS running on QFX5210 switches, so release-note items that apply to QFX5210 switches also apply to Junos on White Box.

### What's New in Release 20.2R3

There are no new features or enhancements to existing features for QFX Series Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features or enhancements to existing features for QFX Series Junos OS Release 20.2R2.

### What's New in Release 20.2R1-S1

#### *Flow-Based and Packet-Based Processing*

- **Support for user-defined flex hashing for MPLS traffic flows (QFX5210; Accton AS7816 running Junos OS on White Box)**—Starting in Junos OS Release 20.2R1-S1, you can configure user-defined flex hashing to load balance MPLS traffic based on TCP or UDP source/destination port information. User-defined flex hashing, which supports protocol versions IPv4 and IPv6, enables you to set byte offsets in packet headers to influence hashing computation. You specify two offsets, each 2 bytes in length, from the first 128 bytes of a packet. Configure the selected bytes to be directly used for hashing or to be used only when the data pattern in these bytes matches with specific values (conditional match). To provide load balancing in spine layers, configure flex hashing and encapsulate the traffic in VXLAN, thus enabling entropy at UDP source ports. At de-encapsulation, configure the **no-inner-payload** statement to load balance based on the outer UDP header.

To configure user-defined flex hashing:

```
set forwarding-options enhanced-hash-key flex-hashing name ethtype mpls num_labels source-port hash-offset
offset1 base_offset1 offset1_value offset1_mask offset2 base_offset2 offset2_value offset2_mask
```

To configure a conditional match (repeat the command below with values for offsets and match data 2-4):

```
set forwarding-options enhanced-hash-key conditional-match name offset1 base_offset1 offset1_value
matchdata1 matchdata1_mask
```

To enable load balancing on VXLAN transit traffic based on the outer UDP header:

```
set forwarding-options enhanced-hash-key vxlan no-inner-payload
```

To troubleshoot, use **show forwarding-options enhanced-hash-key**.

Limitations:

- Use a maximum of two MPLS labels.
- Use only even values for **offset1** and **offset2**.
- If you are using conditional matches, configure the conditions before you attach them to the flex-hashing entry.
- An aggregated Ethernet (AE), or LAG, interface is not supported as an input interface. You can configure input interfaces on LAGs by configuring the same user-defined flex-hashing data and the same conditional-match data on all *member* interfaces of a LAG interface. Use unique flex-data profile names and unique conditional-data profile names for each member interface—for example:
  - ...enhanced-hash-key conditional-match COND\_L1\_V6\_UDP\_SRC\_PORT\_1...
  - ...enhanced-hash-key conditional-match COND\_L1\_V6\_UDP\_SRC\_PORT\_2...

### *Software Installation and Upgrade*

- **Zero touch provisioning (ZTP) with IPv6 support (EX3400, EX4300, QFX5100 and QFX5200 switches, MX-Series routers)**—Starting in Junos OS Release 20.2R1-S1, ZTP supports the DHCPv6 client. During the bootstrap process, the device first uses the DHCPv4 client to request for information regarding image and configuration file from the DHCP server. The device checks the DHCPv4 bindings sequentially. If there is a failure with one of the DHCPv4 bindings, the device will continue to check for bindings until provisioning is successful. If there are no DHCPv4 bindings, however, the device will check for DHCPv6 bindings and follow the same process as for DHCPv4 until the device can be provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

**NOTE:** Only HTTP and HTTPS transport protocols are supported on EX3400, EX4300, QFX5100, and QFX5200 devices.

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

## **What's New in Release 20.2R1**

### *Hardware*

- **New QFX5120-48T Ethernet Switch (QFX Series)**—Starting with Junos OS Release 20.2R1, the QFX5120-48T is a 10GbE/100GbE data center switch offering 48 10GbE RJ-45 ports and six 40GbE/100GbE QSFP28/QFSP+ ports. The 48 copper ports support 1-Gbps and 10-Gbps speeds and the last 6 ports (port 48 to 53) support 40-Gbps and 100-Gbps speeds. By default, the first 48 ports operate at 10-Gbps speed and the last six ports 100-Gbps speed.

QFX5120-48T switches supports both manual and auto-channelization, but manual CLI channelization always takes precedence. [See [Port Settings](#).]

To install the QFX5120-48T switch hardware and perform initial software configuration, routine maintenance, and troubleshooting, see the [QFX5120 Switch Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform.

[Table 2 on page 235](#) summarizes the software features supported in this release.

**Table 2: Features Supported by QFX5120-48T Switches**

Feature	Description
Authentication and Access Control	<ul style="list-style-type: none"> <li>• IEEE 802.1X authentication support. [See <a href="#">User Access and Authentication User Guide</a>.]</li> <li>• IP source guard. [See <a href="#">Configuring IP Source Guard (ELS)</a>.]</li> <li>• Local password authentication support for password change policy.</li> <li>• Storm control support (broadcast, unicast, and multicast). [See <a href="#">Understanding Storm Control</a>.]</li> <li>• Radius and TACACS+ authentication. [See <a href="#">Authentication Order for RADIUS, TACACS+, and Local Password</a>.]</li> <li>• Role-based access control (RBAC), and role-based CLI management.</li> </ul>
BGP	<ul style="list-style-type: none"> <li>• Support for BGP Monitoring Protocol (BMP) Version 3 and IPv6 BGP standards. [See <a href="#">Understanding the BGP Monitoring Protocol</a> and <a href="#">Supported IPv6 Standards</a>.]</li> <li>• BGP advertising aggregate bandwidth across external BGP links for load balancing. [See <a href="#">Load Balancing for a BGP Session</a>.]</li> <li>• Support for BGP large communities, link-state distribution, multipath at global level, and support for 4-byte autonomous system numbers. [See <a href="#">Routing Policies for BGP Communities</a>.]</li> <li>• EBGP route support, multiprotocol BGP (MBGP) extensions, and frequent BGP keepalive messages with a short BGP hold time. [See <a href="#">BGP Overview</a>.]</li> <li>• Routing protocol process (rpd) recursive resolution over multipath. [See <a href="#">BGP Overview</a>.]</li> <li>• BGP labeled-unicast. [See <a href="#">labeled-unicast (Protocols BGP)</a>.]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
Class of Service	<ul style="list-style-type: none"> <li>● Standard class of service (CoS) feature support including configuring classification, rewrite, queuing, shaping, buffering, and scheduling parameters for traffic management. [See <a href="#">CoS Support on QFX Series Switches</a>.]</li> <li>● IEEE 802.1p rewrite and classification.</li> <li>● Class-based queuing with prioritization. [See <a href="#">Understanding CoS Output Queue Schedulers</a>.]</li> <li>● Single-rate two-color marking, single-rate three-color marking, and two-rate three-color marking. [See <a href="#">Overview of Policers</a>.]</li> <li>● Separate unicast and multi-destination classifiers, forwarding classes, and output queues. [See <a href="#">Understanding Junos CoS Components</a>.]</li> <li>● Direct port scheduling. [See <a href="#">Understanding CoS Port Schedulers on QFX Switches</a>.]</li> <li>● Queue shaping using the shaping-rate statement. [See <a href="#">Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)</a>.]</li> <li>● Priority-based flow control (PFC) with 802.3x Ethernet PAUSE and explicit congestion notification (ECN). [See <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC)</a> and <a href="#">Understanding CoS Explicit Congestion Notification</a>.]</li> <li>● CoS support for link aggregation groups (LAGs).</li> <li>● Weighted random early detection (WRED) packet drop profiles and tail drop. [See <a href="#">Understanding CoS Congestion Management and Understanding CoS WRED Drop Profiles</a>.]</li> <li>● Rewrite rule (marking) of bridged packets. [See <a href="#">Understanding Junos CoS Components</a>.]</li> <li>● Policing or rate limiting of traffic to apply limits to traffic flow. [See <a href="#">Overview of Policers</a>.]</li> </ul>
DHCP	<ul style="list-style-type: none"> <li>● Client link-layer address option 79 for DHCPv6. [See <a href="#">mac-address (DHCP Relay Agent)</a>.]</li> <li>● DHCP server, DHCP smart relay configuration, DHCP relay with DHCP server, and DHCP client in separate routing instances. [See <a href="#">DHCP Message Exchange Between DHCP Clients and DHCP Server in Different Virtual Routing Instances</a>.]</li> <li>● DHCP relay with option 82 for Layer 2 VLANs and Layer 3 interface. [See <a href="#">DHCP Relay Agent Information Option (Option 82)</a>.]</li> <li>● DHCP and DHCPv6 snooping. [See <a href="#">DHCP Snooping</a>.]</li> <li>● DHCP static addresses. [See <a href="#">Configuring Static DHCP IP Addresses</a>.]</li> <li>● Extended DHCP (also referred to as virtual router (VR) aware DHCP). [See <a href="#">Legacy DHCP and Extended DHCP</a>.]</li> <li>● Textual interface description using DHCP relay agent option 82 (circuit ID). [See <a href="#">DHCP Relay Agent Information Option (Option 82)</a>.]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (continued)

Feature	Description
EVPN and VXLAN	<ul style="list-style-type: none"> <li>• EVPN proxy ARP and ARP suppression. [See <a href="#">EVPN Proxy ARP and ARP Suppression Proxy</a>.]</li> <li>• EVPN control plane and VXLAN data plane support. [See <a href="#">Understanding EVPN with VXLAN Data Plane Encapsulation</a>.]</li> <li>• EVPN pure type-5 route support. [See <a href="#">EVPN Type-5 Route with VXLAN encapsulation for EVPN-VXLAN</a>.]</li> <li>• LACP in EVPN active-active multihoming. [See <a href="#">Example: Configuring LACP for EVPN VXLAN Active-Active Multihoming</a>.]</li> <li>• Automatically generated Ethernet segment identifiers in EVPN-VXLAN and EVPN-MPLS networks. [See <a href="#">Understanding Automatically Generated and Assigned ESIs in EVPN Networks</a>.]</li> <li>• EVPN-VXLAN support of Virtual Chassis and Virtual Chassis Fabric. [See <a href="#">Integrating a Virtual Chassis Fabric into an EVPN-VXLAN Environment</a>.]</li> <li>• Support for VMTO for ingress traffic. [See <a href="#">Configuring EVPN Routing Instances</a>.]</li> <li>• MAC filtering, storm control, and port mirroring support in EVPN-VXLAN overlay networks. [See <a href="#">MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment</a>.]</li> <li>• Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface. See <a href="#">Understanding Flexible Ethernet Services Support With EVPN-VXLAN</a>.]</li> <li>• Support for multihomed proxy advertisement. [See <a href="#">EVPN Multihoming Overview</a>.]</li> <li>• Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network. [See <a href="#">Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network</a>.]</li> <li>• Support for graceful restart and graceful restart protocol extension support for unicast and type 5 messages on EVPN-VXLAN. [See <a href="#">Graceful Restart in EVPN</a>.]</li> <li>• Standard class-of-service (CoS) features—classifiers, rewrite rules, and schedulers are supported on VXLAN interfaces. [See <a href="#">Understanding CoS on OVSDB-Managed VXLAN Interfaces</a>.]</li> <li>• Firewall filtering and policing on EVPN-VXLAN traffic. [See <a href="#">Understanding VXLANs and Overview of Firewall Filters</a>.]</li> <li>• Configurable VXLAN UDP port.</li> <li>• Support for IGMP snooping for EVPN-VXLAN in a multihomed environment. [See <a href="#">Overview of Multicast Forwarding with IGMP Snooping in an EVPN-VXLAN Environment</a>.]</li> <li>• Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks. [See <a href="#">Supported Protocols on an IRB Interface in EVPN-VXLAN</a> .]</li> <li>• VXLAN Layer 2 gateway (static, OVSDB, EVPN), Q-in-Q tag manipulation, dynamic load balance, and hashing options. [See <a href="#">OVSDB-VXLAN User Guide for QFX Series Switches</a>.]</li> <li>• BPDU protection in EVPN-VXLAN. [See <a href="#">Supported Protocols on an IRB Interface in EVPN-VXLAN</a>.]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
Firewall Filters and Policers	<ul style="list-style-type: none"> <li>• Support for firewall filters on interfaces, VLANs, routed VLAN interfaces (RVIs), link aggregation groups (LAGs), and loopback interfaces. [See <a href="#">Overview of Firewall Filters.</a>]</li> <li>• Single-rate two-color marking, single-rate three-color marking, and two-rate three-color marking. [See <a href="#">Overview of Policers.</a>]</li> <li>• Dynamic allocation of firewall filters.</li> <li>• Enhanced filter classification of CPU-generated packets.</li> <li>• Firewall filter actions. [See <a href="#">Firewall Filter Match Conditions and Actions (QFX and EX Series Switches.)</a>]</li> <li>• Firewall filter flexible match conditions and firewall filters on loopback and management interface. [See <a href="#">Firewall Filter Flexible Match Conditions.</a>]</li> <li>• Port firewall filters (egress and ingress) and routed firewall filters (egress and ingress). [See <a href="#">Firewall Filter Match Conditions and Actions (QFX and EX Series Switches.)</a>]</li> <li>• VLAN firewall filters (egress and ingress). [See <a href="#">Firewall Filter Match Conditions and Actions (QFX and EX Series Switches.)</a>]</li> <li>• TCP/UDP port ranges in classification. [See <a href="#">Firewall Filter Match Conditions and Actions (QFX and EX Series Switches.)</a>]</li> <li>• Filter-based GRE de-encapsulation. [See <a href="#">Configuring a Firewall Filter to De-Encapsulate GRE Traffic.</a>]</li> <li>• Loopback firewall filter scale optimization. [See <a href="#">Planning the Number of Firewall Filters to Create.</a>]</li> </ul>
High Availability (HA) and Resiliency	<ul style="list-style-type: none"> <li>• Automatic recovery for port error disable condition. [See <a href="#">disable-timeout (Port Error Disable).</a>]</li> <li>• Operating system resiliency to recover the Junos OS software using device recovery mode. [See <a href="#">Rescue Configuration.</a>]</li> <li>• Partial resiliency for errors, machine-check exception (MCE), and advanced error reporting (AER).</li> <li>• Ethernet ring protection switching (ERPS). [See <a href="#">Ethernet Ring Protection Switching Overview.</a>]</li> <li>• Graceful protocol restart for BGP and OSPF. [See <a href="#">Understanding Graceful Restart for BGP, graceful-restart (Protocols BGP)</a> and <a href="#">Configuring Graceful Restart for OSPF.</a>]</li> <li>• Nonstop software upgrade (NSSU), Nonstop bridging, and Nonstop active routing (NSR) for IPv6 and OSPFv2.</li> <li>• Virtual Chassis support. [See <a href="#">Understanding QFX Series Virtual Chassis.</a>]</li> <li>• Virtual Chassis with NSSU support. You can interconnect two QFX5120-48T switches into a Virtual Chassis that operates as one logical device managed as a single chassis. [See <a href="#">Virtual Chassis Overview for Switches.</a>]</li> <li>• Network Device Collaborative Protection Profile (NDcPP) certification.</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
Interfaces and Chassis	<ul style="list-style-type: none"> <li>● Dynamic ARP inspection (DAI) and static ARP support. [See <a href="#">Understanding and Using Dynamic ARP Inspection (DAI)</a>.]</li> <li>● Support for dynamic load balancing. [See <a href="#">Understanding Load Balancing for Aggregated Ethernet Interfaces</a>.]</li> <li>● Proxy ARP per VLAN and unrestricted proxy ARP. [See <a href="#">Restricted and Unrestricted Proxy ARP Overview</a>.]</li> <li>● Link protection support on aggregated Ethernet interfaces and updated behavior in static link protection mode.</li> <li>● Automatic detection of MDI and MDIX port connections. Auto MDI/MDIX is enabled by default. [See <a href="#">no-auto-mdix</a>.]</li> <li>● Digital optical monitoring (DOM). [See <a href="#">show interfaces diagnostics optics</a>.]</li> <li>● Support for fiber channel over Ethernet (FCoE), FCoE initialization protocol (FIP), FIP snooping, and up to 2500 total FIP snooping sessions supported on an interface. [See <a href="#">Understanding VN_Port to VF_Port FIP Snooping on an FCoE Transit Switch</a>.]</li> <li>● Filter-based GRE decapsulation.</li> <li>● IPv4 generic routing encapsulation (GRE) support. [See <a href="#">Configuring Generic Routing Encapsulation Tunneling</a>.]</li> <li>● Auto-negotiation and port speed. [See <a href="#">auto-negotiation</a>.]</li> <li>● Configure speed of Gigabit Ethernet copper SFP interfaces. [See <a href="#">Gigabit Ethernet Interface</a>.]</li> <li>● IEEE 802.3ah link fault management (LFM). [See <a href="#">OAM Link Fault Management</a>.]</li> <li>● Interface ranges. [See <a href="#">Interface Ranges</a>.]</li> <li>● Jumbo frames (up to 9216 bytes) and jumbo frames on routed VLAN interfaces (RVIs). [See <a href="#">Configuring Routed VLAN Interfaces on Switches (CLI Procedure)</a>.]</li> <li>● Layer 3 logical interfaces. [See <a href="#">Layer 3 Logical Interfaces</a>.]</li> <li>● Support for network-to-network interface (NNI) and user network interface (UNI) on the same physical interface. [See <a href="#">Configuring Q-in-Q Tunneling</a>.]</li> <li>● Channelizing Ethernet interfaces. [See <a href="#">Channelizing Interfaces Overview</a>.]</li> <li>● Dynamic port swap from 40G to 100G without restarting the Packet Forwarding Engine.</li> <li>● PVLAN and Q-in-Q on the same interface. [See <a href="#">Configuring Q-in-Q Tunneling on QFX Series Switches</a>.]</li> <li>● Link aggregation static and dynamic with LACP (fast and slow LACP), LLDP, and MC-LAG with configuration sync.</li> <li>● Uplink failure detection debounce interval. [See <a href="#">Uplink Failure Detection</a>.]</li> </ul>



Table 2: Features Supported by QFX5120-48T Switches (continued)

Feature	Description
IPv6	<ul style="list-style-type: none"> <li>● BGP support for advertising multiple paths to IPv6 addresses. [See <a href="#">Example: Advertising Multiple Paths in BGP</a>.]</li> <li>● Configure per-interface neighbor discovery protocol (NDP) cache protection. [See <a href="#">Neighbor Discovery Cache Protection Overview</a>.]</li> <li>● IPv6 specific SSH and Telnet.</li> <li>● Support for IPv6 filter-based forwarding. [See <a href="#">Understanding Filter-Based Forwarding</a>.]</li> <li>● Firewall filter support for IPv6 traffic: IPv6 fields for ingress port and VLAN firewall filters and policer action for MPLS firewall filters. [See <a href="#">Firewall Filter Match Conditions for IPv6 Traffic</a>.]</li> <li>● Support for IPv6 L3 forwarding, IPv6 Layer 3 VPNs, IPv6 traceroute, IPv6 tunneling, and IPv6 attributes in RADIUS message and stateless auto configuration.</li> <li>● Support for IPv6 OSPFv3, IPv6 ping, secure IPv6 neighbor discovery protocol (NDP), and IPv6 source guard. [See <a href="#">OSPF Version 3 for IPv6</a> and <a href="#">IPv6 Neighbor Discovery User Guide</a>.]</li> <li>● IPv6 access security (IPv6 neighbor discovery inspection, IPv6 stateless address auto-configuration (SLAAC) snooping, and understanding IPv6 router advertisement guard). [See <a href="#">IPv6 Neighbor Discovery Inspection</a>, <a href="#">IPv6 Stateless Address Auto-configuration (SLAAC) Snooping</a> and <a href="#">Understanding IPv6 Router Advertisement Guard</a>.]</li> <li>● Support for IPv6 over MPLS (6PE), IPv6 over MPLS LSPs, IPv6 static routing, IS-IS for IPv6, path MTU discovery, SNMP, NTP, and DNS. [See <a href="#">Configuring Junos OS for IPv6 Path MTU Discovery</a>.]</li> <li>● Virtual Router Redundancy Protocol (VRRP) and support for VRRP on IPv6 networks. [See <a href="#">VRRP and VRRP for IPv6 Overview</a>.]</li> </ul>
Junos OS XML API and Scripting	<ul style="list-style-type: none"> <li>● Scripts: Python, SLAX, and XSLT commit, event, op, SNMP, and open-source Python modules supported in automation enhancement.</li> <li>● Support for REST API interfaces.</li> <li>● JET for Junos: modern programmatic interface for developers of third-party applications. [See <a href="#">Understanding JET Interaction with Junos OS</a>.]</li> <li>● Configuration management: JSON format for configuration data. [See <a href="#">Defining the Format of Configuration Data to Upload in a Junos XML Protocol Session</a>.]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
Junos Telemetry Interface (JTI)	<ul style="list-style-type: none"> <li>● Support for the Junos Telemetry Interface [See. <a href="#">Understanding OpenConfig and gRPC.</a>]</li> <li>● Sensor level statistics support on Junos Telemetry Interface (JTI). [<a href="#">Guidelines for gRPC and gNMI Sensors.</a>]</li> <li>● gNMI support for routing engine statistics for JTI. [See <a href="#">Guidelines for gRPC and gNMI Sensors.</a>]</li> <li>● Enhancements to the sensor for BGP peer information.</li> <li>● Sensor for network discovery protocol (NDP) and Address Resolution Protocol table state information for IPv6 routes.</li> <li>● Sensor for memory utilization for routing protocol tasks. [See <a href="#">Guidelines for gRPC and gNMI Sensors.</a>]</li> <li>● Sensor for LSP events and properties, LSP statistics, and gRPC streaming for LSP statistics. [See <a href="#">Guidelines for gRPC and gNMI Sensors.</a>]</li> <li>● Packet Forwarding Engine statistics export using gNMI and JTI.</li> <li>● Aggregated Ethernet interfaces configured with the link aggregation control protocol (LACP), Ethernet interfaces configured with the link layer discovery protocol (LLDP), BGP peers, and RSVP interface events. [See <a href="#">Understanding OpenConfig and gRPC on Junos Telemetry Interface.</a>]</li> <li>● OpenConfig LLDP model (v0.1.0). [See <a href="#">OpenConfig Data Model Version.</a>]</li> <li>● OpenConfig to support operational models for VLANs.</li> <li>● OpenConfig Junos OS, OpenConfig, and Network Agent packages are delivered in a single TAR file. [See <a href="#">Installing the OpenConfig Package.</a>]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (continued)

Feature	Description
Layer 2 Features	<ul style="list-style-type: none"> <li>● Data center bridging (DCB) application protocol TLV exchange.</li> <li>● Data Center Bridging Capability Exchange Protocol (DCBX) version support for IEEE DCBX version 1.01. [See <a href="#">Understanding DCBX</a>.]</li> <li>● MAC address filtering, MAC table aging, and static MAC address assignment for interface. [See <a href="#">MAC Addresses</a> and <a href="#">MAC Table Aging</a>.]</li> <li>● Disable MAC learning, persistent MAC learning, MAC address limit per port, MAC limiting, MAC move limiting, MAC notification, and per VLAN (VLAN membership MAC limit). [See <a href="#">Understanding MAC Limiting and MAC Move Limiting for Port Security</a>.]</li> <li>● Enhanced Layer 2 Software (ELS). [See <a href="#">Layer 2 Networking</a>.]</li> <li>● IP directed broadcast traffic forwarding.</li> <li>● VLAN support, Link layer discovery protocol (LLDP), and Q-in-Q tunneling support. [See <a href="#">Configuring Q-in-Q Tunneling</a>.]</li> <li>● Static LAG link protection. [See <a href="#">link-protection (Static LSPs)</a>.]</li> <li>● Redundant trunk groups (link redundancy). [See <a href="#">Understanding Redundant Trunk Links (Legacy RTG Configuration)</a>.]</li> <li>● L2PT, UDLD, 802.1AE/802.1x, Ethernet Local Management Interface (E-LMI), and Multiple MAC Registration Protocol (MMRP). [See <a href="#">layer2-protocol-tunneling</a>.]</li> </ul>
Layer 3 Features	<ul style="list-style-type: none"> <li>● Configuring the GTP-TEID field for GTP traffic. [See <a href="#">Traffic Sampling, Forwarding, and Monitoring User Guide</a>.]</li> <li>● Equal-cost multipath (ECMP) flow-based forwarding: 64 ECMP paths. [See <a href="#">Traffic Sampling, Forwarding, and Monitoring User Guide</a>.]</li> <li>● Support to control traceroute over Layer 3 VPN.</li> <li>● Virtual routing and forwarding (VRF) support in IRB interfaces in a Layer 3 VPN.</li> <li>● Support for VRF-lite, BGP, IGMP, IS-IS, OSPF, PIM, and RIP.</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
MPLS	<ul style="list-style-type: none"> <li>● MPLS support for label edge routers (LER) and label switch routers (LSR). [See <a href="#">MPLS Overview for Switches</a>.]</li> <li>● Support for MPLS signaling protocols LDP and RSVP. [See <a href="#">LDP Overview</a> and <a href="#">RSVP Overview</a>.]</li> <li>● Fast reroute (FRR) support (a component of MPLS local protection for both one-to-one and many-to-one local protection).</li> <li>● Static LSPs. [See <a href="#">LSP Overview</a>.]</li> <li>● MPLS node protection, link protection, and statistics for static LSPs.</li> <li>● MPLS OAM (LSP ping).</li> <li>● MPLS statistics. [See <a href="#">statistics (Protocols MPLS)</a>.]</li> <li>● MPLS automatic bandwidth allocation and dynamic count sizing.</li> <li>● MPLS with RSVP-based LSPs.</li> <li>● Support for IRB interfaces over an MPLS core network. [See <a href="#">Example: Configuring IRB Interfaces on QFX5100 Switches over an MPLS Core Network</a>.]</li> <li>● MPLS stitching for virtual machine connections. [See <a href="#">Using MPLS Stitching with BGP to Connect Virtual Machines</a>.]</li> <li>● MPLS over Layer 3 subinterfaces. [See <a href="#">MPLS Limitations on QFX Series and EX4600 Switches</a>.]</li> <li>● Resource reservation protocol-traffic engineering (RSVP-TE), traffic engineering extensions (OSPF-TE, IS-IS-TE), Path Computation Element Protocol (PCEP), and PCE-initiated LSPs for the PCEP implementation. [See <a href="#">MPLS Applications User Guide</a>.]</li> <li>● Equal-cost multipath (ECMP) operation on MPLS using firewall filters.</li> </ul>
Multichassis Link Aggregation	<ul style="list-style-type: none"> <li>● Resilient hashing support for link aggregation group (LAG) routes. [See <a href="#">Resilient Hashing on LAGs and ECMP groups</a>.]</li> <li>● Keep a link up on a multichassis link aggregation group (MC-LAG) when LACP is not configured on one of the MC-LAG peers. [See <a href="#">Forcing MC-LAG Links or Interfaces with Limited LACP Capability to Be Up</a>.]</li> <li>● Layer 3 unicast and multicast support for MC-LAG. [See <a href="#">Advanced MC-LAG Concepts</a>.]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (continued)

Feature	Description
Network Management	<ul style="list-style-type: none"> <li>● IEEE 802.1ag OAM connectivity fault management. [See <a href="#">Understanding Ethernet OAM Connectivity Fault Management for Switches.</a>]</li> <li>● Port mirroring (local and remote) and remote port mirroring to IP address (GRE). [See <a href="#">Understanding Port Mirroring and Analyzers.</a>]</li> <li>● sFlow technology support. [See <a href="#">Understanding How to Use sFlow Technology for Network Monitoring on a Switch.</a>]</li> <li>● Chef for Junos OS support. [See <a href="#">Chef for Junos OS Getting Started Guide.</a>]</li> <li>● Puppet for Junos OS support. [See <a href="#">Puppet for Junos OS Administration Guide.</a>]</li> <li>● Adding non-native YANG modules to the Junos OS schema. [See <a href="#">Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS.</a>]</li> <li>● Enforcing RFC-compliant behavior in NETCONF sessions. [See <a href="#">Configuring RFC-Compliant NETCONF Sessions.</a>]</li> <li>● Configuring the ephemeral database using the NETCONF and Junos XML protocols. [See <a href="#">Committing an Instance of the Ephemeral Configuration Database Using the NETCONF or Junos XML Protocol.</a>]</li> <li>● Simple network management protocol (SNMP) remote monitoring (RMON) events, alarms, and history. [See <a href="#">SNMP MIB Explorer.</a>]</li> <li>● Real-time performance monitoring (RPM). [See <a href="#">Understanding Real-Time Performance Monitoring on Switches.</a>]</li> </ul>
Open vSwitch Database (OVSDB)	<ul style="list-style-type: none"> <li>● Automatic configuration of OVSDB-managed VXLANs with trunk interfaces. [See <a href="#">Understanding Dynamically Configured VXLANs in an OVSDB Environment.</a>]</li> <li>● BFD in a VMware NSX for vSphere environment with OVSDB and VXLAN. [See <a href="#">Understanding BFD in a VMware NSX Environment with OVSDB and VXLAN.</a>]</li> <li>● CoS on OVSDB-managed VXLAN interfaces. [See <a href="#">Configuring CoS on OVSDB-Managed VXLAN Interfaces.</a>]</li> <li>● Firewall filters on OVSDB-managed interfaces. [See <a href="#">Understanding Firewall Filters on OVSDB-Managed Interfaces.</a>]</li> <li>● MAC limiting on OVSDB managed interfaces. [See <a href="#">Features Supported on OVSDB-Managed Interfaces.</a>]</li> <li>● OVSDB commit failures, schema updates, and support with Contrail.</li> <li>● OVSDB software in Junos OS software package.</li> <li>● OVSDB support with VMware NSX for vSphere. See [<a href="#">Understanding the Junos OS Implementation of OVSDB and VXLAN in a VMware NSX for vSphere Environment.</a>]</li> <li>● Policers and storm control on OVSDB-managed interfaces. [See <a href="#">Understanding Firewall Filters on OVSDB-Managed Interfaces.</a>]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (continued)

Feature	Description
Routing Protocols	<ul style="list-style-type: none"> <li>● Bidirectional forwarding detection (BFD) support for BGP, IS-IS, and PIM. [See <a href="#">Example: Configuring BFD for BGP</a> and <a href="#">Example: Configuring BFD for IS-IS</a>.]</li> <li>● Static routing. [See <a href="#">Protocol-Independent Routing Properties User Guide</a>.]</li> <li>● Unified Forwarding Table (UFT). [See <a href="#">Understanding the Unified Forwarding Table</a>.]</li> <li>● IPv4 over GRE tunnels—encapsulation and de-encapsulation support.</li> <li>● IGMP version (v1/v2/v3), IGMP filter, IGMP snooping, proxy (relay), and querier. [See <a href="#">Understanding IGMP</a>, <a href="#">IGMP Snooping Overview</a>, and <a href="#">igmp-querier</a>.]</li> <li>● Remote support for LDP in IS-IS, static adjacency segment identifier for IS-IS, and alternate loop-free routes and topology-independent loop-free alternate for IS-IS. [See <a href="#">Understanding Remote LFA over LDP Tunnels in IS-IS Networks</a>.]</li> <li>● Multicast Listener Discovery version 1 and 2. [See <a href="#">Configuring MLD</a>.]</li> <li>● Multicast Source Discovery Protocol (MSDP) and multicast-only fast reroute (MoFRR). [See <a href="#">source (Protocols MSDP)</a>.]</li> <li>● IPv6 protocol independent multicast (PIM), PIM Static RP and PIM dense mode (PIM DM), PIM source-specific multicast (PIM SSM), and PIM sparse mode (PIM SM). [See <a href="#">PIM Overview</a>.]</li> <li>● Support for static multicast route leaking for VRF and virtual-router instances. [See <a href="#">Understanding Multicast Route Leaking for VRF and Virtual-Router Instances</a>.]</li> <li>● Virtual routing instances for multicast and unicast protocols. [See <a href="#">Configuring Virtual Router Routing Instances</a>.]</li> <li>● Remote LFA support for LDP tunnels in OSPF and alternate loop-free routes for OSPF and protocol independent multicast (PIM). [See <a href="#">Configuring Loop-Free Alternate Routes for OSPF</a>.]</li> </ul>
Spanning Tree Protocols	<ul style="list-style-type: none"> <li>● Support for IEEE 802.1s Multiple Spanning Tree Protocol (MSTP), IEEE 802.1w rapid spanning tree protocol (RSTP), IEEE 802.1D Spanning Tree Protocol (STP), and IEEE 802.1ak multiple VLAN Registration Protocol (MVRP). [See <a href="#">Spanning-Tree Protocols User Guide</a>.]</li> <li>● VSTP and RSTP and concurrent configuration. [See <a href="#">Configuring VSTP Protocol</a>.]</li> <li>● Bridge protocol data unit (BPDU) protection, loop protection, and root protection. [See <a href="#">BPDU Protection for Spanning-Tree Protocols</a>, <a href="#">Loop Protection for Spanning-Tree Protocols</a> and <a href="#">Understanding Root Protection for STP, RSTP, VSTP, and MSTP</a>.]</li> </ul>
System Logging	<ul style="list-style-type: none"> <li>● Support for forwarding structured system log messages to a remote system log server. [See <a href="#">Directing System Log Messages to a Remote Machine or the Other Routing Engine</a>.]</li> <li>● System logging (syslog) over IPv4 and IPv6.</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
System Management	<ul style="list-style-type: none"> <li>● Automatic software download, fast reboot, configuration and image rollback, commit process split into two steps, and rescue configuration. [See <a href="#">Software Installation and Upgrade Guide</a>.]</li> <li>● Support for Precision Time Protocol (PTP) transparent clock. [See <a href="#">Configuring Transparent Clock Mode for Precision Time Protocol</a>.]</li> <li>● Online insertion and removal (OIR). [See <a href="#">Removing an Expansion Module from a QFX5100 Device</a>.]</li> <li>● Device recovery mode introduced with upgraded FreeBSD. [See <a href="#">How to Recover Junos OS with Upgraded FreeBSD</a>.]</li> <li>● IPv4 support for Telnet. [See <a href="#">Configuring Telnet Service for Remote Access to a Switch</a>.]</li> <li>● Secure boot with system security enhancement: secure boot. [See <a href="#">Software Installation and Upgrade Guide</a>.]</li> <li>● Common BIOS support.</li> <li>● Licensing enhancements. [See <a href="#">Licenses for QFX Series</a>.]</li> <li>● Zero touch provisioning (ZTP). [See <a href="#">Understanding Zero Touch Provisioning</a>.]</li> </ul>
Time Management	<ul style="list-style-type: none"> <li>● Network Time Protocol (NTP). [See <a href="#">Understanding NTP Time Servers</a>.]</li> <li>● Enhancement to NTP authentication method. [See <a href="#">Configuring NTP Authentication Keys</a>.]</li> </ul>
VLANs	<ul style="list-style-type: none"> <li>● Configure tagged VLANs using the 802.1Q standard. [See <a href="#">Configuring Tagged VLANs</a>.]</li> <li>● Default VLAN and multiple VLAN range support, dual VLAN tag translation, routed VLAN interfaces, and jumbo frames.</li> <li>● Support for 4096 VLAN IDs. [See <a href="#">802.1Q VLAN IDs</a>.]</li> <li>● Support to exclude RVIs from state calculations. [See <a href="#">Excluding a Routed VLAN Interface from State Calculations</a>.]</li> <li>● Support for IRB interfaces on Q-in-Q VLANs. [See <a href="#">Configuring Q-in-Q Tunneling and VLAN Q-in-Q Tunneling and VLAN Translation</a>.]</li> <li>● Static MAC address assignment for physical interface.</li> <li>● Support for Private VLANs and Q-in-Q on the same interface. [See <a href="#">Understanding Private VLANs</a>.]</li> <li>● VLAN support for configuration and operational state models in Openconfig. [See <a href="#">OpenConfig Overview</a>.]</li> </ul>

Table 2: Features Supported by QFX5120-48T Switches (*continued*)

Feature	Description
---------	-------------

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

### ***Authentication, Authorization, and Accounting***

- **802.1X authentication on Layer 3 interfaces (QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, and QFX5220)**—Starting in Junos OS Release 20.2R1, 802.1X authentication is supported on Layer 3 interfaces. The 802.1X IEEE standard for port-based network access control authenticates users attached to a LAN port. It blocks all traffic to and from a supplicant (client) at the interface until the supplicant's credentials are presented and matched on the RADIUS authentication server.

[See [802.1X Authentication](#).]

### ***Class of Service***

- **CoS support in EVPN-VXLAN overlay networks (QFX10002, QFX10008, and QFX10016 switches)**—Starting with Junos OS Release 20.2R1, QFX10002, QFX10008, and QFX10016 switches support CoS in EVPN-VXLAN overlay networks, namely ingress and egress classification, scheduling, and rewrite rules based on IEEE 802.1p/DSCP code points.

[See [VXLAN Constraints on QFX Series and EX Series Switches](#).]

### ***EVPN***

- **EVPN-VXLAN multicast support (QFX10002-60C)**—Starting in Junos OS Release 20.2R1, the QFX10002-60C switch supports the following multicast features:
  - Internet Group Management Protocol version 2 (IGMPv2) and IGMP snooping [See [Overview of Multicast Forwarding with IGMP Snooping in an EVPN-VXLAN Environment](#).]
  - Selective multicast forwarding [See [Overview of Selective Multicast Forwarding](#).]
  - Assisted replication [See [Assisted Replication Multicast Optimization in EVPN Networks](#).]

With the support of these multicast features, the QFX10002-60C switch can now perform the following:

- Layer 2 intra-VLAN multicast forwarding
- Layer 3 inter-VLAN multicast routing with:
  - An IRB interface running Protocol Independent Multicast (PIM)
  - A PIM gateway connected through a Layer 2 multicast VLAN (MVLAN) or a Layer 3 interface



- An external multicast router

### *High Availability (HA) and Resiliency*

- **Support for failover configuration synchronization for the ephemeral database (EX Series, MX Series, MX Series Virtual Chassis, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.2R1, when you configure the **commit synchronize** statement at the **[edit system]** hierarchy level in the static configuration database of an MX Series Virtual Chassis or dual Routing Engine device, the backup Routing Engine will synchronize both the static and ephemeral configuration databases when it synchronizes its configuration with the master Routing Engine. This happens, for example, when a backup Routing Engine is newly inserted, comes back online, or changes roles. On a dual Routing Engine system, the backup Routing Engine synchronizes both configuration databases with the master Routing Engine. In an MX Series Virtual Chassis, the master Routing Engine on the protocol backup synchronizes both configuration databases with the master Routing Engine on the protocol master.

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

### *Interfaces and Chassis*

- **Support for 100-Gbps and 40-Gbps ports to operate at 10-Gbps or 1-Gbps speed (QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 20.2R1, you can use the Mellanox pluggable adapter (model number: MAM1Q00A-QSA) to convert quad-lane based ports to a single-lane based port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ cable connector. Use the QSA adapter to convert a 40GbE or a 100GbE port to a 10GbE or a 1GbE port. You can then plug-in an SFP+ transceiver or an SFP transceiver into the QSA adapter which is inserted into the QSFP+ or QSFP ports of the switch. You can use the commands **show chassis hardware** and **show chassis pic fpc-slot slot-number pic-slot slot-number** to view the optics inventory information for the QSFP ports.

With this adapter, the QSFP Ports on QFX10002, QFX10008, and QFX10016 switches support the following transceiver types— 100-Mbps, 1-Gbps, 10-Gbps SFP+: SR, LR, ER, ZR, CWDM, DAC and T-SFP+.

**NOTE:** For this adapter to work on the QSFP+ ports on the QFX10000-36Q line card in the QFX10008, you need to channelize the ports using the CLI command **set fpc fpc-slot pic pic-number port port-number port speed 10G**.

[See [show chassis hardware](#) and [show chassis pic](#).]

- **Support for multiple speeds and autonegotiation (QFX5120-48Y, QFX5110-48S, and QFX5100-48S with the JNP-SFPP-10GE-T transceiver)**—Starting in Junos OS Release 20.2R1, you can configure your switch to operate at multiple speeds when the JNP-SFPP-10GE-T transceiver is installed.

On the QFX5110-48S and QFX5100-48S switches, you can configure 100-Mbps, 1-Gbps, and 10-Gbps speeds on the mge-0/0/z port by using the **set interfaces mge-0/0/z speed (100m|1g|10g)** command.

The switch ports operate at the configured speed and they can also switch to a supported lower speed (automatically) with the same transceiver installed, based on peer capability.

The QFX5120 operates at only two speeds—10 Gbps and 1 Gbps—when this transceiver is installed. By default, the switch comes up with 10-Gbps speed. To operate at 1-Gbps speed, use the **set chassis fpc 0 pic 0 port *port-number* speed 1G** command. Due to hardware limitations, you can configure the *port-number* value only in multiples of four, starting from port 0. You must also configure sets of four consecutive ports (for example, 0-3, 4-7, and so on) to operate at the common speed. After setting 1-Gbps speed, to revert to 10-Gbps speed, simply delete the **1G** speed configuration.

**NOTE:** Only QFX5110-48S and QFX5100-48S switches support the multi-rate Gigabit Ethernet (mge) interface.

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

### *Juniper Extension Toolkit (JET)*

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the **set system scripts language python3** command.

[See [language \(Scripts\)](#), [Develop Off-Device JET Applications](#), and [Develop On-Device JET Applications](#).]

### *Junos Telemetry Interface*

- **Network instance (policy) statistics and OpenConfig configuration enhancements on JTI (ACX1100, ACX2100, ACX5448, ACX6360, EX4300, MX240, MX480, MX960, MX10003, PTX10008, PTX10016, QFX5110, and QFX10002)**—Junos OS Release 20.2R1 provides enhancements to support the OpenConfig data models **openconfig-local-routing.yang** and **openconfig-network-instance.yang**.

[See [Mapping OpenConfig Routing Policy Commands to Junos Configuration](#) and [Mapping OpenConfig Network Instance Commands to Junos Operation](#).]

- **ON-CHANGE BGP peer information statistics support for JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 provides BGP peer sensor support using Junos telemetry interface (JTI) and remote procedure call (gRPC) services or gRPC Network Management Interface (gNMI) services. ON\_CHANGE statistics are sent to an outside collector.

The following resource paths are supported:

- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/active** (ON\_CHANGE)
- **/network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes** (ON\_CHANGE)

- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/received (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/sent (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/afi-safis/afi-safi/state/prefixes/rejected (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/admin-state (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/established-transitions (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/last-established (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/received/notification (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/messages/received/update (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/notification (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/messages/sent/update (stream)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/session-state (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/neighbors/neighbor/state/supported-capabilities (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/transport/state/local-address (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-address (ON\_CHANGE)
- /network-instances/network-instance/protocols/protocol/bgp/transport/state/remote-port (ON\_CHANGE)

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

- **EVPN statistics export using JTI (QFX5100, QFX5110, QFX5120, QFX5200, QFX10002-60C, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and using remote procedure call (gRPC) services to export EVPN statistics from devices to an outside collector.

Use the following sensors to export EVPN statistics:

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

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

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

- **CPU statistics support on JTI (MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, QFX5100, and QFX5200)**—Junos OS Release 20.2R1 supports streaming various CPU statistics and process parameters using remote procedure call (gRPC) or gRPC Network Management Interface (gNMI) services and Junos telemetry interface (JTI). You can stream CPU usage per process (statistics are similar to output from the `show system process detail` operational mode command), as well as CPU usage per Routing Engine core.

This feature supports the private data model `openconfig-procmon.yang`.

To stream statistics to an outside collector, include the following resource paths in a gRPC or gNMI subscription:

- Individual process level information (resource path `/system/processes/process`)
- Individual Routing Engine core information (resource path `/components/component/cpu/`)

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

- **Packet Forwarding Engine sensor support with INITIAL\_SYNC on JTI (MX960, MX2008, MX2010, MX2020, PTX1000, PTX5000, PTX10000 line of routers, QFX5100, and QFX5200)**—Starting in Junos OS Release 20.2R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export Packet Forwarding Engine statistics from devices to an outside collector using gNMI submode `INITIAL_SYNC`. When an external collector sends a subscription request for a sensor with `INITIAL_SYNC` (`gnmi-submode 2`), the host sends all supported target leaves (fields) under that resource path at least once to the collector with the current value. This is valuable because:

- The collector has a complete view of the current state of every field on the device for that sensor path.
- Event-driven data (`ON_CHANGE`) is received by the collector at least once before the next event is seen. In this way, the collector is aware of the data state before the next event happens.
- Packet Forwarding Engine sensors that contain zero counter values (zero-suppressed) that normally do not show up in streamed data are sent, ensuring that all fields from each line card (also referred to as source) are known to the collector.

**NOTE:** `ON_CHANGE` data is not available for native (UDP) Packet Forwarding Engine Sensors.

INITIAL\_SYNC submode requires that at least one copy to be sent to the collector; however, sending more than one is acceptable.

INITIAL\_SYNC submode is supported for the following sensors:

- Sensor for CPU (ukernel) memory (resource path `/junos/system/linecard/cpu/memory/`)
- Sensor for firewall filter statistics (resource path `/junos/system/linecard/firewall/`)
- Sensor for physical interface traffic (resource path `/junos/system/linecard/interface/`)
- Sensor for logical interface traffic (resource path `/junos/system/linecard/interface/logical/usage/`)
- Sensor for physical interface queue traffic (resource path `/junos/system/linecard/interface/queue/`)
- Sensor for physical interface traffic except queue statistics (resource path `/junos/system/linecard/interface/traffic/`)
- Sensor for NPU memory (resource path `/junos/system/linecard/npu/memory/`)
- Sensor for NPU utilization (resource path `/junos/system/linecard/npu/utilization/`)
- Sensor for packet statistics (resource path `/junos/system/linecard/packet/usage/`)
- Sensor for software-pollled queue-monitoring statistics (resource path `/junos/system/linecard/qmon-sw/`)

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

### Layer 2 Features

- **L2PT support (EX4650 and QFX5120-48Y switches, and QFX5100 and QFX5110 switches and Virtual Chassis)**—Starting in Junos OS Release 20.2R1, you can configure Layer 2 protocol tunneling (L2PT) to tunnel any of the following Layer 2 protocols: CDP, E-LMI, GVRP, IEEE 802.1X, IEEE 802.3AH, LACP, LLDP, MMRP, MVRP, STP (including RSTP and MSTP), UDLD, VSTP, and VTP.

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

### Multicast

- **Static multicast route leaking for VRF and virtual router instances (EX4650 and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can configure the switch to statically share (leak) IPv4 multicast routes for IGMPv3 (S,G) traffic among different virtual router or virtual routing and forwarding (VRF) instances. You can only leak static multicast routes per group, not per source and group. The destination prefix length must be 32.

To configure multicast route leaking to the VRF or virtual router instance *routing-instance-name*, configure the **next-table *routing-instance-name.inet.0*** statement at the **[edit routing-instances *routing-instance-name* routing-options static route destination-prefix/32]** hierarchy level.

[See [Understanding Multicast Route Leaking for VRF and Virtual Router Instances](#).]

- **Multicast-only fast reroute (MoFRR) (EX4650 and QFX5120-48Y)**—Starting in Junos OS Release 20.2R1, you can configure MoFRR to minimize multicast packet loss in PIM domains when link failures occur. With MoFRR enabled, the switch maintains primary and backup traffic paths, forwarding traffic from the primary path and dropping traffic from the backup path. If the primary path fails, the switch can quickly start forwarding the backup path stream (which becomes the primary path). The switch creates a new backup path if it detects available alternative paths. MoFRR applies to all multicast (S,G) streams by default, or you can configure a policy for the (S,G) entries where you want MoFRR to apply.

[See [Understanding Multicast-Only Fast Reroute.](#)]

### **Network Management and Monitoring**

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

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

### **Routing Policy and Firewall Filters**

- **Support for MPLS firewall filter on loopback interface (EX4650, QFX5120-32C, and QFX5120-48Y)**—Starting with Junos OS Release 20.2R1, you can apply an MPLS firewall filter to a loopback interface on a label-switching router (LSR). For example, you can configure an MPLS packet with `ttl=1` along with MPLS qualifiers such as `label`, `exp`, and Layer 4 `tcp/udp` port numbers. Supported actions include `accept`, `discard`, and `count`.

You configure this feature at the `[edit firewall family mpls]` hierarchy level. You can only apply a loopback filters on `family mpls` in the ingress direction.

[See [Overview of MPLS Firewall Filters on Loopback Interface.](#)]

### **Virtual Chassis**

- **Virtual Chassis with NSSU support (QFX5120-48T)**—Starting in Junos OS Release 20.2R1, you can interconnect two QFX5120-48T switches into a Virtual Chassis that operates as one logical device managed as a single chassis. The Virtual Chassis:

- Has both switches in Routing Engine role (one master and one backup)
- Supports 100GbE QSFP28 or 40GbE QSFP+ ports (48 through 53) as Virtual Chassis ports (VCPs)
- Supports NSSU

A QFX5120-48T Virtual Chassis supports the same protocols and features as a standalone switch in Junos OS Release 20.2R1 except for the following:

- EVPN-VXLAN
- Junos telemetry interface (JTI)
- Multichassis link aggregation (MC-LAG)
- Priority-based flow control (PFC)

Configuration parameters and operation are the same as for other non-mixed QFX Series Virtual Chassis.

[See [Virtual Chassis Overview for Switches.](#)]

- **802.1X authentication, Layer 2 port security, and MPLS support in a Virtual Chassis (QFX5120-48Y Virtual Chassis)**—Starting in Junos OS Release 20.2R1, the following protocol features are supported on a QFX5120-48Y Virtual Chassis:
  - IEEE 802.1X authentication
  - Layer 2 port security features, including IP source guard, IPv6 router advertisement (RA) guard, DHCP, and DHCP snooping
  - MPLS

Configuration and operation are the same on the Virtual Chassis as on the standalone switch.

[See [802.1X Authentication](#), [MPLS Overview](#), [DHCP Snooping](#), [Understanding DHCP Snooping \(ELS\)](#), [Understanding IP Source Guard for Port Security on Switches](#), and [Understanding IPv6 Router Advertisement Guard.](#)]

## SEE ALSO

[What's Changed | 256](#)

[Known Limitations | 260](#)

[Open Issues | 262](#)

[Resolved Issues | 267](#)

[Documentation Updates | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)



## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 256](#)
- [What's Changed in Release 20.2R2 | 257](#)
- [What's Changed in Release 20.2R1 | 258](#)

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

### What's Changed in Release 20.2R3

#### *General Routing*

- **Support only for manual channelization on QSFP-100G-SR4-T2 optics (QFX5120-48T and QFX5120-32C)**— We recommend that you use the active optical cable (AOC) for auto-channelization. The QSFP-100G-SR4-T2 cables do not support auto-channelization. To use the QSFP-100G-SR4-T2 optics with an external breakout cable, you must configure the channelization manually by running the **channel-speed** statement at the **edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)** hierarchy level.

[See [channel-speed](#).]

#### *Junos XML API and Scripting*

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

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

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

releases in which the root UI\_LOGIN\_EVENT and UI\_LOGOUT\_EVENT messages are included in system log files.

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

### *Network Management and Monitoring*

- **Changes to <commit> RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:
  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.
  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and emits only an `<ok>` or `<rpc-error>` element.

[See [Configuring RFC-Compliant NETCONF Sessions](#).]

### *User Interface and Configuration*

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

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

## **What's Changed in Release 20.2R2**

### *Platform and Infrastructure*

- **Priority-based flow control (PFC) support (QFX5120-32C)**—Starting in Junos OS 20.2R2, we provide support for priority-based flow control (PFC) using Differentiated Services code points (DSCPs) at Layer 3 for untagged traffic.
- **IPv6 address in the prefix TIEs displayed correctly**—The IPv6 address in the prefix TIEs are displayed correctly in the `show rift tie` output.

### **Routing Protocols**

- Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.
- 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.

### **What's Changed in Release 20.2R1**

#### **General Routing**

- **Support for full inheritance paths of configuration groups to be built into the database by default (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting with Junos OS Release 20.2R1, the `persist-groups-inheritance` option at the `[edit system commit]` hierarchy level is enabled by default. To disable this option, use `no-persist-groups-inheritance`.

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

- **Priority-based flow control (PFC) support (QFX5120-32C)**—We provide support for priority-based flow control (PFC) using Differentiated Services code points (DSCPs) at Layer 3 for untagged traffic.

#### **Interfaces and Chassis**

- **Autonegotiation status displayed correctly (QFX5120-48Y)**—In Junos OS Release 20.2R1, 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 Extension Toolkit**

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the `PASS` keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

[See [Develop Off-Device JET Applications](#) and [Develop On-Device JET Applications](#).]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document

additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH\_USAGE\_INVALID**.

[See [Juniper EngNet](#).]

### **Network Management and Monitoring**

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

### **Routing Protocol**

- **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 | 231](#)

---

[Known Limitations | 260](#)

---

[Open Issues | 262](#)

---

[Resolved Issues | 267](#)

---

[Documentation Updates | 280](#)

---

[Migration, Upgrade, and Downgrade Instructions | 281](#)

## Known Limitations

### IN THIS SECTION

- [Class of Service \(CoS\) | 260](#)
- [Layer 2 Features | 260](#)
- [Layer 2 Ethernet Services | 260](#)
- [Platform and Infrastructure | 260](#)
- [Routing Protocols | 261](#)

Learn about known limitations in Junos OS Release 20.2R3 for QFX Series Switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

### Class of Service (CoS)

- On the QFX5100 devices, ISSU does not support Junos OS Release 20.1 and later. [PR1479439](#)

### Layer 2 Features

- On the QFX5000 devices with storm control, significant difference between the configured rate and actual rate is observed. [PR1526906](#)

### Layer 2 Ethernet Services

- If the configuration or image file name has special characters such as #, %, or @, ZTP over HTTP or HTTPS does not work. [PR1503588](#)

### Platform and Infrastructure

- After configuring and deleting the Ethernet loopback configuration, the interface goes down and does not come up. [PR1353734](#)
- The QFX5000 device gets stuck in the database prompt state after rebooting. [PR1411826](#)
- On the QFX10000 line of switches, the analyzer does not mirror after adding the child member to an aggregated Ethernet interface. [PR1417694](#)

- On the QFX5120 line of switches, one of the VCP ports of the throughput test result for most of the frame sizes is not close to 100 percent. [PR1453709](#)
- After changing the VLAN name on the trunk interface, the local host MAC learning does not hold for more than 30 seconds. [PR1454274](#)
- On the QFX5120-48T device, convergence delay for the link-protected MPLS LSP is more than 50 minutes. [PR1478584](#)
- On the QFX5120 device, the following error message is observed while performing NSSU: **syntax error: request-package-validate message**. [PR1479753](#)
- There is no option to upgrade firmware for the backup Routing Engine. [PR1479925](#)
- The output of the **show snmp mib walk jnxFruName** command has an extra entry for the Routing Engine. [PR1483384](#)
- On the QFX5120 Virtual Chassis, the output of the **show chassis alarm** command displays incorrect PEM status after multiple GRES events. [PR1486736](#)
- On the QFX10000 devices, traffic drop for more than 50 minutes is observed on bringing down the aggregated Ethernet interface. [PR1486853](#)
- A 100 percent Layer 2 MAC scaling traffic loss is observed in the QFX10002-60C switch after loading the EVPN-VXLAN collapsed profile configurations. [PR1489753](#)
- Data corruption might occur while abrupt power cycles are performed. [PR1507750](#)
- Changing the scaled firewall profiles on the fly does not release the TCAM resources as expected. [PR1512242](#)
- On the QFX10000 device, the interface encapsulation ethernet-bridge for EVPN is not supported. [PR1538852](#)
- On the QFX5000 device, microburst absorption is limited. [PR1545046](#)

## Routing Protocols

- The multicast route and pim (s,g) are incorrectly populated. [PR1483732](#)
- On QFX5100 devices not running the QFX-5E codes (non TVP architecture), when image with Broadcom SDK upgrade (6.5.X) is installed, the CPU utilization might go up by around 5 percent. [PR1534234](#)
- On the QFX10002 device, the S,G convergence on the remote PE devices are very slow, taking around 30 minutes to converge completely. [PR1542675](#)

SEE ALSO

| [What's New | 231](#)

---

[What's Changed | 256](#)

---

[Open Issues | 262](#)

---

[Resolved Issues | 267](#)

---

[Documentation Updates | 280](#)

---

[Migration, Upgrade, and Downgrade Instructions | 281](#)

## Open Issues

### IN THIS SECTION

- [EVPN | 262](#)
- [High Availability \(HA\) and Resiliency | 263](#)
- [Infrastructure | 263](#)
- [Interfaces and Chassis | 263](#)
- [Layer 2 Features | 263](#)
- [Layer 2 Ethernet Services | 263](#)
- [Platform and Infrastructure | 264](#)
- [Routing Protocols | 266](#)
- [Virtual Chassis | 267](#)

Learn about open issues in Junos OS Release 20.2R3 for QFX Series Switches. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

### EVPN

- In the ERB scale setup powering up, a leaf might cause ingress traffic loss upto 250 seconds. [PR1544204](#)
- After changing VNID, it takes about 7 minutes for the control plane to populate remote VTEPs in the VLAN. [PR1550163](#)

## High Availability (HA) and Resiliency

- On the QFX5200-32C devices, the reboot time is degraded from 205 seconds in Junos OS Release 20.2R1 to 260 seconds in Junos OS Release 20.3R1. [PR1511607](#)

## Infrastructure

- The following error message is seen during FTP: **ftpd[14105]: bl\_init: connect failed for /var/run/blacklistd.sock(No such file or directory)**. [PR1315605](#)
- Device goes to database prompt with **panic: ffs\_valloc: dup alloc** during powering on of the device. [PR1480185](#)

## Interfaces and Chassis

- On the QFX5110 MC-LAG, flooding of the multicast packets for around 16 to 20 seconds is observed after disabling and enabling a member link of ICL after reboot. [PR1422473](#)

## Layer 2 Features

- On the QFX5000 Virtual Chassis, multicast traffic gets flooded even when the IGMP report times out. [PR1431893](#)
- New tenant addition and deletion leads to intra-VNI traffic drop for a few milliseconds. [PR1455654](#)
- On QFX5110 and QFX5120 platforms, changing lo0 IP address might sometimes either result in stale entry of IP in mpls\_entry table or missing IP entry, which results in traffic drop for VXLAN traffic. [PR1472333](#)
- Traffic does not get load balanced by QFX5000 platforms over ESI links with EVPN\_VXLAN configured. [PR1551543](#)

## Layer 2 Ethernet Services

- The DHCP decline packets are not forwarded to the DHCP server when forward-only is set within dhcp-reply. [PR1429456](#)
- ZTP not getting activated after returning the device to zero was observed once or twice. [PR1529246](#)



## Platform and Infrastructure

- When an RSVP path is deleted (because of LSP deletion or switchover to a new path), the reservation state block (RSB) data structure has to be deleted to free up memory. When RSB deletion is performed, LSP attribute object in RSB is not deleted by rpd. This causes build-up of rpd memory usage over a period of time (memory leak). Build-up of rpd memory is proportional to the frequency of RSB deletes. [PR1115686](#)
- On the QFX5100-48T-6Q devices, port LEDs might not work. [PR1317750](#)
- On the QFX10000 devices, source MAC and TTL values are not updated for routed multicast packets in EVPN-VXLAN. [PR1346894](#)
- The backup Routing Engine might crash after GRES occurs continuously for more than 10 times. [PR1348806](#)
- On the QFX10000 line of switches, the Aruba wireless access point (AP) heartbeat packets get dropped. As a result, the Aruba wireless AP cannot work. [PR1352805](#)
- USB upgrade of network operating system image is not supported. [PR1373900](#)
- Due to the transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Those reported events had been disabled to prevent alarms and possibly unnecessary hardware replacements. [PR1384435](#)
- The DRAM and buffer utilization fields are not correct. [PR1394978](#)
- CPU performance might become slow. [PR1399369](#)
- uRPF in the **Strict** mode does not work. [PR1417546](#)
- The IPv6 communication issue might be observed after passing through the QFX10002-60C devices. [PR1424244](#)
- When spine underlay is tagged and untagged, the inner packet comes over the TYPE-2 tunnel and goes over the TYPE-2 tunnel resulting in IPv4 to silently discard traffic on PECHIP. [PR1435864](#)
- On the QFX5200 line of switches, the ISSU might fail. [PR1438690](#)
- On the QFX5000 devices, the port qualifier is not supported. [PR1440980](#)
- On the QFX10000 line of switches, removal of the EVPN-VXLAN Layer 3 gateway on the IRB interface from the spine switches might cause traffic to be silently discarded. [PR1446291](#)
- The vhostd application fails to generate a minor alarm. [PR1448413](#)
- On the QFX5000 line of switches, misleading ISSU logs are printed during the NSSU process even when the box does not perform ISSU. [PR1451375](#)
- Interface sends mirrored traffic out even after it is removed from the output VLAN. [PR1452459](#)
- 9.51 percent of degradation with commit time and 12 percent of degradation with VLAN commit convergence are observed while comparing 19.4DCB with 19.3DCB. [PR1457939](#)

- **storm-control** does not rate-limit ARP packets. [PR1461958](#)
- On the QFX5110 line of switches, the VXLAN VNI (mcast) scaling causes traffic issue. [PR1462548](#)
- On the QFX10002-60C line of switches, the Packet Forwarding Engine installation or deletion, and link flap convergence time are reduced in Junos OS Release 19.4 compared to Junos OS Releases 19.3R1 and 19.2R1. [PR1464572](#)
- On the QFX5120-48T devices, finding discrepancy in the output of the **show chassis environment pem** command can be seen in the backup member as well. [PR1474520](#)
- On the QFX5220 devices, the lo0 firewall filter might affect the Layer 3 forwarding traffic. [PR1475620](#)
- On the QFX10000 devices, the loopback-based filter with decap GRE does not work as expected. [PR1479613](#)
- The output of the **app-engin** command displays a command that does not display information about the backup member. [PR1479900](#)
- On the QFX5120-48T devices, the JTI exports in the fan state as **Online** for a failed fan module. [PR1480259](#)
- On the QFX5110 and QFX5120 devices, the ICMP redirect messages are not generated. [PR1481020](#)
- On the QFX5000 device, dcpfe does not come up in an abrupt power-off or power-on situation. [PR1481176](#)
- Disabled interfaces might still transmit power after the device reboots. [PR1487554](#)
- On the QFX5120-48T devices, commit fails on the backup device of the Virtual Chassis while removing storm control with HA configured. Warning messages are also observed as patch removes the statement that is not empty. [PR1488847](#)
- Interface on platforms using Broadcom chipset might have an abnormal status. [PR1495564](#)
- The interfaces on the EX4600-EM-8F device expansion module do not come up on the QFX5100-24Q device with the QFX5E image. [PR1502237](#)
- On the QFX5100 devices, degradation is observed during the system reboot time and FPC online time. [PR1513540](#)
- On the QFX10002-60C devices, degradation during system reboot time is observed. [PR1516086](#)
- The dcpfe process generates the core file after adding IRB in the same routing instance as that of the underlay VTEP interface. [PR1519651](#)
- SNMP trap of power failure might not be sent out. [PR1520144](#)
- Higher token allocation with the **arp-enhanced-scale** command due to kernel global token leakage is observed. [PR1530947](#)
- On QFX5100 which is working on 5e image, LED is not working well on 40G port and channelized port. [PR1536395](#)

- The BFD neighborship fails with the EVPN\_VXLAN configuration after the Layer 2 learning restarts. [PR1538600](#)
- On the QFX5000 devices, route leaking does not work for the IPv4 routes if mask is less than 16 and for the IPV6 routes if mask is less than 64. [PR1538853](#)
- On the QFX10002-60C devices, ARP or token scale is lower than the QFX10002 and QFX10008 devices that causes the dcpfe process to generate the core file at a high scale. [PR1541686](#)
- On the QFX5000 Virtual Chassis fan, traffic loss might be seen after swapping the primary and backup Routing Engines. [PR1544353](#)
- BD creation fails for few VLANs while switching from the script configuration to profile configuration. [PR1545517](#)
- Need to move WRL7 to RCPL31 for the QFX-10-M and QFX-10-F devices. [PR1547565](#)
- After 12 hours of longevity with events, the Layer 3 traffic with destination to local host is dropped. [PR1548740](#)
- Traffic does not get load-balanced by the QFX10002 device over ESI links with EVPN-VXLAN configured. [PR1550305](#)
- PRBS (psuedorandom binary sequence) test on the QFX5200 device fails for 100GbE interfaces with the default settings. [PR1560086](#)
- On the QFX5100 Virtual Chassis, the following continuous message is observed: **agentd-pfe-proxy\_telemetry\_publisher**. [PR1566528](#)
- On the QFX5100 device, the following internal comment is displayed: **Placeholder for QFX platform configuration**. [PR1567037](#)
- The Packet Forwarding Engine might produce error messages while deleting an interface in configurations with IRB interfaces. [PR1054798](#)
- 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](#)
- Upgrading satellite devices may lead to some SDs in SyncWait state. Cascade port flap not causing the issue. [PR1556850](#)

## Routing Protocols

- On the QFX5100 Virtual Chassis, instability issues due to disabling DDoS protection is observed. [PR1238875](#)
- On the QFX5100 Virtual Chassis or Virtual Chassis fan, the following error is observed in the hardware with the mini-PDT base configurations: **BRCM\_NH-,brcm\_nh\_bdvlan\_ucast\_uninstall(), 128:I3 nh 6594 unintsall failed**. [PR1407175](#)

- The remaining BFD sessions of the aggregated Ethernet interface flap continuously if one of the BFD sessions is deleted. [PR1516556](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- Sometimes when we perform **deactivate protocols bgp** on the QFX5000 RIOT devices, we may see **BRCM-VIRTUAL,brcm\_vxlan\_riot\_destroy\_nh(),1494:Failed to delete egr\_if(400138) err-Operation still running** error messages during arp\_ndp clean up stage and these are harmless. [PR1529240](#)
- BFD for BGP protocol flaps with sub-second timers with certain events performed in the fabric. [PR1539085](#)

## Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. [PR1490552](#)

## SEE ALSO

[What's New | 231](#)

[What's Changed | 256](#)

[Known Limitations | 260](#)

[Resolved Issues | 267](#)

[Documentation Updates | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)

## Resolved Issues

### IN THIS SECTION

- [Resolved Issues: 20.2R3 | 268](#)
- [Resolved Issues: 20.2R2-S2 | 272](#)
- [Resolved Issues: 20.2R2 | 272](#)
- [Resolved Issues: 20.2R1 | 275](#)

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

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

## Resolved Issues: 20.2R3

### *EVPN*

- On the QFX5000 device used on EVPN-VXLAN scenarios, load-balancing traffic (inter VLAN) might not work for multiple ESI-VTEP pairs with the underlay aggregated Ethernet interface between leaf and spines. [PR1512253](#)
- All the ARP reply packets toward some address are flooded across the entire fabric. [PR1535515](#)
- EVPN-VXLAN registers MAC-move counters under **system statistics bridge** even though there is no actual MAC-move for the multihome clients. [PR1538117](#)
- The l2ald process might generate core file if the EVPN-VXLAN configuration is changed. [PR1541904](#)
- The l2ald daemon might crash when **forwarding-options evpn-vxlan shared-tunnels** is configured. [PR1548502](#)
- The EVPN-VXLAN MAC-IP aging test fails. [PR1562925](#)

### *Forwarding and Sampling*

- The l2ald process might crash due to next-hop issue in the EVPN-MPLS. [PR1548124](#)

### *Infrastructure*

- The output of the **show interfaces extensive** command might display 0 temporarily during a race condition when SNMP query is issued. [PR1533314](#)

### *Interfaces and Chassis*

- The logical interface might flap after the addition or deletion of the native VLAN configuration. [PR1539991](#)
- MAC entry remains as DR after MC-LAG failover. [PR1562535](#)

### *Layer 2 Features*

- Traffic might be forwarded incorrectly on an interface with VXLAN enabled and the **hold-time up xxx** command statement configured. [PR1550918](#)
- On the QFX5120 devices, packets with VLAN ID 0 are dropped. [PR1566850](#)

### Layer 2 Ethernet Services

- DHCP packet drop may be seen when DHCP relay is configured on leaf device. [PR1554992](#)

### Platform and Infrastructure

- On the QFX5000 line of switches, the number of egress ACL filter entries is only 512 in Junos OS Release 19.4R1. [PR1472206](#)
- On the QFX10000 devices, the chassisd process might generate core files on the backup Routing Engine after committing due to **CHASSISD\_MAIN\_THREAD\_STALLED** for 200 seconds. [PR1481143](#)
- SNMP index in the Packet Forwarding Engine reports as 0, causing sFlow to report either IIF or OIF (not both) as 0 in the sFlow record data at the collector. [PR1484322](#)
- IRB MAC is not be programmed in hardware when the MAC persistence timer expires. [PR1484440](#)
- Slow response might be observed if the **show | compare** or **commit check** action in a large-scale configuration environment is committed. [PR1500988](#)
- The output VLAN push might not work. [PR1510629](#)
- On the QFX5000 line of switches, multicast traffic loss is observed due to few multicast routes missing in the spine node. [PR1510794](#)
- The DHCP traffic might not be forwarded correctly while sending the DHCP unicast packets. [PR1512175](#)
- Channelized interfaces might fail to come up. [PR1512203](#)
- In a Virtual Chassis environment, the output of the **show chassis forwarding-options** command displays incorrect value when num-65-127-prefix value is configured for the FPC that is not local (backup and line card members of the Virtual Chassis). [PR1512712](#)
- On the QFX5100 devices, cprod timeout triggers high CPU utilization. [PR1520956](#)
- The output interface index in the sFLOW packet is zero when the transit traffic is observed on the IRB interface with VRRP enabled. [PR1521732](#)
- On the QFX10000 devices, channelizing the 40GbE port to 10GbE port might bring down another interface. [PR1527814](#)
- Packet loss is observed while validating the policer after restarting the chassis control. [PR1531095](#)
- QFX10k2 / Firewall log incorrectly populating from Packet Forwarding Engine. [PR1533814](#)
- High rate of ARP or NS packets might be observed between a device that runs Junos OS and host when the device that runs Junos OS receives an ARP or NS packet on an interface in transition. [PR1534796](#)
- The following Packet Forwarding Engine error message is seen:  
**BRCM-VIRTUAL,brcm\_virtual\_tunnel\_port\_create(),489:Failed NW vxlan port token(45) hw-id(7026) status(Entry not found).** [PR1535555](#)
- Software recovery or installation using the **Bootable USB Flash Drive** option might fail. [PR1536799](#)
- The interfaces on QFX5100-48T switch might stay up when the peer device is rebooting. [PR1538071](#)

- On the QFX5100-48T devices, interfaces are not created after channel-speed 10Gbps is applied across ports 48 to 53. [PR1538340](#)
- The **Management Ethernet link down** alarm is seen while verifying the system alarms in a Virtual Chassis setup. [PR1538674](#)
- ARP request might be dropped in the leaf in the EVPN-VXLAN scenario. [PR1539278](#)
- The rpd memory leak might be observed on the backup Routing Engine due to link flaps. [PR1539601](#)
- Not able to take RSI properly due to the authentication error. [PR1539654](#)
- FPC might not be recognized after power cycle (hard reboot). [PR1540107](#)
- On the QFX5100 Virtual Chassis, the **End segment Not Present** message is not reported for the ping overlay function with the local host MAC. [PR1542226](#)
- On the QFX5000 devices running EVPN-VXLAN, the Packet Forwarding Engine related error message might be observed: **bd\_platform\_irk\_ifl\_attach\_detach: platform specific irb ifl attach/detach failed (-1)**. [PR1543812](#)
- The Broadcom chip FPC might crash during system bootup. [PR1545455](#)
- OSPFv3 session may keep flapping and OSPFv3 hellos might be dropped in the host-path. [PR1547032](#)
- On the QFX10000 devices, traffic might get dropped while changing the configuration to **set routing-options forwarding-table no-ecmp-fast-reroute** with 128 ECMP entries. [PR1547457](#)
- On the QFX5100 Virtual Chassis, the backup Routing Engine clears the reporting alarm for a PEM failure intermittently for a missing power source. [PR1548079](#)
- The 40GbE interface might be channelized after the Virtual Chassis member restarts. [PR1548267](#)
- Neighbor Solicitation might be dropped from the peer device. [PR1550632](#)
- Interface filter with source-port 0 matches everything instead of just port 0. [PR1551305](#)
- On the QFX5110 and QFX5120 devices, the DHCPv6 traffic received over VTEP might not be forwarded. [PR1551710](#)
- The **action-shutdown** command of storm control does not work for the ARP broadcast packets. [PR1552815](#)
- The traffic might not be passed because VLAN tag 2 is added while passing through the Virtual Chassis port. [PR1555835](#)
- Traffic might be dropped when a firewall filter rule uses the **then vlan** action. [PR1556198](#)
- Analyzer might cause traffic storm due to the flapping of the link. [PR1557274](#)
- Licenses for the VRRP, CFM, QINQ, VXLAN, MCLAG, ESI-LAG, LFM/Ethernet-OAM features might incorrectly show as invalid licenses. [PR1558017](#)
- On the QFX5000 devices, the firewall filter might fail to work. [PR1558320](#)

- Amber LEDs are observed for fan modules in the QFX5120 devices after upgrading to Junos OS Release 20.2R1. [PR1558407](#)
- Few IPv6 ARP resolutions might fail after loading the base configurations. [PR1560161](#)
- When configuring the static MAC and static ARP on the EVPN core aggregate interface the underlay next-hop programming might not be updated in the Packet Forwarding Engine. [PR1561084](#)
- On the QFX5110-48S-4C devices, the PTP lock status gets stuck at the **Acquiring** state instead at the **Phase aligned** state. [PR1561372](#)
- On the QFX5000 devices, port mirroring might not work as expected. [PR1562607](#)
- On the QFX5120 devices, **storm control** with IRB interface might not work correctly. [PR1564020](#)
- QFX10K: Firewall log incorrectly populating from PFE for IPv6 traffic. [PR1569120](#)

### ***Routing Policy and Firewall Filters***

- The policy configuration might be mismatched between the rpd and mgd processes when deactivate policy-options prefix-list is involved in the configuration sequence. [PR1523891](#)

### ***Routing Protocols***

- On the QFX 5100-48T-6Q Virtual Chassis or Virtual Chassis fan, the following error message is observed while copying the image to the Virtual Chassis fan member and trying to downgrade the image: **rpf for member 14, failed**. [PR1486632](#)
- Traffic might be silently discarded when the clear bgp neighbor all command is executed on a router and also on the corresponding Route reflector in succession. [PR1514966](#)
- The dcpfe process might crash while updating VRF instances for multicast routes during IRB uninit. [PR1546745](#)
- BGP LU session might flap when the Accumulated Interior Gateway protocol is used. [PR1558102](#)
- On the QFX5110-32Q device, the following syslog error message is observed after loading the NC T5 EVPN-VXLAN configuration: **LBCM-L2,pfe\_bcm\_l2\_sp\_bridge\_port\_tpid\_set() Config TPID New/Old (8100:8100) Other-Tpid's ba49, 4aa0, 80f**. [PR1558189](#)
- The dcpfe process might crash when the size of the Local Bias Filter Bitmap string exceeds 256 characters. [PR1568159](#)
- On the QFX5210-64C device, ping does not work while verifying the native VLAN behavior on the Q-n-Q interface. [PR1568533](#)



### ***User Interface and Configuration***

- The config under groups stanza is not inherited properly. [PR1529989](#)

### **Resolved Issues: 20.2R2-S2**

- On the QFX5120-48Y line of switches, amber LED lights are on continuously displayed on the fan modules even though there are no faults in the fan after upgrading to Junos OS Release 20.2R1 and later. [PR1558407](#)

### **Resolved Issues: 20.2R2**

#### ***Class of Service (CoS)***

- The PFC feature is not supported with the QFX5120 Virtual Chassis due to chip limitation. [PR1431895](#)
- Traffic might be forwarded to the incorrect queue when a fixed classifier is used. [PR1510365](#)

#### ***EVPN***

- EVPN-VXLAN core isolation is not working when the system is rebooted or the routing is restarted. [PR1461795](#)
- Unable to create a new VTEP interface. [PR1520078](#)
- ARP table might not be updated after performing VMotion or a network loop. [PR1521526](#)
- All the ARP reply packets towards some address are flooded across the entire fabric. [PR1535515](#)

#### ***Infrastructure***

- OID ifOutDiscards reports zero and sometimes shows valid value. [PR1522561](#)

#### ***Interfaces and Chassis***

- The dcpfe might crash when the ICL is disabled and then enabled. [PR1525234](#)

#### ***Layer 2 Ethernet Services***

- EX/QFX device sometimes doesn't obtain default-route or route listing gets delayed. [PR1504931](#)
- The aggregated Ethernet interface sometimes might not come up after switch is rebooted. [PR1505523](#)

#### ***Layer 2 Features***

- Flow control is enabled in PFE irrespective of interface configuration and the fix causes a very small amount of packet loss when a parameter related to an interface such as "interface description" on any port is changed. [PR1496766](#)
- On the QFX5000 line of switches, traffic imbalance might be observed if **hash-params** is not configured. [PR1514793](#)

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

### **Platform and Infrastructure**

- The PMTUD might not work for both IPv4 and IPv6 if the ingress Layer 3 interface is an IRB. [PR1442587](#)
- On the QFX5000 line of switches, the dcpfe process crashes due to the usage of data that is not null getting terminated. [PR1454527](#)
- On the QFX5100 switches, the interface output counter is double counted for self-generated traffic. [PR1462748](#)
- The sFlow could not work correctly if the received traffic goes out of more than one interface. [PR1475082](#)
- Egress port mirroring might not work when the analyzer port and mirrored port belong to a different FPC. [PR1477956](#)
- QFX5100: If more than one UDF filter/term is configured, then only the first filter/term will be programmed in hardware. This is due to SDK 6.5.16 upgrade. [PR1487679](#)
- Junos OS: EX2300 Series: High CPU load due to receipt of specific multicast packets on layer 2 interface (CVE-2020-1668). [PR1491905](#)
- ARP might not get refreshed after timeout. [PR1497209](#)
- Virtual Chassis is not stable with 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces. [PR1497563](#)
- Outbound SSH connection flaps or memory leaks during the push configuration to ephemeral database with high rate. [PR1497575](#)
- Traffic might get dropped if the aggregated Ethernet member interface is deleted or added, or a SFP of the aggregated Ethernet member interface is unplugged or plugged. [PR1497993](#)
- BFD sessions flap after deactivating or activating the aggregated Ethernet interface or executing GRES. [PR1500798](#)
- On the QFX5000 switches, ERPS might not work correctly. [PR1500825](#)
- The following error message might be observed during MPLS route add, change, or delete operation: mpls\_extra NULL. [PR1502385](#)
- The interface becomes physically down after changing to the FEC-none mode. [PR1502959](#)
- LLDP is not acquired when native-vlan-id and tagged VLAN-ID are the same on a port. [PR1504354](#)
- "Media type" in **show interface** command is displayed as "Fiber" for SFP-10G-T. [PR1504630](#)
- The l2cpd process might crash if the ERP configuration is added or removed, and the l2cpd process is restarted. [PR1505710](#)
- The archival function might fail in certain conditions. [PR1507044](#)
- The fxpc may crash and restart with a fxpc core file created while installing image through ZTP. [PR1508611](#)

- Traffic might be affected on QFX10002/QFX10008/QFX10016 platform. [PR1509220](#)
- ARP replies might be flooded through the EVPN-VXLAN network as unknown unicast ARP reply. [PR1510329](#)
- The output VLAN push might not work. [PR1510629](#)
- On the QFX5000 line of switches, multicast traffic loss is observed due to few multicast routes missing in the spine node. [PR1510794](#)
- The QFX10000-36Q line card used on QFX10008/QFX10016 platforms may fail to detect any QSFP. [PR1511155](#)
- In the VXLAN configuration, the firewall filters might not be loaded into the TCAM with the following message due to TCAM overflow after upgrading to Releases 18.1R3-S1, 18.2R1, and later : DFWE ERROR DFW: Cannot program filter. [PR1514710](#)
- The routes update might fail upon the HMC memory issue and traffic impact might be seen. [PR1515092](#)
- The 100-Gigabit Ethernet AOC non-breakout port might be auto-channelized to other speed. [PR1515487](#)
- The MAC learning might not work properly after multiple MTU changes on the access port in the VXLAN scenario. [PR1516653](#)
- The dcpfe process might crash due to memory leak. [PR1517030](#)
- The vgd process might generate a core file when the OVSDDB server restarts. [PR1518807](#)
- Traffic forwarding might be affected when adding, removing, or modifying the VLAN or VNI configurations such as VLAN-ID, VNI-ID, and Ingress-Replication command. [PR1519019](#)
- Output interface index in sFLOW packet are zero when transit traffic are observed on the IRB interface with VRRP enabled. [PR1521732](#)
- On the QFX10002, QFX10008, and QFX10016 line of switches, the following error message is observed during specific steps while clearing and loading the scaled configuration again:  
**PRDS\_SLU\_SAL:jprds\_slu\_sal\_update\_lrnrcnt(),1379:jprds\_slu\_sal\_update\_lrnrcnt call failed.** [PR1522852](#)
- Sampling with the rate limiter command enabled, crosses the sample rate 65535. [PR1525589](#)
- Packet loss is observed while validating the policer after restarting the chassis control. [PR1531095](#)
- High rate of ARP or NS packets might be observed between a device that runs Junos OS and host when the device that runs Junos OS receives an ARP or NS packet on an interface in transition. [PR1534796](#)
- Management Ethernet link down alarm seen while verifying system alarms in Virtual Chassis setup. [PR1538674](#)

### ***Routing Protocols***

- On the QFX 5100-48T-6Q Virtual Chassis or Virtual Chassis fan, the following error message is observed while copying image to the Virtual Chassis fan member and trying to downgrade the image: rcp for member 14, failed. [PR1486632](#)
- EX4300-MP/EX4600/QFX5000 Series: High CPU load due to receipt of specific layer 2 frames in EVPN-VXLAN deployment. (CVE-2020-1687) & High CPU load due to receipt of specific layer 2 frames when deployed in a Virtual Chassis configuration (CVE-2020-1689). [PR1495890](#)
- Scale of filters with egress-to-ingress command is enabled. [PR1514570](#)
- The rpd might report 100% CPU usage with BGP route damping enabled. [PR1514635](#)
- Enabling Ipv6 flow based Packet forwarding Engine hashing gives commit error. [PR1519018](#)
- Firewall "sample" configuration gives the warning as unsupported on QFX10002-36q and will not work. [PR1521763](#)
- On the QFX5000 line of switches, the fxpc process might crash if the VXLAN interface flaps. [PR1528490](#)

### ***User Interface and Configuration***

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

### ***Virtual Chassis***

- On QFX5120 and QFX5210 platforms unexpected storm control events might happen. [PR1519893](#)

## **Resolved Issues: 20.2R1**

### ***EVPN***

- The ESI of IRB interfaces does not update after autonomous-system number change if the interface is down. [PR1482790](#)
- QFX10002-60C EVPN-VXLAN multicast: The **show** command issued for the VTEP interface did not show mesh-group id. [PR1498052](#)
- The VXLAN function might be broken due to a timing issue. [PR1502357](#)

### ***Class of Service (CoS)***

- Traffic might be forwarded to an incorrect queue when fixed classifier is used. [PR1510365](#)

### ***General Routing***

- The following error message is generated while booting: **CMQFX: Error requesting SET BOOLEAN, illegal setting 66**. [PR1385954](#)
- The configuration statement **show chassis errors active detail** is not supported for QFX5000 platforms. [PR1386255](#)

- The 10G fiber interfaces might flap frequently when they are connected to other vendor's switch. [PR1409448](#)
- The statement **show interface** indicates Media type: Fiber on QFX5100-48T running '-qfx-5e-' Junos OS image. [PR1419732](#)
- A vmcore is seen on QFX Series Virtual Chassis. [PR1421250](#)
- SFP-LX10 stay down until autonegotiate is disabled. [PR1423201](#)
- The default logical interfaces on channelized physical interfaces might not be created after ISSU/ISSR. [PR1439358](#)
- CRC error might be seen on the VCPs of the QFX5100 Virtual Chassis. [PR1449406](#)
- On QFX5000 no warning or error is shown when dual VLAN tag feature is configured on physical interface. [PR1450455](#)
- Members might stay disconnected from a QFX5120-32C and QFX5120-48T Virtual Chassis after a full-stack reboot. [PR1453399](#)
- Changing the VLAN name associated with access ports might prevent MAC addresses from being learned in an EVPN-VXLAN scenario. [PR1454095](#)
- The cosd crash might be observed if forwarding-class-set is directly applied on the child interface of an aggregated Ethernet interface. [PR1455357](#)
- Telemetry traffic might not be sent out when the telemetry server is reachable through a different routing instance. [PR1456282](#)
- Link up delay and traffic drop might be seen on mixed SP L2/L3 and EP L2 type configurations. [PR1456336](#)
- QFX5110 QSFP-100GBASE-SR4 made by the third party cannot link up. [PR1457266](#)
- An FPC might restart during runtime on the QFX10000 line of devices. [PR1464119](#)
- EPR iCRC errors in QFX10000 platforms might cause protocols to go down. [PR1466810](#)
- A few of DHCP INFORM packets specific to a particular VLAN might be taking the wrong resolve queue. [PR1467182](#)
- Traffic loss might be seen with framing errors or runts if MACsec is configured on EX4600/QFX5100 platforms. [PR1469663](#)
- The speed 10m might not be configured on the GE interface. [PR1471216](#)
- The traffic loss might occur when VTEP source interface is configured in multiple routing instances. [PR1471465](#)
- Egress ACL filter entries will be only 512 in Junos OS Release 19.4R1 on QFX5000. [PR1472206](#)
- The shaping of CoS does not work after reboot. [PR1472223](#)
- DSCP marking might not work as expected if the fixed classifiers are applied to interfaces on QFX5000/EX4600 platforms. [PR1472771](#)

- The detached interface in LAG might process the xSTP BPDUs. [PR1473313](#)
- On QFX5000, the **global-mac-table-aging-time** statement behavior with multi-homed EVPN-VXLAN ESI. [PR1473464](#)
- ERP might not come up properly when MSTP and ERP are enabled on the same interface. [PR1473610](#)
- The RIPv2 packets forwarded across a L2 circuit connection might be dropped. [PR1473685](#)
- Continuous error log messages might be raised on QFX5000 platforms in EVPN-VXLAN scenario. [PR1474545](#)
- L2 circuit might fail to communicate through VLAN 2 on QFX5000 platforms. [PR1474935](#)
- On QFX Series platforms the system might stop new MAC learning and have impact on Layer 2 traffic forwarding. [PR1475005](#)
- DAC cables are not being properly detected in Packet Forwarding Engine in QFX5200. [PR1475249](#)
- There might be a traffic drop on QFX5110 and QFX5120 switches acting as leaf switches in a multicast environment with VXLAN. [PR1475430](#)
- FPC major error is seen after system boot up or FPC restart. [PR1475851](#)
- QFX Series platforms are exhibiting invalid Packet Forwarding Engine PG counter pairs to copy, src 0xfffff80, dst 0. [PR1476829](#)
- Continuous error logs on the device: **prds\_ptc\_wait\_adoption\_status: PECHIP[1] PTC[1]: timeout on getting adoption valid bit[8] asserted.** [PR1477192](#)
- The default Virtual Chassis MAC persistence timer is incorrectly set to 20 seconds instead of 20 minutes. [PR1478905](#)
- The remaining interface might be still in down state even though the number of channelized interfaces is no more than 5. [PR1480480](#)
- ARP request packets for unknown host might get dropped in remote PE device in EVPN-VXLAN scenario. [PR1480776](#)
- On QFX10000 and QFX5000, in SP style configuration, BUM traffic incorrectly gets blocked, while disabling or enabling a different logical interface. [PR1482202](#)
- On QFX5110, whenever the autonegotiation is toggled on the interface, explicitly set the link-mode as well as the speed for the configuration to take effect. [PR1484715](#)
- The dcpfe core file might be seen with non-oversubscribed mode. [PR1485854](#)
- The 10GbE VCP ports will not be active in a QFX5100 Virtual Chassis scenario. [PR1486002](#)
- Virtual Chassis ports might go down in a mixed Virtual Chassis setup of QFX5100-24Q-2P/EX4300 and EX4600/EX4300. [PR1489985](#)
- After ISSU/ISSR, a port using SR4/LR4 optics might not come up. [PR1490799](#)
- BFD sessions start to flap when the firewall filter in the loopback0 is changed. [PR1491575](#)

- Traffic loss could be observed in a mixed Virtual Chassis setup of QFX5100 and EX4300. [PR1493258](#)
- Traffic loss could be seen in a MC-LAG scenario on QFX5120/EX4650. [PR1494507](#)
- SNMP polling for CPU utilization and CPU state of backup Routing Engine does not show in a two-member Virtual Chassis. [PR1495384](#)
- ARP do 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. [PR1497380](#)
- An lcmd core file might be generated on QFX52100-64C. [PR1497947](#)
- Traffic might get dropped if aggregated Ethernet member interface is deleted and then added or a SFP of the aggregated Ethernet member interface is unplugged/plugged. [PR1497993](#)
- On QFX5210, unexpected behavior is seen for Port LED after upgrade. [PR1498175](#)
- Inter-VNI/VRF and intra-VNI/VRF traffic is dropped between the CE devices when the interfaces connected between TOR and multihomed PE devices are disabled. [PR1498863](#)
- The l2cpd crash might be seen while adding or deleting ERP configuration and then restarting l2cpd. [PR1505710](#)
- ARP replies might be flooded through the EVPN-VxLAN network as unknown unicast ARP reply. [PR1510329](#)

#### ***High Availability (HA) and Resiliency***

- Unified ISSU will not be supported for QFX5000 for some versions. [PR1472183](#)

#### ***Interfaces and Chassis***

- The MC-LAG configuration-consistency ICL-config might fail after committing some changes. [PR1459201](#)
- Executing commit might hang up because dcd process gets stuck. [PR1470622](#)
- Commit error is not thrown when member link is added to multiple aggregation group with different interface specific options. [PR1475634](#)
- MC-LAG consistency check fails if multiple IRB units are configured with the same VRRP group. [PR1488681](#)
- Error message is not getting generated while verifying GRE limitation. [PR1495543](#)

#### ***Junos Fusion for Enterprise***

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

#### ***Layer 2 Ethernet Services***

- EVPN-VXLAN ERB - dhcp relay-source lo0.1 is not used when enabled with anycast legacy IRB. [PR1455076](#)

- Member links state might be asynchronized on a connection between PE and CE devices in an EVPN A/A scenario. [PR1463791](#)
- Issues with DHCPv6 relay processing confirm and reply packets. [PR1496220](#)

### **Layer 2 Features**

- MAC learning might not work correctly on QFX5120. [PR1441186](#)
- The LLDP function might fail when a Juniper Networks device connects to a non-Juniper one. [PR1462171](#)
- A few MAC addresses might be missing from the MAC table in software on QFX5000 platform. [PR1467466](#)
- On QFX5120 switches QinQ, the third VLAN tag is not pushed onto the stack and SWAP is being done instead. [PR1469149](#)
- Traffic might be affected if composite next hop is enabled. [PR1474142](#)
- On QFX5200, MAC learning rate is degraded by 88 percent. [PR1494072](#)

### **MPLS**

- Traffic might silently get dropped or discarded on the PE device when the CE device sends traffic to the PE device and the destination is resolved with two LSPs through one upstream interface. [PR1475395](#)
- The traffic might be lost over QFX5100 switch acting as a transit PHP node in the MPLS network. [PR1477301](#)
- BGP session might keep flapping between two directly connected BGP peers because of the incorrect TCP-MSS in use. [PR1493431](#)

### **Platform and Infrastructure**

- The SLAX script might be lost after upgrading software. [PR1479803](#)
- Traceroute monitor with mtr version v.69 shows a false 10 percent loss. [PR1493824](#)

### **Routing Protocols**

- OSPF VRF sessions take a long time to come up when the host table is full and host routes are in LPM table. [PR1358289](#)
- BGP IPv4 or IPv6 convergence and RIB install/delete time degraded in Junos OS Release 19.1R1 and later mainline releases. [PR1414121](#)
- PIM (S,G) joins can cause MSDP to incorrectly announce source-active messages in some cases. [PR1443713](#)
- CRC errors might be seen on QFX5100 Virtual Chassis. [PR1444845](#)
- The core files might occur during adding or removing EVPN Type 5 routing instance. [PR1455547](#)
- [pfe\_loadbalance] [pfeloadtag] flows not falling back to single link when inactivity-interval is set higher than IFG. [PR1471729](#)



- Traffic might not be forwarded over ECMP link in EVPN-VXLAN scenario. [PR1475819](#)
- ARP packets are always sent to CPU regardless of whether the storm-control is activated. [PR1476708](#)
- GRE transit traffic is not forwarded in VRRP scenario. [PR1477073](#)
- MUX State in LACP interface does not go to "collecting and distributing" and remains attached after enabling the ae interface. [PR1484523](#)
- FPC might go to "NotPrnt" state after upgrading with non-QFX5100-24Q image in a Virtual Chassis/Virtual Chassis fabric setup. [PR1485612](#)
- CPU port queue gets full due to excessive pause frames being received on interfaces. This causes control packets from the CPU to all ports to be dropped. [PR1487707](#)
- The BGP route-target family might prevent RR from reflecting L2 VPN and L3 VPN routes. [PR1492743](#)
- The rpd might crash on QFX10000 due to rpd resolver problem of INH. [PR1494005](#)
- Firewall filter might not work in certain conditions under Virtual Chassis setup. [PR1497133](#)
- Traffic drop might be observed after modifying FBF firewall filter. [PR1499918](#)
- Change in x-path output for value "input-updates" in `show bgp neighbors`. [PR1504399](#)

#### SEE ALSO

[What's New | 231](#)

[What's Changed | 256](#)

[Known Limitations | 260](#)

[Open Issues | 262](#)

[Documentation Updates | 280](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 documentation for the QFX Series Switches.

#### SEE ALSO

[What's New | 231](#)

[What's Changed | 256](#)

[Known Limitations | 260](#)

[Open Issues | 262](#)

[Resolved Issues | 267](#)

[Migration, Upgrade, and Downgrade Instructions | 281](#)

## Migration, Upgrade, and Downgrade Instructions

### IN THIS SECTION

- [Upgrading Software on QFX Series Switches | 281](#)
- [Installing the Software on QFX10002-60C Switches | 284](#)
- [Installing the Software on QFX10002 Switches | 284](#)
- [Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches | 285](#)
- [Installing the Software on QFX10008 and QFX10016 Switches | 287](#)
- [Performing a Unified ISSU | 291](#)
- [Preparing the Switch for Software Installation | 292](#)
- [Upgrading the Software Using Unified ISSU | 292](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 294](#)

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

### Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.

3. Select **20.2** in the Release pull-down list to the right of the Software tab on the Download Software page.

4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.2 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Download the software to a local host.

8. Copy the software to the device or to your internal software distribution site.

9. Install the new jinstall package on the device.

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

Customers in the United States and Canada use the following command:

```
user@host> request system software add
source/jinstall-host-qfx-5-x86-64-20.2-R3.n-secure-signed.tgz reboot
```

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

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

- `http://hostname/pathname`
- `scp://hostname/pathname` (available only for Canada and U.S. version)

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

Rebooting occurs only if the upgrade is successful.

**NOTE:** After you install a Junos OS Release 20.2 **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-.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-.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-x86-64-20.2R3.n-secure-signed.tgz reboot
```

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

For example:

```
user@switch> request system software add
ftp://ftpserverserver/directory/jinstall-host-qfx-10-f-x86-64-20.2R3.n-secure-signed.tgz reboot
```

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

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

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

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

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

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

To install the software on re0:

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

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

For example:

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

To install the software on re1:

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

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

For example:

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

Reboot both Routing Engines.

For example:

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

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

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

### **Installing the Software on QFX10008 and QFX10016 Switches**



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

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



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

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

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

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

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

4. Disable nonstop-bridging:

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

5. Save the configuration change on both Routing Engines:

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

6. Exit the CLI configuration mode:

```
user@switch# exit
```

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

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

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

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

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

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.2R3.n-secure-signed.tgz
```

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

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

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

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

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

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

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

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

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

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

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

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

12. Transfer routing control to the backup Routing Engine:

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

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

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

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

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

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

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.2R3.n-secure-signed.tgz
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

## Performing a Unified ISSU

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

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

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 292](#)
- [Upgrading the Software Using Unified ISSU on page 292](#)

## Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:

**NOTE:** If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the **request system snapshot** command.

## Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:

- On the switch, enter:

```
user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, `jinstall-host-qfx-10-f-x86-64-20.1R2.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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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

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

#### SEE ALSO

- [What's New | 231](#)

---

- [What's Changed | 256](#)

---

- [Known Limitations | 260](#)

---

- [Open Issues | 262](#)

---

- [Resolved Issues | 267](#)

---

- [Documentation Updates | 280](#)

## Junos OS Release Notes for SRX Series

#### IN THIS SECTION

- [What's New | 296](#)
- [What's Changed | 306](#)
- [Known Limitations | 315](#)
- [Open Issues | 317](#)
- [Resolved Issues | 319](#)
- [Documentation Updates | 328](#)
- [Migration, Upgrade, and Downgrade Instructions | 328](#)



These release notes accompany Junos OS Release 20.2R3 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

## What's New

### IN THIS SECTION

- [What's New in Release 20.2R3 | 296](#)
- [What's New in Release 20.2R2 | 296](#)
- [What's New in Release 20.2R1 | 296](#)

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

### What's New in Release 20.2R3

There are no new features in Junos OS Release 20.2R3 for the SRX Series devices.

### What's New in Release 20.2R2

There are no new features in Junos OS Release 20.2R2 for the SRX Series devices.

### What's New in Release 20.2R1

#### *Application Security*

- **AppQoE multihoming with active/active deployment (NFX150, NFX250, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and vSRX)**—Starting In Junos OS Release 20.2R1, AppQoE is enhanced to support multihoming with active/active deployment. Previously, AppQoE supported multihoming with active/standby deployment.

In active/active deployment, the spoke device connects to multiple hub devices. Application traffic can transit through any of the hub devices if the link to the hub device meets SLA requirements. Application traffic can switch seamlessly between the hub devices in case of service-level agreement (SLA) violation or the active hub device is not responding.

To support active/active mode, you must enable the BGP multipath to allow the device to select multiple equal-cost BGP paths to reach a given destination.

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

- **Packet capture of unknown application traffic (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 20.2R1, we've added new capability to your security device that allows you to capture unknown application traffic.

Once you have configured the packet capture options on your security device, the unknown application traffic information is gathered and stored on the device in a packet capture file (**.pcap**). You can use the packet capture of an unknown application to define a new custom application signature. You can use this custom application signature in a security policy to manage the application traffic more efficiently.

You can also send the **.pcap** file to Juniper Networks in cases where the traffic is incorrectly classified, or to request for the creation of an application signature.

[See [Application Identification](#).]

- **Application Quality of Experience (SRX4600)**—Starting in Junos OS Release 20.2R1, the SRX4600 supports AppQoE functionality. AppQoE enhances the user experience at the application level by monitoring the performance of business-critical applications. Based on the score, AppQoE selects the best possible link for that application traffic to meet performance requirements specified in the service-level agreement (SLA).

The SRX4600 supports AppQoE in both the hub-and-spoke and the full mesh topologies.

AppQoE support is already available on SRX300, SRX320, SRX340, SRX345, SRX550HM, SRX1500, SRX4100, SRX4200, and vSRX.

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

### ***Authentication and Access Control***

- **Support to view user identify information in JIMS Active Directory (SRX Series)**— Starting in Junos OS Release 20.2R1, you can search and view user identity information such as logged users, connected devices and group list from Juniper Identity Management Service (JIMS) and Active Directory (AD) domain. The SRX Series device relies on JIMS to obtain user identity information.

You can search the user identity information and validate the authentication source to provide access to the device. You can request JIMS to retrieve the group list for the Active Directory domain for identity information of an individual user.

[See [Configure Juniper Identity Management Service to Obtain User Identity Information.](#)]

### ***Flow-Based and Packet-Based Processing***

- **IOC NP-cache scaling increased (SRX4600, SRX5000 line of devices)**—Starting in Junos OS Release 20.2R1, we have increased the number of hash table entries for IOC3 from 2 million to 20 million wings, for IOC4 from 2 million to 10 million wings on SRX5000 line of devices and for IOC on SRX4600 from 2 million to 5 million wings.

[See [Express Path.](#)]

### ***General Packet Radio Switching (GPRS)***

- **Support for Must-IE check and IE removal for GTPv1 and GTPv2 (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Release 20.2R1, Junos OS supports the following information element (IE) enforcement functions for GTPv1 and GTPv2:
  - **Must-IE check:** Use this function to check for the presence of IEs in GTPv1-C and GTPv2-C messages that helps to verify message integrity. The device check for the presence of Must-IEs of specific GTP messages and forwards the messages only if Must-IEs are present.
  - **IE removal:** Use this function to remove IEs from GTPv1-C and GTPv2-C. This function helps to retain interoperability between Second-Generation Partnership Project (2GPP) and Third-Generation Partnership Project (3GPP) networks.

[See [Example: Configure Must-IE check for GTPv1 and GTPv2](#), and [Example: Configure IE removal for GTPv1 and GTPv2.](#)]

### ***Intrusion Detection and Prevention (IDP)***

- **Policy-based threat profile for IDP (SRX Series)**—Starting from Junos OS Release 20.2R1, you can configure IDP rules with threat profiles to define attacker IP and target IP feeds.

When traffic matches the feed data, IDP provides feed update to add the IP information in the Security Intelligence (SecIntel) module.

This feature allows the SRX Series device to identify threats, and propagate intelligence for real-time enforcement and provides the ability to perform endpoint classification.

[See [IDP Policy Rules and IDP Rule Bases](#), [security-intelligence](#), and [Encrypted Traffic Analysis Overview.](#)]

- **Signature Language Constructs (SRX Series)**—Starting in Junos OS 20.2R1, the following signature language constructs are supported in the IDP engine code to write more efficient signatures that help reduce false attacks:
  - Byte extract
  - Byte test
  - Byte jump
  - Byte math
  - Is-data-at
  - Detection filter

[See [IDP Signature Language Enhancements](#).]

### *Junos Telemetry Interface*

- **Packet Forwarding Engine and Routing Engine sensor support on JTI (SRX5400, SRX5600, and SRX5800)**—Junos OS Release 20.2R1 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 \(Junos Telemetry Interface\)](#).]

### *Juniper Extension Toolkit (JET)*

- **Python 3 support for JET (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS can use Python 3 to execute JET scripts. To enable unsigned JET Python applications that support Python 3 to run on devices running Junos OS, use the `set system scripts language python3` command.

[See [language \(Scripts\)](#), [Develop Off-Device JET Applications](#), and [Develop On-Device JET Applications](#).]

### *J-Web*

- **Improved VPN usability (SRX Series)**—Starting in Junos OS Release 20.2R1, we've refreshed the IPsec VPN page. You can see a new improved site-to-site VPN workflow configuration.

[See [About the IPsec VPN Page](#).]

- **Pass-through tunnel inspection is supported in TAP mode (SRX 300 line of devices, SRX550M, SRX1500, SRX4100, and SRX4200)**—Starting in Junos OS Release 20.2R1, the J-Web Setup Wizard TAP mode supports pass-through tunnel inspection. This allows the SRX Series device to inspect pass-through traffic over an IP-IP tunnel or GRE tunnel.

[See [Start J-Web](#).]

- **HTTP X-Forwarded for header support in IDP (SRX Series)**—Starting in Junos OS Release 20.2R1, IDP supports the HTTP X-Forwarded option. When you enable this option, during traffic flow, IDP saves the source IP addresses (IPv4 or IPv6) from the HTTP and SMTP traffic contexts and displays them in the attack logs.

[See [About the Sensor Page](#).]

- **Enhancements to custom application signatures (SRX Series)**—Starting in Junos OS Release 20.2R1, we've enhanced custom applications signatures with the following:
  - By default, the priority for the custom application is set to Low. This allows a predefined application to take precedence. If you want to override a predefined application, you must set the priority to High.
  - Depth option is supported. Use this byte limit for Application Identification (App ID) to identify custom application patterns for applications running over TCP or UDP or Layer 7 applications.
  - Custom Application Byte Limit is supported in Global Settings. This byte limit helps in understanding when to stop the identification of custom applications.

[See [Add Application Signatures](#) and [Global Settings](#).]

## ATP Cloud

- **Support for adaptive threat profiling**—Starting in Junos OS Release 20.2R1, you can configure adaptive threat profiling in Juniper Sky ATP. Adaptive Threat Profiling allows SRX Series devices to generate, propagate, and consume threat feeds based on their own advanced detection and policy-match events. You can generate adaptive threat profiling feeds with traditional policies, unified policies with application identification (AppID) or URL-based match criteria, and IDP. Navigate to **Configure > Adaptive Threat Profiling** in Juniper Sky ATP UI to configure adaptive threat profiling.

[See [Adaptive Threat Profiling Overview](#) and [Add Threat Feed for Adaptive Threat Profiling](#).]

- **Support for encrypted traffic analysis**—Starting in Junos OS Release 20.2R1, encrypted traffic analysis is supported in Juniper Networks Sky ATP. Encrypted traffic analysis helps you to detect malicious threats that are hidden in encrypted traffic without intercepting and decrypting the traffic. Navigate to **Monitor > Encrypted Traffic** in Juniper Sky ATP UI to view detailed information about encrypted traffic analysis-based detections. To configure encrypted traffic analysis, use the **security-metadata-streaming** command at **[edit services]** hierarchy level. Use the **show services security-metadata-streaming statistics** command to view the statistics of the sessions.

[See [Encrypted Traffic Analysis Overview](#) and [Encrypted Traffic Analysis Details](#).]

## Logical Systems and Tenant Systems

- **Support for user firewall UAC authentication entries in shared mode for logical systems and tenant systems (SRX Series)**—Starting in Junos OS Release 20.2R1, logical systems and tenant systems support user firewall authentication with Unified Access Control (UAC).

[See [Understanding Integrated User Firewall Support in a Tenant System](#).]

- **User authentication support for tenant systems (SRX Series)**—Starting in Release 20.2R1, Junos OS introduces the following authentication support for tenant systems:
  - **address-assignment pools**: Creates centralized IPv4 and IPv6 address pools independent of the client applications that use the pools.
  - **access profiles**: Runs authentication and accounting requests.
  - **clear network-access aaa subscribers**: Clears AAA subscriber statistics and logs out subscribers. You can log out subscribers based on the username or on the subscriber session identifier.

[See [Firewall Authentication for Tenant Systems](#).]

### **Multicast**

- **Strict packet order for multicast traffic (SRX345 and SRX1500)**—Starting in Junos OS Release 20.2R1, we have introduced a new mechanism to maintain multicast traffic order and resolve packet drop issue. Use the **strict-packet-order** command at the **[edit security flow]** hierarchy level to maintain the packet order.

As part of this enhancement, you can configure the multicast route next-hop resolve attempts. When a multicast route next-hop resolve is unsuccessful, the SRX Series device attempts to resolve the next-hop route based on the specified retry counts. Use the **multicast-nh-resolve-retry** command at the **[edit security flow]** hierarchy level to specify the number of retry counts.

[See [flow](#).]

### **Network Address Translation (NAT)**

- **Increased port block allocation size (SRX5000 line of devices with SPC2 and SPC3 cards)**—we've increased the port block allocation size so you can store more log files in the log server.
  - When you disable **interim log**, you can increase the size of port block allocation from 64 to 8 .
  - When you enable **interim log**, you can increase the size of port block allocation from 128 to 8.

If you configure the port block allocation size less than 8, the system displays the warning message **warning: To save system memory, the block size is recommended to be no less than 8.**

[See [Guidelines for Configuring Secured Port Block Allocation](#) and [Configure Port Block Allocation Size](#).]

### **Network Management and Monitoring**

- **NETCONF sessions over outbound HTTPS (EX Series, MX Series, PTX1000, PTX3000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX Series, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.2R1, the Junos OS with upgraded FreeBSD software image includes a Juniper Extension Toolkit (JET) application that supports establishing a NETCONF session using outbound HTTPS. The JET application establishes a persistent HTTPS connection with a gRPC server over a TLS-encrypted gRPC session and authenticates the NETCONF client using an X.509 digital certificate. A NETCONF session over outbound HTTPS enables you to remotely manage devices that might not be accessible through other protocols, for example, if the device is behind a firewall.

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

- **Python 3 support for YANG scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. Junos OS does not support using Python 2.7 to execute YANG Python scripts as of this release.

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

- **Traffic log enhancement (SRX Series)**—Starting in Junos OS Release 20.2R1, we've enhanced the traffic log by supporting:

- Escape in stream log forwarding and on-box reporting to avoid parsing errors. Stream mode supports escape in **sd-syslog** and **binary** format. Event mode supports escape only in **binary** format.
- Different security log transport options for different streams.
- Stream-event mode.
- Increased maximum length of the stream mode **sd-syslog** format syslog message to 4\*1472 bytes.
- Different source addresses for different streams.
- Year and millisecond in timestamps.

[See [log \(Security\)](#) and [mode \(Security Log\)](#).]

- **CPU usage monitoring (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.2R1, you can use the following operational commands to monitor the average CPU usage information for the last minute, hour, or day of an SPC3 card:
  - ***show security monitoring performance spu summary fpc fpc-slot-number pic pic-slot-number***
  - ***show security monitoring performance spu summary fpc fpc-slot-number pic pic-slot-number thread thread-number***

You can monitor the CPU usage information only when the PIC is online.

We've introduced the new SNMP MIBs **jnxJsSPUMonitoringSPUThreadsNumber**, **jnxJsSPUMonitoringSPUThreadIndex**, **jnxJsSPUMonitoringSPUThreadLastMinUsage**, **jnxJsSPUMonitoringSPUThreadLastHourUsage**, and **jnxJsSPUMonitoringSPUThreadLastDayUsage** to monitor the CPU usage information of an SPC3 card.

[See [show snmp mib](#) and [show security monitoring performance spu](#).]



### Platform and Infrastructure

- **Support for Application Quality of Experience (AppQoE) (SRX4600)**—Starting in Junos OS Release 20.2R1, AppQoE is supported on SRX4600 devices along with SRX300, SRX320, SRX340, SRX345, SRX550M, SRX4100, and SRX4200 devices.

[See [Security Policy for Controlling Traffic for VRF Routing-Instance](#), [Flow Management in SRX Series Devices Using VRF Routing-Instance](#), [Understanding ALG Support for VRF Routing-Instance](#), and [Network Address Translation for VRF Routing-Instance](#).]

### Port Security

- **Media Access Control Security (MACsec) (SRX380)**—Starting in Junos OS Release 20.2R1, MACsec is supported on high availability (HA) control and fabric ports of SRX380 devices in chassis cluster mode. MACsec provides secure communication for almost all types of Layer 2 traffic on Ethernet links. MACsec is capable of identifying and preventing most security threats at Layer 2 and can be used in combination with other security protocols to provide end-to-end network security. MACsec is standardized in IEEE 802.1AE.

[See [Media Access Control Security \(MACsec\) on Chassis Cluster](#).]

### Security

- **Support for security feeds in security policies (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, you can add source and destination addresses to the security intelligence (SecIntel) profiles to generate security feeds in a security policy. You can accomplish this by configuring the **security-intelligence** configuration statements. After the feeds are generated, you can configure other security policies to use the feeds as a **dynamic-address** to match designated traffic and perform policy actions.

You can configure the **security-intelligence** configuration statements as permit, deny, or reject match conditions in a security policy at the following hierarchy levels:

```
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit
  application-services]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then deny application-services]
[edit security policies from-zone zone-name to-zone zone-name policy policy-name then reject application-services]
```

[See [security-intelligence](#) and [Encrypted Traffic Analysis Overview](#).]

- **Enhancements to configuring security policies (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, we have added advanced connection tracking options to security policies.

You can configure the **advanced-connection-tracking** command at the **[edit security zones security-zone zone name]** hierarchy levels to generate a connection track table using source IP, destination IP (optional), and destination port (optional) during session creation stage when traffic enters a given zone. This connection track mapping table also appears on the backup node in high availability (HA) pair.

You can configure the **advanced-connection-tracking** option under **[edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit]** to mandate that traffic matching given policy do a lookup in the *to-zone*'s connection track mapping table using the new session's key information. If there is no match, a new connection is not created.

[See [advanced-connection-tracking](#).]

### **Software Installation and Upgrade**

- **Zero-touch provisioning (ZTP) enhancements to support both DHCP options and phone-home client (SRX300, SRX320, SRX340, SRX345, SRX550 HM, and SRX1500)**—Starting in Junos OS Release 20.2R1, you can use zero-touch provisioning with DHCP options or the phone-home client to provision your device. As part of the factory default configuration, both ZTP and the phone-home client are included and are running at the same time when the device boots up in factory-default mode. ZTP with DHCP options is the first priority for provisioning. The device checks for DHCP bindings, and if there are DHCP bindings, but the DHCP bindings are not given the necessary ZTP-related options, (such as file server, and at least one image file or configuration file) the phone-home client will take over the provisioning process.

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

### **Unified Threat Management (UTM)**

- **UTM CLI test commands for Web Filtering and antispam feature (SRX Series)**— Starting in Release 20.2R1, Junos OS introduces the following test commands that help you to configure the Enhanced Web Filtering:
  - **test security utm enhanced-web-filtering url-check <test-url>**: Checks the category of a test string.
  - **test security utm web-filtering profile <profile-name><test-url>**: Checks the reputation of a test string.

Junos OS introduces the following test command for the antispam feature:

- **test security utm anti-spam ip-check <test-IP>**: Checks whether the IP address is a spam source.

[See [Unified Threat Management User Guide](#).]

- **CDF mode and inline-tap mode for AV**—Starting in Release 20.2R1, Junos OS introduces continuous delivery function (CDF) and inline-tap mode at the existing **[edit security utm default-configuration anti-virus]** hierarchy level. Continuous delivery function holds the last packet and sends out the other packets. This reduces system memory usage and speeds up the traffic. Inline-tap mode permits the traffic even if it is infected. Use inline-tap mode to check the antivirus feature without blocking or modifying the traffic.

[See [Unified Threat Management User Guide](#).]

- **Safe search enhancement for Web filtering (SRX Series and vSRX)**—Starting in Junos OS Release 20.2R1, we've introduced safe search UTM Web filtering on well-known search engines. This safe search enhancement enforces the safest Web browsing mode available, by default. You can disable the safe search option at the Web filtering-level and profile-level configurations. You can also block search engine

cache on the well-known search engines. By blocking the search engine cache, you can hide your Web-browsing activities from other users if you are a part of an organization that has multiple Web users in educational, financial, health-care, banking, and corporate segments.

[See [Safe Search Enhancement for Web Filtering](#), [feature-profile](#), [websense-redirect](#), and [juniper-local](#).]

## SEE ALSO

[What's Changed | 306](#)

[Known Limitations | 315](#)

[Open Issues | 317](#)

[Resolved Issues | 319](#)

[Documentation Updates | 328](#)

[Migration, Upgrade, and Downgrade Instructions | 328](#)

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 306](#)
- [What's Changed in Release 20.2R2 | 308](#)
- [What's Changed in Release 20.2R1-S1 | 310](#)
- [What's Changed in Release 20.2R1 | 310](#)

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

### What's Changed in Release 20.2R3

#### *Flow-Based and Packet-Based Processing*

- On SRX Series devices in earlier releases, when the session table was full there was no alarm set to indicate this. Starting from this release, when the percent of flow session table utilization is 95% on FPC and PIC, an alarm message ? Flow session table is almost full on FPC <number> PIC <number>? is set. Similarly, when the percent of DCP session table utilization is 95% on FPC and PIC, an alarm message ? DCP session table is almost full on FPC <number> PIC <number>? is set.

- **Self-generated IKE packets chooses outgoing interface matching source IP Address (SRX Series)** – A self-generated Internet Key Exchange (IKE) packet always select the ECMP outgoing interface that matches source IP address. Note that filter-based forwarding for self-generated traffic with rerouting is not supported.

#### *Junos OS XML API and Scripting*

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

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

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

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

#### *Network Management and Monitoring*

- **Changes to `<commit>` RPC responses in RFC-compliant NETCONF sessions (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—When you configure the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server's response for `<commit>` operations includes the following changes:
  - If a successful `<commit>` operation returns a response with one or more warnings, the warnings are redirected to the system log file, in addition to being omitted from the response.
  - The NETCONF server response emits the `<source-daemon>` element as a child of the `<error-info>` element instead of the `<rpc-error>` element.
  - If you also configure the `flatten-commit-results` statement at the `[edit system services netconf]` hierarchy level, the NETCONF server suppresses any `<commit-results>` XML subtree in the response and emits only an `<ok>` or `<rpc-error>` element.

[See [Configuring RFC-Compliant NETCONF Sessions.](#)]

### **User Interface and Configuration**

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

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

## **What's Changed in Release 20.2R2**

### **J-Web**

- **Change in the J-Web browser tab title (SRX Series)**—The J-Web browser tab title displays the device model and the hostname. The same details are 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 is not configured, you can see the host URL or IP address in the J-Web browser tab title. 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 20.2R2, you can configure the port block allocation size of 1 through 64512. To save system memory, the recommended port block allocation size is 64. If you configure the port block allocation with a size 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 earlier releases, you can configure port block allocation size of 1 through 64512 on SRX5400, SRX5600, and SRX5800 devices only.

[See [Configure Port Block Allocation Size.](#)]

### **Platform and Infrastructure**

- **Support for fully qualified domain name (FQDN) for log server (SRX Series)**—Starting in Junos OS Release, you can configure TTL value for a DNS server cache with hostname or IP address.

[See [Configuring the TTL Value for DNS Server Caching.](#)]

### **Routing Protocols**

- **Advertising 32 secondary loopback addresses to traffic engineering database as prefixes (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—We've made changes to export multiple loopback addresses to the `Isdist.0` and `Isdist.1` routing tables as prefixes. This eliminates the issue of advertising secondary loopback addresses as router IDs instead of prefixes. In earlier releases, multiple secondary loopback addresses in the traffic engineering database were added to the `Isdist.0` and `Isdist.1` routing tables as part of node characteristics and advertised them as the router ID.

### **System Log**

- **Support fully qualified domain name (FQDN) for log server (SRX Series)**—In Junos OS, you can configure TTL value for a DNS server cache with hostname or IP address.

[See [Configuring the TTL Value for DNS Server Caching.](#)]

### **VPNs**

- **The junos-ike package installed by default (SRX5000 Series devices)**— For SRX5000 Series devices with RE3 installed, the `junos-ike` package is installed by default. As a result, `iked` and `ikemd` process runs on the Routing Engine by default instead of IPsec key management daemon (`kmd`). In earlier Junos OS Releases, `junos-ike` package is an optional package for SRX5000 Series devices with RE3 and IPsec Key Management Daemon (`KMD`) runs by default.

[See [Enabling IPsec VPN Feature Set on SRX5K-SPC3 Services Processing Card.](#)]

- **IKE Index displayed in show security ipsec security-associations detail Output (SRX5400,SRX5600, SRX5800)**— When you execute the `show security ipsec security-associations detail` command, a new output field **IKE SA Index** corresponding to every IPsec Security Association (SA) within a tunnel is displayed under each IPsec SA information.

[See [show security ipsec security-associations.](#)]

## What's Changed in Release 20.2R1-S1

### Network Address Translation (NAT)

- **Port block allocation support (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS 20.2R2, you can configure the port block allocation size of 1 through 64512. To save system memory, the recommended port block allocation size is 64. If you configure the port block allocation with a size 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 earlier releases, you can configure port block allocation size of 1 through 64512 on SRX5400, SRX5600, and SRX5800 devices only.

[See [Configure Port Block Allocation Size](#).]

## What's Changed in Release 20.2R1

### Application Security

- Junos OS Release 20.2R1 introduces a new CLI configuration statement **depth** under **set services application-identification application *application-name* over application signature *signature-name* member *number*** hierarchy. You can use this configuration statement to specify the byte limit for application identification (AppID) to identify the custom application pattern for the applications running over TCP or UDP or Layer 7 applications.

Starting in Junos OS Release 20.2R1, you can display the configured **depth** value in J-Web using the **show services application-identification application detail** command.

```
user@host> show services application-identification application detail application-1
```

```
Application Name: test
Application type: application-1
Description: N/A
Application ID: 16777221
Priority: high
Order: 65500
Disabled: No
Cacheable: No
Activation Date: N/A
Last Modified: N/A
Underlying consolidated Protocols/ports application is dependent on:
  Protocols:
    Protocol: junos:HTTP / 67
    Protocol: junos:UDP / 216
    Protocol: junos:TCP / 205
```

```

Protocol: junos:NET-PROXY / 2629
Protocol: junos:SPDY / 1469
Protocol: junos:SSL / 199
Protocol: junos:LIBJINGLE-PSEUDOTCP / 3237
Protocol: junos:STUN / 201
Protocol: junos:HTTPS / 68
Protocol: junos:HTTP / 67
Protocol: junos:HTTP2 / 2553
Protocol: junos:HTTP-TUNNEL / 750
Protocol: junos:HTTP-PROXY / 2956
Protocol: junos:HAPROXY / 3331
Protocol: junos:COTP / 22
TCP Ports:
Port: 80
Port: 3128
Port: 8000
Port: 8080
Layer-7 Immediate Protocol(s):
Protocol: HTTP / 67
Signature: fgsm
Port range: N/A
Member(s): 1
Member m01
    Depth: 4
Context: http-get-url-parsed-param-parsed
Pattern: ads
Direction: CTS

```

In the above sample, you can see the configured value of the depth is displayed as 4.

[See [Application Identification](#)].

- Starting in Junos OS Release 20.2R1, the syntax of the commands used for displaying the SLA profile details is changed as following:

Syntax in Junos OS Release Prior to 20.2R1	Syntax in Junos OS Release 20.2R1 or Later
<pre>show security advance-policy-based-routing sla profile sla-profile-name application application-name destination-group-name destination-group-name status</pre>	<pre>show security advance-policy-based-routing sla profile profile-name application application-name next-hop next-hop-id status</pre>
<pre>show security advance-policy-based-routing sla profile sla-profile-name application application-name destination-group-name destination-group-name</pre>	<pre>show security advance-policy-based-routing sla profile profile-name application application-name next-hop next-hop-id</pre>



[See *show security advance-policy-based-routing sla profile (Application Name)*, *show security advance-policy-based-routing sla profile (Next-Hop)*, and *show security advance-policy-based-routing sla profile (Status)*.]

### Class of Service (CoS)

- We've corrected the output of the **show class-of-service interface | display xml** command that appeared as `<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>` to `<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>`

### Flow-Based and Packet-Based Processing

- **ECMP load balancing in chassis cluster (SRX Series)**—Starting in Junos OS Release 20.2R1, in a chassis cluster setup, to avoid reroute flapping between primary and secondary sessions, add a logic to skip the reroute for backup sessions. But reroute can change the chassis interface of a flow session, so the session can be changed from backup session to primary session after reroute. You cannot skip reroute for such a session.

When you change the logic, the session reroute skips only the packets received from the chassis interface. So we can make sure the session continues as the backup session even after you reroute and change the out-going interface. Otherwise, reroute cannot be skipped for backup sessions.

- **Simplified HA (SRX Series)**—Starting in Junos OS Release 20.2R1, on SRX Series devices in a simplified HA setup, when you clear the session using the **clear security flow session** command, some warm sessions exist for an extended duration. To clear these warm sessions, a new CLI command **clear security flow session session-state warm** is introduced.

`clear security flow session all`

### Juniper Extension Toolkit (JET)

- **PASS keyword required for Python 3 JET applications (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you are writing a JET application using Python 3, include the **PASS** keyword in the Exception block of the script. Otherwise, the application throws an exception when you attempt to run it.

[See [Develop Off-Device JET Applications](#) and [Develop On-Device JET Applications](#).]

- **Updates to IDL for RIB service API bandwidth field (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The IDL for the RouteGateway RIB service API has been updated to document additional rules for the **bandwidth** field. You must set **bandwidth** only if a next hop has more than one gateway, and if you set it for one gateway on a next hop, you must set it for all gateways. If you set **bandwidth** when there is only a single usable gateway, it is ignored. If you set **bandwidth** for one or more gateways but not all gateways on a next hop, you see the error code **BANDWIDTH\_USAGE\_INVALID**.

[See [Juniper EngNet](#).]

### Juniper Sky ATP

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 20.2R1, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Sky Advanced Threat Prevention (ATP).

### Network Management and Monitoring

- **Request support information for IPsec VPN (SRX Series)**—Starting in Junos OS Release 20.2R1, we've introduced the CLI `ipsec-vpn` option to the `request support information security-components` command. This new option displays all the configuration, states, and statistics information necessary for debugging IPsec VPN related issues.

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

- **Junos OS only supports using Python 3 to execute YANG Python scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.2R1, Junos OS uses Python 3 to execute YANG action and translation scripts that are written in Python. In earlier releases, Junos OS uses Python 2.7 to execute these scripts.

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

### VPNs

- **New vendor ID for Internet Key Exchange (SRX Series)**—In Junos OS Release 20.2R1, we've introduced a new vendor ID **Juniper Networks** for Internet IKEv1 and IKEv2 which is advertised to the peer.

[See [Understanding IKE and IPsec Packet Processing](#).]

- **Change in CLI options help text description (SRX Series)**—Starting in Junos OS Release 20.2R1, we've changed the help text description as **NOT RECOMMENDED** for the following CLI options under `[edit security ike proposal proposal-name]`, `[edit security ike policy policy-name]`, `[edit security ipsec proposal proposal-name]`, and `[edit security ipsec policy policy-name]` hierarchies.

Hierarchy	CLI Options	Help Text Description
[edit security ike proposal <i>proposal-name</i> authentication-algorithm]	md5	NOT RECOMMENDED
	sha1	NOT RECOMMENDED
[edit security ike proposal <i>proposal-name</i> encryption-algorithm]	3des-cbc	NOT RECOMMENDED
	des-cbc	NOT RECOMMENDED

Hierarchy	CLI Options	Help Text Description
[set security ike proposal <i>proposal-name</i> dh-group]	group1	NOT RECOMMENDED
	group14	NOT RECOMMENDED
	group2	NOT RECOMMENDED
	group5	NOT RECOMMENDED
[edit security ike proposal <i>proposal-name</i> authentication-method]	dsa-signatures	NOT RECOMMENDED
[edit security ike policy <i>policy-name</i> proposal-set]	basic	NOT RECOMMENDED
	compatible	NOT RECOMMENDED
	standard	NOT RECOMMENDED
[edit security ipsec policy <i>policy-name</i> proposal-set]	basic	NOT RECOMMENDED
	compatible	NOT RECOMMENDED
	standard	NOT RECOMMENDED
[edit security ipsec proposal <i>proposal-name</i> encryption-algorithm]	3des-cbc	NOT RECOMMENDED
	des-cbc	NOT RECOMMENDED
[edit security ipsec proposal <i>proposal-name</i> authentication-algorithm]	hmac-md5-96	NOT RECOMMENDED
	hmac-sha1-96	NOT RECOMMENDED
[edit security ipsec policy <i>policy-name</i> perfect-forward-secrecy keys]	group1	NOT RECOMMENDED
	group2	NOT RECOMMENDED
	group5	NOT RECOMMENDED
	group14	NOT RECOMMENDED

[See [authentication-algorithm \(Security IPsec\)](#) and [encryption-algorithm \(Security IKE\)](#).]

- **Change in thread ID configuration (SRX Series)**—Starting in Junos OS Release 20.2R1, when you add, change, or delete the thread ID from distribution profile at [**edit security distribution-profile *profile-name* fpc slot-number pic slot-number thread-id**], all tunnels part of modified distribution profile anchored on modified SPU member of distribution profile are torn down and re-negotiated.

[See [distribution-profile](#).]

#### SEE ALSO

[What's New | 296](#)

[Known Limitations | 315](#)

[Open Issues | 317](#)

[Resolved Issues | 319](#)

[Documentation Updates | 328](#)

[Migration, Upgrade, and Downgrade Instructions | 328](#)

## Known Limitations

### IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 315](#)
- [J-Web | 316](#)
- [VPNs | 316](#)

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 the Routing Engine and Packet Forwarding Engine, some packets get missed in the PCAP files while using the JDPI unknown packet capture feature. [PR1491919](#)

- Committing a large number of custom applications with a single member, a single context, and a varying pattern might result in significant time taken for completion of commit. Commit status can be checked using `show services application-identification commit-status`. [PR1493127](#)

## J-Web

- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is autorefreshed. As a workaround, create a dynamic application as the last action while modifying the policy rule and click the Save button to avoid loss of configuration changes made to the policy rule. [PR1460214](#)
- For a spoke device in a hub-and-spoke topology, J-Web shows the VPN topology as Site to Site. [PR1495973](#)

## 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 is sent to the peer if traffic flows over even one of the tunnels for the given VPN object. [PR1366585](#)
- On the SRX5000 line of devices with an SPC3 card, sometimes IKE SA is not seen on the device when the st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). [PR1441411](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, with 60,000 tunnels up, when RGO failover happens while an IPsec and/or IKE rekey is in progress, those rekeying tunnels might go down and traffic loss might be seen until the tunnel is reestablished. [PR1471499](#)
- In SPC2 and SPC3 mixed-mode HA deployments, tunnel per second (TPS) is getting affected while dead peer detection (DPD) is being served on existing tunnels. This limitation is due to a large chunk of CPU being occupied by infrastructure (gencfg) used by IKED to synchronize its DPD state to the backup nodes. [PR1473482](#)
- On SRX Series devices, the accounting stop message is not being sent after deactivating the access profile under the security IKE gateway. [PR1485732](#)

## SEE ALSO

---

[What's New | 296](#)

---

[What's Changed | 306](#)

---

[Open Issues | 317](#)

---

[Resolved Issues | 319](#)

---

Documentation Updates | [328](#)

Migration, Upgrade, and Downgrade Instructions | [328](#)

## Open Issues

### IN THIS SECTION

- [Flow-Based and Packet-Based Processing | 317](#)
- [J-Web | 318](#)
- [Routing Policy and Firewall Filters | 318](#)
- [VPNs | 318](#)

Learn about open issues in this release for SRX Series devices.

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

### Flow-Based and Packet-Based Processing

- Use an antireplay window size of 512 for IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores \* 32 packets in one batch). Hence, there are no out-of-order packets with 512 antireplay window size. [PR1470637](#)
- You need to configure the default IPv6 route (egress is fxp0) if you use IPv6 GRE or IP-IP tunnel and dynamic route protocol (BGP, OSPF, and so on) in Layer 3 HA. Use the following configuration example (2010::1 is in the same sub network with fxp0):

- `set groups global routing-options rib inet6.0 static route 0::0/0 next-hop 2010::1`  
`set groups global routing-options rib inet6.0 static route 0::0/0 retain`  
`set groups global routing-options rib inet6.0 static route 0::0/0 no-readvertise`  
[PR1482616](#)

## J-Web

- On the SRX5000 line of devices, J-Web might not be responsive sometimes when you commit configuration changes after adding a new dynamic application while creating a new firewall rule. J-Web displays a warning while validating the configuration due to dynamic application or any other configuration changes. As a workaround, refresh the J-Web page. [PR1460001](#)
- Configuration of global settings options of IPsec VPN such as TCP encap profile, IPsec power mode and IKE package installation are not supported from J-Web. [PR1496439](#)

## Routing Policy and Firewall Filters

- When the cli `show security match-policy` command is used with `url-category` as a match item and the destination IP address cannot be divided by 3, an incorrect result may be returned. [PR1483251](#)

## VPNs

- In the output of the `show security ipsec inactive-tunnels` command, Tunnel Down Reason is not displayed as this functionality is not supported in Junos OS Release 18.2R2 and later. [PR1383329](#)
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, a new behavior has been introduced that differs from the behavior on the older SPC2 card. The SRX Series device with AutoVPN configuration can now accept multiple IPsec tunnels from a peer device (with the same source IP address and port number) using different IKE IDs. [PR1407356](#)
- On the SRX5000 line of devices with an SPC3 card, sometimes IKE SA is not seen on the device when the `st0` binding on the VPN configuration object is changed from one interface to another (for example, `st0.x` to `st0.y`). [PR1441411](#)
- Tunnel debugging configuration is not synchronized to the backup node. It needs to be configured again after RGO failover. [PR1450393](#)
- On the SRX5000 line of devices with SPC3 and SPC2 mixed mode, with a very large number of IKE peers (60,000) with dead peer detection (DPD) enabled, IPsec tunnels might flap in some cases when IKE and IPsec rekeys are happening at the same time. [PR1473523](#)
- Some TCP connections going through IPsec tunnels are getting stuck after RG1 failover. [PR1477184](#)

- During 10,000 tunnel ramp-up, sometimes, IKED generates a core file. [PR1479548](#)
- The SRX5000 line of devices with SPC3 was not supporting simultaneous IKE negotiation in Junos OS Releases 19.2, 19.3, 19.4, and 20.1. [PR1497297](#)

#### SEE ALSO

[What's New | 296](#)

[What's Changed | 306](#)

[Known Limitations | 315](#)

[Resolved Issues | 319](#)

[Documentation Updates | 328](#)

[Migration, Upgrade, and Downgrade Instructions | 328](#)

## Resolved Issues

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

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

### Resolved Issues: 20.2R3

#### *Chassis Clustering*

- Disabled node on SRX chassis cluster sends out ARP request packets. [PR1548173](#)

#### *Flow-Based and Packet-Based Processing*

- The `rst-invalidate-session` command does not work if configured together with the `no-sequence-check` command. [PR1541954](#)

#### *General Routing*

- The TCP packet might be dropped if syn-proxy protection is enabled. [PR1521325](#)
- On SRX Series devices with chassis cluster, high CPU usage might be seen due to the `llmd` process. [PR1521794](#)
- Certificate validation might fail when OCSP is used and the OCSP server is a dual-stack device. [PR1525924](#)
- On the SRX1500 device, the traffic rate shown in the CLI command is not accurate. [PR1527511](#)
- The MAC table is null in Layer 2 mode after one pass-through session is created successfully. [PR1528286](#)



- On SRX4100 and SRX4200 devices, four out of eight fans might not work. [PR1534706](#)
- The firewall filter SA and DA tags are not in the log messages as expected in port details. [PR1539338](#)
- Packet drop might be seen when a packet with destination port 0 is received on the SRX380 device. [PR1540414](#)
- The nsd process might crash when DNS-based allowlisting is configured under SSL proxy. [PR1542942](#)
- Need syslog to indicate signature download completion. [PR1545580](#)
- The flowd process might generate core files when the user changes the flow mode configuration to packet mode. [PR1546653](#)
- On SRX4100 and SRX4200, if PEM0 is removed, the output of jnxOperatingDescr.2 might be incomplete. [PR1547053](#)
- Advanced anti-malware file or email statistics does not get incremented with the latest PB version. [PR1547094](#)
- Continuous "LCC: ch\_cluster\_lcc\_set\_context:564: failed to lock chassis\_vmx mutex 11" chassisd process logs are generated. [PR1547953](#)
- Lcmd log "gw\_cb\_presence:136: PEM(slot = 0): error detecting presence ( fruid = 15, drv\_id = 30, status = -11 )" generates every second on the SRX4100 and SRX4200 devices. [PR1550249](#)
- The speed mismatch error is seen while trying to commit reth0 with together-options. [PR1553888](#)
- An IPFD core might be generated when using adaptive threat profiling. [PR1554556](#)
- The dumpdisklabel command fails with message "ERROR: Unknown platform srx550m". [PR1557311](#)
- The outbound-ssh routing-instance command output shows as unsupported. [PR1558808](#)
- Application identification unknown packet capture utility does not function on SRX Series devices when the enhanced-services mode is enabled. [PR1558812](#)
- High CPU usage on pkid process might be seen when the device is unable to connect to a particular CRL URL. [PR1560374](#)

### ***Interfaces and Chassis***

- When SRX Series devices receive proxy ARP requests on VRRP interfaces, SRX Series devices send ARP replies with the underlying interface MAC address. [PR1526851](#)
- Backup Routing Engine or backup node may stuck in bad status with improper "backup-router" configuration. [PR1530935](#)

### ***Intrusion Detection and Prevention (IDP)***

- The flowd or srxpfe process might generate core files during the idpd process commit on SRX Series devices. [PR1521682](#)
- IDP policy load might fail post image upgrade for 15.1x49 releases. [PR1546542](#)

### **J-Web**

- The "+" button is not shown in the J-Web interface menu. [PR1550755](#)

### **Platform and Infrastructure**

- Syslog reporting "PFE\_FLOWD\_SELFPING\_PACKET\_LOSS: Traffic impact: Selfping packets loss/err: 300 within 600 second" error messages in node 0 and node 1 control panel. [PR1522130](#)
- The commit might not fail as expected when the reth interface is deleted. [PR1538273](#)

### **Routing Policy and Firewall Filters**

- The flowd or srxpfe process might stop when an SRX Series or NFX Series device running Junos OS Release 18.2R1 or later supports the unified policy feature. [PR1544554](#)
- Traffic might be dropped unexpectedly when the url-category match condition is used on a security policy. [PR1546120](#)
- Global policies working with multi-zones cause high PFE CPU utilization. [PR1549366](#)
- On the SRX5000 line of devices, the secondary node might get stuck in performing ColdSync after a reboot, upgrade, or if ISSU is performed. [PR1558382](#)
- The traffic might be dropped due to inserting one global policy above others on SRX Series devices. [PR1558827](#)

### **Subscriber Access Management**

- Incorrect counter type (counter instead of gauge) is specified for some values in MIB jnxUserAAAMib. [PR1533900](#)

### **Unified Threat Management (UTM)**

- Stream buffer memory leak might happen when UTM is configured under unified policies. [PR1557278](#)
- UTM license expiry event loss may cause the device to not quit the advanced service mode and maximum-sessions is decreased by half. [PR1563874](#)

### **VPNs**

- IPsec SA is missing the keyword NULL after RG failover. [PR1507270](#)
- IPsec traffic might get dropped after RGO failover. [PR1522931](#)
- On all SRX Series devices using IPsec with NAT traversal, MTU size for the external interface might be changed after IPsec SA is re-established. [PR1530684](#)
- After the IPsec tunnel using policy-based VPN is overwritten by another VPN client, traffic using this IPsec tunnel will be dropped. [PR1546537](#)
- Traffic going through policy-based IPsec tunnel might be dropped after RGO failover. [PR1550232](#)
- A session might be closed when the session is created during the IPsec rekey. [PR1564444](#)

## Resolved Issues: 20.2R2

### *Application Layer Gateways (ALGs)*

- The srxpfe or mspmand process might crash if FTPS is enabled in a specific scenario. [PR1510678](#)

### *Flow-Based and Packet-Based Processing*

- The show security group-vpn server statistics |display XML is not in expected format. [PR1349959](#)
- With the 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](#)
- ECMP load balancing does not happen when RG1 node 0 is secondary. [PR1475853](#)
- On SRX4100 and SRX4200 devices with chassis cluster in transparent mode, when a failover occurs for RG1, the interface on the new secondary node flaps as expected to let the switch update its MAC address table. [PR1490291](#)
- Not able to clear the warm sessions on the peer SRX Series devices. [PR1493174](#)
- Outbound SSH connection flap or memory leak issue might be observed while pushing the configuration to the ephemeral DB with a high rate. [PR1497575](#)
- The srxpfe or flowd process might stop due to memory corruption within JDPI. [PR1500938](#)
- The downloads might permanently get stuck or not complete when TCP proxy is used on SRX Series devices. [PR1502977](#)
- Fabric interface might be monitored down after chassis cluster reboot. [PR1503075](#)
- SOF asymmetric scenario is not working with the phase 1 solution. [PR1507865](#)
- TAP mode behavior has been improved and the configuration has been greatly simplified. [PR1521066](#)
- In a dual CPE scenario, if the rule match is completed before application identification is done, AppQoE moves the session to other node. [PR1514973](#)
- VRRP does not work on the redundant Ethernet interface with a VLAN ID greater than 1023. [PR1515046](#)
- PCAP file generated using packet capture was improper on the SRX5000 line of devices. [PR1515691](#)
- A logic issue was corrected in SSL proxy that could lead to an srxpfe or flowd core file under load. [PR1516903](#)
- The PPPoE session does not come up after return to zero on SRX Series devices. [PR1518709](#)
- FQDN-based security log stream does not dynamically update the IP address. [PR1520071](#)
- Adaptive Threat Profiling would stop submitting new IP addresses to a feed after a limit of 10,000 has been reached. [PR1524284](#)

### ***Interfaces and Chassis***

- PPO IPv6 route does not work. [PR1495839](#)

### ***Intrusion Detection and Prevention (IDP)***

- 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***

- While creating a firewall policy rule, the list of available dynamic applications is empty in HA on the Select Dynamic Application page. [PR1490346](#)
- 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](#)

### ***Layer 2 Ethernet Services***

- DHCP might not work after performing request system zeroize or load factory-default on SRX Series devices. [PR1521704](#)

### ***Network Address Translation (NAT)***

- NAT PBA size 1 on SRX Series devices. [PR1525822](#)

### ***Platform and Infrastructure***

- Packets get dropped when the next hop is IRB over the LT interface. [PR1494594](#)

### ***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](#)
- The show security dynamic-address feed-name command could not list secprofiling feed. [PR1537714](#)

### ***Unified Threat Management (UTM)***

- UTM causes emails from outside to inside to not be received. [PR1523222](#)

### ***VPNs***

- On a 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](#)

## Resolved Issues: 20.2R1

### *Application Layer Gateways (ALGs)*

- RTSP data sessions are cleared unexpectedly during cold sync. [PR1468001](#)
- The flowd or srpxfe process might stop when an ALG creates a gate with an incorrect protocol value. [PR1474942](#)
- SIP messages that need to be fragmented might be dropped by SIP ALG. [PR1475031](#)
- FTPS traffic might get dropped on SRX Series or MX Series devices if FTP ALG is used. [PR1483834](#)

### *Authentication and Access Control*

- SRX Series: Unified Access Control (UAC) bypass vulnerability (CVE-2020-1637). [PR1475435](#)

### *Flow-Based and Packet-Based Processing*

- Command **show security pki local-certificate logical-system all** is not showing any output. [PR1414628](#)
- The trusted-ca and root-ca names or IDs should not be the same within an SSL proxy configuration. [PR1420859](#)
- Introduction of default inspection limits for application identification to optimize CPU usage and improve resistance to evasive applications. [PR1454180](#)
- TCP session might not time out properly upon receiving TCP RESET packet. [PR1467654](#)
- RPM test probe fails to show that round-trip time has been exceeded. [PR1471606](#)
- Support LLDP protocol on reth interface. [PR1473456](#)
- Certificate error when configuration is validated during Junos OS upgrade. [PR1474225](#)
- An unhealthy node might become primary in SRX4600 devices with chassis cluster scenario. [PR1474233](#)
- Packet drop might be observed on the SRX300 line of devices when adding or removing an interface from MACsec. [PR1474674](#)
- Stateful firewall rule configuration deletion might lead to memory leak. [PR1475220](#)
- The flowd or srpxfe process might stop when deleting user firewall local authentication table entry. [PR1477627](#)
- MPCs might stop when there is bulk route update failure in a corner case. [PR1478392](#)
- The nsd process pause might be seen during device reboots if dynamic application groups are configured in policy. [PR1478608](#)
- The flowd process core files might be seen when there is mixed NAT-T traffic or non-NAT-T traffic with PMI enabled. [PR1478812](#)
- When SRX5K-SPC3s or MX-SPC3s are installed in slots 0 or 1 in SRX5800 or MX960 devices, EMI radiated emissions are observed to be higher than regulatory compliance requirements. [PR1479001](#)

- The show mape rule statistics command might display negative values. [PR1479165](#)
- The wl-interface stays in ready status after you execute request chassis fpc restart command in Layer 2 mode. [PR1479396](#)
- Recent changes to JDPI's classification mechanism caused a considerable performance regression (more than 30 percent). [PR1479684](#)
- The flowd or srpxfe process might stop when advanced anti-malware service is used. [PR1480005](#)
- On Web proxy, memory leak in association hash table and DNS hash table. [PR1480760](#)
- The jsqsyncd process synchronizes its databases every second even there is no change. [PR1482428](#)
- The firewall Web authentication graphics have been updated. [PR1482433](#)
- IMAP curl sessions get stuck in the active state if AAMW IMAP block mode is configured. [PR1484692](#)
- The show chassis temperature-thresholds command displays extensive FPC 0 output. [PR1485224](#)
- The configuration **set chassis psu redundancy n-plus-n** needs support on in high availability (HA) mode. [PR1486746](#)
- Commit does not work after the installation through boot loader. [PR1487831](#)
- If a cluster ID of 16 or multiples of 16 is used, the chassis cluster might not come up. [PR1487951](#)
- CPU board inlet increases after OS upgrade from Junos OS Release 15.1X49 to Junos OS Release 18.x. [PR1488203](#)
- All interfaces remain in the down status after the SRX300 line of devices power up or reboot. [PR1488348](#)
- There is a risk of service interruption on all SRX Series devices with a dual stacked CA server. [PR1489249](#)
- GRE or IPSec tunnel might not come up when **set security flow no-local-favor-ecmp** command is configured. [PR1489276](#)
- Sometimes multiple flowd core files are generated on both nodes of chassis cluster at the same time when changing media MTU. [PR1489494](#)
- Continuous drops seen in control traffic, with high data queues in one SPC2 PIC. [PR1490216](#)
- Phone client stop seen while doing SRX345 device ZTP with CSO. [PR1496650](#)
- 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 capital characters for source-identity when using **show security match-policies** command. [PR1499090](#)
- J-Flow version 9 does not display correct outgoing interface for APBR traffic. [PR1502432](#)

- AppQoE support for dynamic-application. [PR1503400](#)
- The cfmd core observed when LTM is triggered for the session configured on ethernet-switching interface without bridge domain configuration. [PR1503696](#)

#### ***Intrusion Detection and Prevention (IDP)***

- Configuring anomaly occurs in CLI. [PR1490437](#)

#### ***J-Web***

- You cannot configure redundant PSU and power budget statistics on the SRX380 device that is in high availability (HA) mode through J-Web. [PR1493713](#)
- The J-Web users might not be able to configure PPPoE using PPPoE wizard. [PR1502657](#)

#### ***Layer 2 Ethernet Services***

- Member links state might be asynchronized on a connection between PE and CE devices in an EVPN active/active mode. [PR1463791](#)

#### ***Multiprotocol Label Switching (MPLS)***

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

#### ***Network Address Translation (NAT)***

- Issuing the `show security nat source paired-address` command might return an error. [PR1479824](#)

#### ***Network Management and Monitoring***

- The flowd or srpxfe process might stop immediately after committing the J-Flow version 9 configuration or after upgrading to affected releases. [PR1471524](#)
- SNMP trap coldStart agent-address becomes 0.0.0.0. [PR1473288](#)

#### ***Platform and Infrastructure***

- Modifying the REST configuration might cause the system to become unresponsive. [PR1461021](#)
- On SRX1500 and the SRX4000 line of devices, physically disconnecting the cable from fxp0 interface causes hardware monitor failure and redundancy group failover, when the device is the primary node in a chassis cluster. [PR1467376](#)
- The RGx might fail over after RG0 failover in a rare case. [PR1479255](#)
- The `/usr/libexec/ui/yang-pkg` and `/usr/libexec/ui/pyang` files not found in SRX Series devices during YANG installation. [PR1496577](#)

### ***Routing Policy and Firewall Filters***

- If a huge number of policies are configured on SRX Series devices and some policies are changed, the traffic that matches the changed policies might be dropped. [PR1454907](#)
- Support for dynamic tunnels on SRX Series devices was mistakenly removed. [PR1476530](#)
- TCP proxy was mistakenly engaged in unified policies when Web filtering was configured in potential match policies. [PR1492436](#)
- Traffic fails to hit the policies with matching source-end-user-profiles. [PR1505002](#)

### ***Routing Protocols***

- The rpd might stop when both instance-import and instance-export policies contain as-path-prepend action. [PR1471968](#)

### ***Unified Threat Management (UTM)***

- The utmd process might pause after deactivating UTM configuration with predefined category upgrading used. [PR1478825](#)

### ***VPNs***

- IKE SA does not get cleared and is showing very long lifetime. [PR1439338](#)
- IKED is treating all re-transmission of first IKE\_INIT request packets as new connections when acting as responder. [PR1460907](#)
- The iked might crash when the IKE SA expires and the IPsec tunnel of expired IKE SAs still exists. [PR1463501](#)
- The newly configured IPsec tunnels might be stuck in VPNM verify-path state in a tunnel scaled scenario. [PR1464353](#)
- IPsec tunnels might flap when one secondary node is coming online after reboot in SRX Series high availability environment. [PR1471243](#)
- The kmd process might crash continually after the chassis cluster failover in the IPsec ADVPN scenario. [PR1479738](#)
- On SRX4200 device, 35 percent of drop is seen in all TPS cases. [PR1481625](#)
- 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 ID to avoid confusion. [PR1493368](#)
- Issue with XML rpc **show security ipsec tunnel-distribution summary** output. [PR1494274](#)

SEE ALSO



---

[What's New | 296](#)

---

[What's Changed | 306](#)

---

[Known Limitations | 315](#)

---

[Open Issues | 317](#)

---

[Documentation Updates | 328](#)

---

[Migration, Upgrade, and Downgrade Instructions | 328](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 20.2R3 documentation for the SRX Series.

### SEE ALSO

---

[What's New | 296](#)

---

[What's Changed | 306](#)

---

[Known Limitations | 315](#)

---

[Open Issues | 317](#)

---

[Resolved Issues | 319](#)

---

[Migration, Upgrade, and Downgrade Instructions | 328](#)

## 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 19.3, 19.4, and 20.1 are EEOL releases. You can upgrade from Junos OS Release 19.3 to Release 19.4 or from Junos OS Release 19.3 to Release 20.1.

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

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

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

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

#### SEE ALSO

[What's New | 296](#)

[What's Changed | 306](#)

[Known Limitations | 315](#)

[Open Issues | 317](#)

[Resolved Issues | 319](#)

[Documentation Updates | 328](#)

## Junos OS Release Notes for vMX

#### IN THIS SECTION

- [What's New | 330](#)
- [What's Changed | 331](#)
- [Known Limitations | 331](#)
- [Open Issues | 331](#)
- [Resolved Issues | 331](#)

- Licensing | 332
- Upgrade Instructions | 332

These release notes accompany Junos OS Release 20.2R3 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

- What's New in Release 20.2R3 | 330
- What's New in Release 20.2R2 | 330

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

### What's New in Release 20.2R3

There are no new features for vMX in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features for vMX in Junos OS Release 20.2R2.

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 331](#)
- [What's Changed in Release 20.2R2 | 331](#)

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

### What's Changed in Release 20.2R3

There are no changes in behavior or syntax for vMX in Junos OS Release 20.2R3.

### What's Changed in Release 20.2R2

There are no changes in behavior or syntax for vMX in Junos OS Release 20.2R2.

## Known Limitations

There are no known behaviors and limitations for vMX in Junos OS Release 20.2R3.

## Open Issues

There are no open issues for vMX in Junos OS Release 20.2R3.

## Resolved Issues

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

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

## Resolved Issues: 20.2R3

There are no resolved issues for vMX in Junos OS Release 20.2R3.

## Resolved Issues: 20.2R2

### *Platform and Infrastructure*

- Configuring the ranges statement for autosensed VLANs might not work on the vMX platforms.  
[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 | 333](#)
- [What's Changed | 334](#)
- [Known Limitations | 334](#)
- [Open Issues | 335](#)
- [Resolved Issues | 335](#)

These release notes accompany Junos OS Release 20.2R3 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

- [What's New in Release 20.2R3 | 333](#)
- [What's New in Release 20.2R2 | 334](#)

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

### What's New in Release 20.2R3

There are no new features for vRR in Junos OS Release 20.2R3.

To learn about common BGP or routing Junos features supported on vRR for Junos OS 20.2R3, see [What's New](#) for MX Series routers.

## What's New in Release 20.2R2

To learn about common BGP or routing Junos features supported on vRR for Junos OS 20.2R2, see [What's New](#) for MX Series routers.

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3](#) | 334
- [What's Changed in Release 20.2R2](#) | 334

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

### What's Changed in Release 20.2R3

There are no changes in behavior or syntax for vRR in Junos OS Release 20.2R3.

To learn more about common BGP or routing changes in behavior or syntax in Junos OS 20.2R3, see [What's Changed](#) for MX Series routers.

### What's Changed in Release 20.2R2

There are no changes in behavior or syntax for vRR in Junos OS Release 20.2R2.

To learn more about common BGP or routing changes in behavior or syntax in Junos OS 20.2R2, see [What's Changed](#) for MX Series routers.

## Known Limitations

Learn about known limitations in this release for vRR.

To learn more about common BGP or routing known limitation in Junos OS 20.2R3, 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.

## Open Issues

There are no known issues for vRR in Junos OS Release 20.2R3.

To learn more about common BGP or routing open issues in Junos OS 20.2R3, 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.

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

### Resolved Issues: 20.2R3

To learn more about common BGP or routing resolved issues in Junos OS 20.2R3, see [Resolved Issues](#) for MX Series routers.

#### CLI

- If `output-queue-priority expedited update-tokens` is configured, rpd might crash upon BGP flapping. [PR1545837](#)
- Six PE device prefixes might not be removed from RIB upon reception of withdrawal from a BGP neighbor when the RIB sharding is enabled. [PR1556271](#)

# Junos OS Release Notes for vSRX

#### IN THIS SECTION

- [What's New | 336](#)
- [What's Changed | 337](#)
- [Known Limitations | 338](#)
- [Open Issues | 338](#)



- Resolved Issues | 339
- Migration, Upgrade, and Downgrade Instructions | 341

These release notes accompany Junos OS Release 20.2R3 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

- What's New in Release 20.2R3 | 336
- What's New in Release 20.2R2 | 336

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

### What's New in Release 20.2R3

There are no new features for vSRX in Junos OS Release 20.2R3.

### What's New in Release 20.2R2

There are no new features for vSRX in Junos OS Release 20.2R2.

## What's Changed

### IN THIS SECTION

- [What's Changed in Release 20.2R3 | 337](#)
- [What's Changed in Release 20.2R2 | 337](#)

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

### What's Changed in Release 20.2R3

There are no changes in behavior or syntax for vSRX in Junos OS Release 20.2R3.

### What's Changed in Release 20.2R2

#### *Platform and Infrastructure*

- **Repetition of WALinuxAgent logs causing file size increase (vSRX 3.0)**—The Azure WALinuxAgent performs the provisioning job for the vSRX instances. When a new vSRX instance is deployed, the continued increasing size of the waagent log file might cause the vSRX to stop.

If the vSRX is still operating, then delete the `/var/log/waagent.log` directly or run the `clear log waagent.log all` command to clear the log file.

Or you can run the `set groups azure-provision system syslog file waagent.log archive size 1m` and `set groups azure-provision system syslog file waagent.log archive files 10` commands to prevent the growing of the waagent logs. These configurations will cause the rotation of log of waagent with the size bigger than 1MB and set a maximum of 10 backups.

See [vSRX with Microsoft Azure](#).

- **vSRX 3.0 instances with AWS Key Management Service (KMS)**—On vSRX 3.0 instances with AWS Key Management Service (KMS), if the MEK is changed, then the keypairs will be re-encrypted using the newly set Master Encryption Key (MEK).

## Known Limitations

### IN THIS SECTION

- [J-Web | 338](#)

Learn about known limitations in Junos OS Release 20.2R3 for vSRX Series. For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

### J-Web

- When a dynamic application is created for an edited policy rule, the list of services is blank when the Services tab is clicked and then the policy grid is autorefreshed. As a workaround, create a dynamic application as the last action while modifying the policy rule and click the Save button to avoid loss of configuration changes made to the policy rule. [PR1460214](#)
- For a spoke device in a hub-and-spoke topology, J-Web shows the VPN topology as Site to Site. [PR1495973](#)

## Open Issues

### IN THIS SECTION

- [Intrusion Detection and Prevention \(IDP\) | 339](#)
- [J-Web | 339](#)
- [Platform and Infrastructure | 339](#)

Learn about open issues in Junos OS Release 20.2R3 for vSRX Series. 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)

- IDP database file format or convention has changed in Junos OS Release 15.1X49 and later releases. So, if the IDP configuration contains some predefined attacks or attack-groups related configurations, then the system will go to amnesiac mode after upgrade. This is due to the failure in IDP configuration commit. [PR1455125](#)

## J-Web

- Configuration of global settings options of IPsec VPN such as TCP encap profile, IPsec power mode, and IKE package installation are not supported from J-Web. [PR1496439](#)

## Platform and Infrastructure

- On vSRX 3.0 on Azure, with Microsoft Azure Hardware Security Module (HSM) enabled, keypair generation fails if the user reuses the certificate ID for creating a new keypair, even if the previous keypair has been deleted. [PR1490558](#)
- When using Juniper vSRX deployment script `deploy-azure-vsrx.sh` to create new vSRX instance, if the same user was defined in both `parameter.json` file and YAML file (using `write_files` module), both passwords will be configured in different configuration groups in the running configuration of vSRX. The password defined in the YAML file will be considered. [PR1491074](#)
- vSRX instances starts to support using cloud feed as source address or destination address in the security policy. Due to the dynamic nature of cloud provisioning, we use warning instead of error when the policy's source address or destination address is not found. [PR1521739](#)

## Resolved Issues

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

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

## Resolved Issues: 20.2R3

### *Intrusion Detection and Prevention (IDP)*

- The flowd or srxpfe process might generate core files during the idpd process commit. [PR1521682](#)

### *Platform and Infrastructure*

- SRX series devices or vSRX instances fail to download dynamic-address feed from Security Director. [PR1442248](#)
- The control link might be broken when there is excessive traffic load on the control link in vSRX cluster deployment. [PR1524243](#)
- The master-password configuration is rejected if master encryption password is not set. [PR1537251](#)
- The srxpfe process might crash when Application Identification Packet-Capture functionality is enabled. [PR1538991](#)
- Configuration integrity mismatch is observed error in vSRX3.0 running on Azure with key-vault integrated. [PR1551419](#)
- High CPU usage on pkid process might be seen when the device is unable to connect to a particular CRL URL. [PR1560374](#)

## Resolved Issues: 20.2R2

### *Intrusion Detection and Prevention (IDP)*

- When adaptive threat profiling is configured within an IDP rule base and logging is enabled, on the vSRX instances the Packet Forwarding Engine process might stop and generate the core file. [PR1532737](#)

### *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*

- On Microsoft Azure deployments, SSH public key authentication is not supported for vSRX 3.0 CLI and portal deployment. [PR1402028](#)
- The vSRX may restart unexpectedly. [PR1479156](#)
- Changes to the configuration command for assigning more vCPUs to the Routing Engine. [PR1505724](#)
- 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](#)
- The flowd or srxpfe process might crash when SSL proxy and AppSecure process traffic simultaneously. [PR1516969](#)

### ***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](#)

### ***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 a IPsec VPN scenario. [PR1517262](#)

## **Migration, Upgrade, and Downgrade Instructions**

### **IN THIS SECTION**

- [Upgrading Software Packages | 342](#)
- [Validating the OVA Image | 348](#)

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 20.2R3 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 20.2R3 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 20.2R3 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-x86-64-20.2R3-2021-02-02.0_RELEASE_20.2R3_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-x86-64-20.2R3-2021-02-02.0_RELEASE_20.2R3_THROTTLE.tgz
no-copy no-validate reboot
Verified junos-vsrx-x86-64-20.2R3-2021-02-02.0_RELEASE_20.2R3_THROTTLE signed by
PackageDevelopmentEc_2020 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING: This package will load JUNOS 20.2R3 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-20.2R3-2021-02-02.0_RELEASE_20.2R3_THROTTLE.tgz

```

```

Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31
junos-srx-mr-vsrx-20.2R3-2021-02-02.0_RELEASE_20.2R3_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31
junos-srx-mr-vsrx-20.2R3-2021-02-02.0_RELEASE_20.2R3_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-20.2R3-2021-02-02.0_RELEASE_20.2R3_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-20.2R3-2021-02-02.0_RELEASE_20.2R3_THROTTLE-linux.tgz
...
upgrade_platform: Input package
/var/tmp/junos-srx-mr-vsrx-20.2R3-2021-02-02.0_RELEASE_20.2R3_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-20.2R3-2021-02-02.0_RELEASE_20.2R3_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-20.2R3-2021-02-02.0_RELEASE_20.2R3_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 20.2R3 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 20.2R3-2021-02-02.0_RELEASE_20.2R3_THROTTLE Kernel 64-bit
JNPR-11.0-20210202.170745_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: 20.2R3-2020-9-10.0_RELEASE_20.2R3_THROTTLE
JUNOS OS Kernel 64-bit [20210202.170745_fbsd-builder_stable_11]
JUNOS OS libs [20210202.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20210202.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20210202.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20210202.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20210202.170745_fbsd-builder_stable_11]
JUNOS py extensions [20201017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20201017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20210202.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20210202.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities
[20201017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20201017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20201017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20201017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package
[20201017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20201017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20201017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20201017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp libs [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srx Data Plane Crypto Support
[20201017.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20201017.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20201017.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20201017.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20210202.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20201017.110007_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](http://prsearch.juniper.net).
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. [apps.juniper.net/hct/home](http://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/](http://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

1 April 2021—Revision 1, Junos OS Release 20.2R3— ACX Series, cSRX, EX Series, JRR Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

25 March 2021—Revision 1, Junos OS Release 20.2R2-S3— MX Series.

22 February 2021—Revision 1, Junos OS Release 20.2R2-S2— MX Series and QFX Series.

13 January 2021—Revision 3, Junos OS Release 20.2R2— ACX Series, cSRX, EX Series, JRR Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

10 December 2020—Revision 2, Junos OS Release 20.2R2— ACX Series, cSRX, EX Series, JRR Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

9 November 2020—Revision 1, Junos OS Release 20.2R2— ACX Series, cSRX, EX Series, JRR Series, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

8 October 2020—Revision 7, Junos OS Release 20.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 September 2020—Revision 6, Junos OS Release 20.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 September 2020—Revision 5, Junos OS Release 20.2R1— ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 August 2020—Revision 1, Junos OS Release 20.2R1-S1— EX Series, MX Series, and QFX Series.

30 July 2020—Revision 4, Junos OS Release 20.2R1- ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 July 2020—Revision 3, Junos OS Release 20.2R1- ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 July 2020—Revision 2, Junos OS Release 20.2R1- ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

30 June 2020—Revision 1, Junos OS Release 20.2R1- ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.