

Release Notes

Published
2022-03-31

Junos[®] OS 20.4R1 Release Notes

SUPPORTED ON

- ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX

SOFTWARE HIGHLIGHTS

- Support for mobility on Junos Multi-Access User Plane (MX204, MX240, MX480, MX960, MX10003)
- Static VXLAN at VLAN or bridge domain level (MX5, MX10, MX40, MX80, MX150, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016 routers and QFX5120-32C, QFX5120-48T, and QFX5120-48Y switches)
- Support for cRPD in SONiC (PTX10008)
- Phone-home client (EX4300-48MP Virtual Chassis)
- RADIUS attributes for dynamic VLAN assignment on colorless ports (EX2300, EX2300-MP, EX3400, EX4300, and EX4300-MP)
- ZTP with DHCPv6 client support (EX3400, EX4300, PTX1000, PTX5000, PTX10002-60C, PTX10008, QFX5100, QFX5200, QFX10002, and QFX10002-60C)
- Support for express segments to establish end-to-end segment routing path (MX Series and PTX Series)
- MAC-VRF with EVPN-VXLAN (MX Series and vMX routers; QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002-60C, QFX10008, and QFX10016 switches)
- Support for tunneling applications in unified policies (NFX Series and SRX Series)
- Support for unidirectional session refreshing (SRX Series)

- Support for captive portal on Wi-Fi Mini-Physical Interface Module (SRX320, SRX340, SRX345, SRX380, and SRX550HM)
- Support for Annex J and G.Fast with specialized SFP (SRX380, SRX300, SRX320, SRX340, and SRX345)
- Security policy support for security inspection on VXLAN tunnels (SRX4100, SRX4200, SRX4600, and vSRX)
- AWS Key Management Service (KMS) Integration support (vSRX 3.0)

IN FOCUS GUIDE

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

Day One+

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

Release Notes: Junos[®] OS Release 20.4R1 for the ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX

31 March 2022

Contents	Introduction 16
	Junos OS Release Notes for ACX Series 16
	What's New 17
	Hardware 18
	High Availability (HA) and Resiliency 21
	Junos Telemetry Interface 21
	Routing Protocols 22
	Timing and Synchronization 22
	What's Changed 23
	General Routing 23
	MPLS 24
	Network Management and Monitoring 24
	Routing Protocols 24
	User Interface and Configuration 24
	Known Limitations 25
	General Routing 25
	Timing and Synchronization 25

Open Issues | 27**Class of Service (CoS) | 27****General Routing | 27****Platform and Infrastructure | 28****VPNs | 28****Resolved Issues | 29****Forwarding and Sampling | 29****General Routing | 29****Interfaces and Chassis | 32****Layer 2 Features | 32****Routing Protocols | 32****Documentation Updates | 32****Migration, Upgrade, and Downgrade Instructions | 33****Upgrade and Downgrade Support Policy for Junos OS Releases | 33****Junos OS Release Notes for cRPD | 34****What's New | 34****Platform and Infrastructure | 35****What's Changed | 35****Junos Telemetry Interface | 36****Known Limitations | 36****Open Issues | 36****Resolved Issues | 36****Routing Policy and Firewall Filters | 37****Junos OS Release Notes for cSRX | 37****What's New | 37****What's Changed | 37****Platform and Infrastructure | 38****Known Limitations | 38****Open Issues | 39****Resolved Issues | 39****Junos OS Release Notes for EX Series | 39****What's New | 40****Authentication, Authorization, and Accounting | 40****EVPN | 40**

Interfaces and Chassis	43
Junos OS XML, API, and Scripting	43
Network Management and Monitoring	44
Routing Protocols	44
Software Installation and Upgrade	45
Subscriber Management and Services	46
What's Changed	46
Test	47
MPLS	49
Network Management and Monitoring	49
Platform and Infrastructure	49
User Interface and Configuration	49
Known Limitations	50
EVPN	51
Platform and Infrastructure	51
Open Issues	51
Infrastructure	52
Juniper Extension Toolkit (JET)	53
Platform and Infrastructure	53
Routing Policy and Firewall Filters	54
Routing Protocols	54
User Interface and Configuration	54
Resolved Issues	55
Authentication and Access Control	55
EVPN	55
Infrastructure	55
Layer 2 Features	56
Network Management and Monitoring	56
Platform and Infrastructure	56
Routing Protocols	57
User Interface and Configuration	57
Virtual Chassis	57
Documentation Updates	58

Migration, Upgrade, and Downgrade Instructions | 58

Upgrade and Downgrade Support Policy for Junos OS Releases | 59

Junos OS Release Notes for JRR Series | 59

What's New | 60

Routing Protocols | 61

What's Changed | 61

Known Limitations | 62

Routing Protocols | 62

Open Issues | 63

Resolved Issues | 63

Resolved Issues: 20.4R1 Release | 63

Documentation Updates | 64

Migration, Upgrade, and Downgrade Instructions | 64

Upgrade and Downgrade Support Policy for Junos OS Releases | 65

Junos OS Release Notes for Juniper Secure Connect | 66

What's New | 66

What's Changed | 66

Known Limitations | 66

Open Issues | 67

Juniper Secure Connect Client | 67

Resolved Issues | 67

Junos OS Release Notes for Junos Fusion for Enterprise | 67

What's New | 68

What's Changed | 68

Known Limitations | 69

Open Issues | 69

Resolved Issues | 70

Resolved Issues: Release 20.4R1 | 70

Documentation Updates | 71

Migration, Upgrade, and Downgrade Instructions | 71

Basic Procedure for Upgrading Junos OS on an Aggregation Device | 71

Upgrading an Aggregation Device with Redundant Routing Engines | 73

Preparing the Switch for Satellite Device Conversion | 74

Converting a Satellite Device to a Standalone Switch | 75

Upgrade and Downgrade Support Policy for Junos OS Releases | 75

Downgrading Junos OS | 76

Junos OS Release Notes for Junos Fusion for Provider Edge | 77

What's New | 77

Hardware | 78

What's Changed | 78

Known Limitations | 79

Open Issues | 79

Resolved Issues | 80

Documentation Updates | 80

Migration, Upgrade, and Downgrade Instructions | 81

Basic Procedure for Upgrading an Aggregation Device | 81

Upgrading an Aggregation Device with Redundant Routing Engines | 84

Preparing the Switch for Satellite Device Conversion | 84

Converting a Satellite Device to a Standalone Device | 86

Upgrading an Aggregation Device | 88

Upgrade and Downgrade Support Policy for Junos OS Releases | 88

Downgrading from Junos OS Release 20.1 | 89

Junos OS Release Notes for MX Series | 89

What's New | 90

Hardware | 91

EVPN | 93

High Availability (HA) and Resiliency | 94

Interfaces and Chassis | 95

Juniper Extension Toolkit (JET) | 95

Junos OS, XML, API, and Scripting | 96

Junos Telemetry Interface | 96

MPLS | 98

Network Management and Monitoring | 99

Routing Policy and Firewall Filters | 100

Routing Protocols | 100

Services Applications | 102

Software Defined Networking | 102

Software Installation and Upgrade | 104

Software Licensing	104
Subscriber Management and Services	104
System Management	105
System Logging	106
What's Changed	106
Class of Service (CoS)	107
EVPN	107
General Routing	108
High Availability (HA) and Resiliency	109
Interfaces and Chassis	109
J-Web	110
MPLS	110
Network Management and Monitoring	111
Platform and Infrastructure	111
Routing Protocols	111
User Interface and Configuration	111
Known Limitations	112
General Routing	112
Interfaces and Chassis	113
MPLS	113
Network Management and Monitoring	113
Open Issues	113
Class of Service (CoS)	114
EVPN	114
Forwarding and Sampling	114
General Routing	115
Infrastructure	117
Interfaces and Chassis	117
Juniper Extension Toolkit (JET)	117
Layer 2 Ethernet Services	118
MPLS	118
Platform and Infrastructure	118
Routing Policy and Firewall Filters	119
Routing Protocols	119

User Interface and Configuration	119
VPNs	119
Resolved Issues	120
EVPN	121
Forwarding and Sampling	122
General Routing	122
Infrastructure	128
Interfaces and Chassis	129
Intrusion Detection and Prevention (IDP)	130
Juniper Extension Toolkit (JET)	130
J-Web	130
Layer 2 Ethernet Services	130
Layer 2 Features	130
MPLS	130
Network Address Translation (NAT)	131
Network Management and Monitoring	131
Platform and Infrastructure	131
Routing Policy and Firewall Filters	132
Routing Protocols	132
Services Applications	134
Subscriber Access Management	134
User Interface and Configuration	134
VPNs	134
Documentation Updates	135
Migration, Upgrade, and Downgrade Instructions	135
Basic Procedure for Upgrading to Release 20.4R1	136
Procedure to Upgrade to FreeBSD 11.x-Based Junos OS	137
Procedure to Upgrade to FreeBSD 6.x-Based Junos OS	139
Upgrade and Downgrade Support Policy for Junos OS Releases	141
Upgrading a Router with Redundant Routing Engines	141
Downgrading from Release 20.4R1	142

Junos OS Release Notes for NFX Series | 142

What's New | 143

- Application Security | 143
- High Availability | 145
- Flow-Based and Packet-Based Processing | 145
- Logical Systems and Tenant Systems | 145
- Routing Protocols | 145
- Security | 146

What's Changed | 146

- Junos OS XML API and Scripting | 147

Known Limitations | 147

- Interfaces | 148

Open Issues | 148

- Interfaces | 149
- Platform and Infrastructure | 149
- Virtual Network Functions (VNFs) | 149

Resolved Issues | 149

- High Availability | 150
- Interfaces | 150
- Platform and Infrastructure | 150

Documentation Updates | 150

Migration, Upgrade, and Downgrade Instructions | 151

- Upgrade and Downgrade Support Policy for Junos OS Releases | 151
- Basic Procedure for Upgrading to Release 20.4 | 152

Junos OS Release Notes for PTX Series | 153

What's New | 154

- Junos OS XML, API, and Scripting | 154
- Junos Telemetry Interface | 155
- MPLS | 157
- Network Management and Monitoring | 158
- Routing Policy and Firewall Filters | 158
- Routing Protocols | 159
- Software Installation and Upgrade | 161

System Logging	161
What's Changed	162
Class of Service (CoS)	162
General Routing	162
MPLS	163
Network Management and Monitoring	163
User Interface and Configuration	163
Known Limitations	164
General Routing	164
Routing Protocols	164
Open Issues	165
General Routing	165
Layer 2 Ethernet Services	167
MPLS	167
Platform and Infrastructure	167
Routing Protocols	167
Resolved Issues	168
General Routing	168
Infrastructure	169
Interfaces and Chassis	169
MPLS	169
Network Management and Monitoring	169
Routing Protocols	170
Documentation Updates	170
Migration, Upgrade, and Downgrade Instructions	171
Basic Procedure for Upgrading to Release 20.4	171
Upgrade and Downgrade Support Policy for Junos OS Releases	174
Upgrading a Router with Redundant Routing Engines	174
Junos OS Release Notes for the QFX Series	175
What's New	175
Hardware	176
Class of Service (CoS)	190
EVPN	191
Flow-Based and Packet-Based Processing	194

High Availability (HA) and Resiliency	197
Interfaces and Chassis	197
IP Tunneling	197
Juniper Extension Toolkit	197
Junos OS XML, API, and Scripting	198
Junos Telemetry Interface	198
Network Management and Monitoring	198
Platform and Infrastructure	200
Routing Policy and Firewall Filters	200
Routing Protocols	201
Software Defined Networking (SDN)	202
Software Installation and Upgrade	203
System Management	204
System Logging	204
What's Changed	205
Class of Service (CoS)	205
General Routing	205
MPLS	206
Network Management and Monitoring	206
User Interface and Configuration	206
Known Limitations	207
General Routing	207
Layer 2 Features	209
Routing Protocols	209
Open Issues	210
EVPN	211
General Routing	211
High Availability (HA) and Resiliency	213
Layer 2 Ethernet Services	213
Layer 2 Features	213
Platform and Infrastructure	213
Routing Policy and Firewall Filters	214
Routing Protocols	214
Virtual Chassis	214

Resolved Issues | 215

Resolved Issues: 20.4R1 Release | 215

Documentation Updates | 219

Migration, Upgrade, and Downgrade Instructions | 220

Upgrading Software on QFX Series Switches | 220

Installing the Software on QFX10002-60C Switches | 223

Installing the Software on QFX10002 Switches | 223

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 | 224

Installing the Software on QFX10008 and QFX10016 Switches | 226

Performing a Unified ISSU | 230

Preparing the Switch for Software Installation | 231

Upgrading the Software Using Unified ISSU | 231

Upgrade and Downgrade Support Policy for Junos OS Releases | 233

Junos OS Release Notes for SRX Series | 234

What's New | 235

Application Layer Gateways (ALGs) | 236

Application Security | 236

ATP Cloud | 237

Authentication and Access Control | 238

Chassis Clustering | 238

Flow-Based and Packet-Based Processing | 240

Interfaces and Chassis | 241

Intrusion Detection and Prevention | 242

Juniper Extension Toolkit (JET) | 243

Junos OS XML and API Scripting | 244

J-Web | 244

Layer 2 Features | 246

Logical Systems and Tenant Systems | 246

Multinode High Availability | 246

Network Management and Monitoring | 247

Securing GTP and SCTP Traffic | 248

Security | 249

Unified Threat Management (UTM) | 250

VPNs	250
What's Changed	251
Class of Service (CoS)	252
Flow-Based and Packet-Based Processing	252
Intrusion Detection and Prevention (IDP)	253
Interfaces and Chassis	253
J-Web	253
Network Address Translation (NAT)	254
Network Management and Monitoring	254
Platform and Infrastructure	254
Securing GTP and SCTP Traffic	254
User Interface and Configuration	255
VPNs	255
Known Limitations	256
Class of Service (CoS)	257
Flow-Based and Packet-Based Processing	257
J-Web	257
VPNs	258
Open Issues	258
Flow-Based Packet-Based Processing	259
Interfaces and Chassis	259
J-Web	259
Protocols	259
Routing Policy and Firewall Filters	259
VPNs	260
Resolved Issues	260
Application Layer Gateways (ALGs)	261
Flow-Based and Packet-Based Processing	261
Interfaces and Chassis	262
Intrusion Detection and Prevention (IDP)	262
J-Web	262
Layer 2 Ethernet Services	263
Network Address Translation (NAT)	263
Platform and Infrastructure	263

Routing Policy and Firewall Filters	263
Routing Protocols	264
Subscriber Access Management	264
Unified Threat Management (UTM)	264
VPNs	264
Documentation Updates	264
Migration, Upgrade, and Downgrade Instructions	265
Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases	265
Junos OS Release Notes for vMX	266
What's New	267
EVPN	267
Juniper Extension Toolkit (JET)	267
Junos OS XML ,API, and Scripting	268
Network Management and Monitoring	268
Routing Protocols	269
What's Changed	269
Licensing	270
Known Limitations	270
Open Issues	270
Resolved Issues	270
Interfaces and Chassis	270
Network Management and Monitoring	271
Licensing	271
Upgrade Instructions	271
Junos OS Release Notes for vRR	272
What's New	272
Routing Protocols	273
What's Changed	273
Known Limitations	273
Open Issues	273
Resolved Issues	274

Junos OS Release Notes for vSRX | 274

What's New | 274

- ATP Cloud | 275
- Flow-Based Packet-Based Processing | 275
- High Availability | 276
- Juniper Extension Toolkit (JET) | 276
- Junos OS XML ,API, and Scripting | 277
- Network Management and Monitoring | 277
- Platform and Infrastructure | 278
- Routing Protocols | 278
- VPNs | 278

What's Changed | 278

- Platform and Infrastructure | 279

Known Limitations | 279

Open Issues | 279

- J-Web | 280
- Platform and Infrastructure | 280

Resolved Issues | 280

- Application Security | 280
- Chassis Clustering | 281
- CLI | 281
- Flow-Based and Packet-Based Processing | 281
- Install and Upgrade | 281
- Interfaces and Chassis | 281
- Intrusion Detection and Prevention (IDP) | 281
- Platform and Infrastructure | 282
- Routing Policy and Firewall Filters | 282
- User Access and Authentication | 282
- VPNs | 282

Migration, Upgrade, and Downgrade Instructions | 282

- Upgrading Software Packages | 284
- Validating the OVA Image | 289

Upgrading Using ISSU | 289

Licensing | 290

Compliance Advisor | 290

Finding More Information | 290

Documentation Feedback | 291

Requesting Technical Support | 291

Self-Help Online Tools and Resources | 292

Creating a Service Request with JTAC | 292

Revision History | 293

Introduction

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

These release notes accompany Junos OS Release 20.4R1 for the ACX Series, Containerized Routing Protocol Process (cRPD), cSRX Container Firewall (cSRX), EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise, Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, virtual MX Series router (vMX), Virtual Route Reflector (vRR), and vSRX Virtual Firewall (vSRX). They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

- [In Focus guide](#)—We have a document called In Focus that provides details on the most important features for the release in one place. We hope this document will quickly get you to the latest information about Junos OS features. Let us know if you find this information useful by sending an e-mail to techpubs-comments@juniper.net.
- **Important Information:**
 - [Upgrading Using ISSU on page 289](#)
 - [Licensing on page 290](#)
 - [Compliance Advisor on page 290](#)
 - [Finding More Information on page 290](#)
 - [Documentation Feedback on page 291](#)
 - [Requesting Technical Support on page 291](#)

Junos OS Release Notes for ACX Series

IN THIS SECTION

- [What's New | 17](#)
- [What's Changed | 23](#)
- [Known Limitations | 25](#)
- [Open Issues | 27](#)
- [Resolved Issues | 29](#)

- Documentation Updates | 32
- Migration, Upgrade, and Downgrade Instructions | 33

These release notes accompany Junos OS Release 20.4R1 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Hardware | 18
- High Availability (HA) and Resiliency | 21
- Junos Telemetry Interface | 21
- Routing Protocols | 22
- Timing and Synchronization | 22

This section describes the new features or enhancements to existing features in Junos OS Release 20.4R1 for the ACX Series.

Hardware

- We've added the following features to the ACX5448 in Junos OS Release 20.4R1.

Table 1: Features Supported by the ACX5448 Routers

Feature	Description
Authentication, Authorization and Accounting	<ul style="list-style-type: none"> • Support for 802.1X authentication on Layer 3 interfaces. 802.1X is an IEEE standard for port-based network access control that authenticates users connected to a LAN port. [See 802.1X Authentication.]
Automation	<ul style="list-style-type: none"> • Support for either WAN interfaces or management interfaces to automatically download and install the appropriate software and the configuration file on your device during the ZTP bootstrap process. [See Zero Touch Provisioning.]
Class of service (CoS)	<ul style="list-style-type: none"> • Support for up to three levels of hierarchical scheduling (physical interfaces, logical interfaces, and queues). Configurable buffer support is also added. By default, all interfaces on the ACX5448 use port-based scheduling (eight queues per physical port). To enable hierarchical scheduling, set the hierarchical-scheduler statement at the [edit interfaces interface-name] hierarchy level. [See Hierarchical Class of Service in ACX Series Routers.]
Ethernet OAM	<ul style="list-style-type: none"> • Support for Ethernet OAM CFM. You can now synchronize local-interface status between two connected devices with remote interface up/down trigger with OAM CFM. CFM provides end-to-end signals even if the two devices are not directly connected. [See Introduction to OAM Connectivity Fault Management (CFM).]
EVPN	<ul style="list-style-type: none"> • Support for EVPNs and Interfaces. In EVPN-MPLS and MC-LAG environments, the configuration of anycast gateways on ACX5448 routers that are multihomed in all-active mode is supported. [See Anycast Gateways.]
Layer 2 features	<ul style="list-style-type: none"> • Support for pseudowire redundancy in MC-LAG. ACX5448 routers support pseudowire redundant Layer 2 circuits in MC-LAG routers. VPLS is not supported. [See Understanding Pseudowire Redundancy Mobile Backhaul Scenarios.]
Layer 3 features	<ul style="list-style-type: none"> • Support for Layer 3 VPN in MC-LAG chassis. ACX5448 routers support Layer 3 VPN in VRRP over IRB interfaces in MC-LAG routers. Layer 3 routing and Layer 3 VPN are not directly supported on the MC-LAG interfaces. [See Understanding VRRP and Understanding Layer 3 VPNs.]

Table 1: Features Supported by the ACX5448 Routers (*continued*)

Feature	Description
Network Security	<ul style="list-style-type: none"> Support for control plane DDoS protection, which is enabled by default on ACX5448 routers for many Layer 2 and Layer 3 protocols. Control Plane DDoS protection uses firewall filters and policers to discard or rate-limit control plane traffic at the Routing Engine level, which prevents malicious traffic from interfering with device operations. You can disable this feature or change the default policer parameters for supported protocol groups. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview]
Software installation and upgrade	<ul style="list-style-type: none"> Support for the ACX5448-M-LT, a top-of-rack router that supports only Junos Limited image. The Junos Limited image does not have data-plane encryption and is intended only for countries in the Eurasian Customs Union because these countries have import restrictions on software containing data-plane encryption. Unlike the JunosWorldwide image, the Junos Limited image supports control plane encryption through Secure Shell (SSH) and Secure Sockets Layer (SSL), thus allowing secure management of the system. [See ACX5448 System Overview.]
Timing and synchronization	<ul style="list-style-type: none"> Support for Precision Time Protocol (PTP) G.8275.2 enhanced profile with PTP over IPv4 and IPv6 unicast traffic. [See Understanding the PTP G.8275.2 Enhanced Profile (Telecom Profile).]

- **Support for SFP-1GE-LH-ET transceivers (ACX1100 and ACX2100)**—Starting in Junos OS Release 20.4R1, the ACX1100 and ACX2100 Universal Metro Routers support the SFP-1GE-LH-ET transceivers. [See the [Hardware Compatibility Tool \(HCT\)](#) for details.]
- **Support for SFP-GE80KT14R15 and SFP-GE80KT15R14 transceivers (ACX5448, ACX5448-D, and ACX5448-M)**—Starting in Junos OS Release 20.4R1, the ACX5448, ACX5448-D, and ACX5448-M Universal Metro Routers support the SFP-GE80KT14R15 and SFP-GE80KT15R14 transceivers. [See the [Hardware Compatibility Tool \(HCT\)](#) for details.]
- **Support for SFPP-10GE-DWDM-IT transceivers (ACX5448, ACX5448-D, and ACX5448-M)**—Starting in Junos OS Release 20.4R1, the ACX5448, ACX5448-D, and ACX5448-M Universal Metro Routers support the SFPP-10GE-DWDM-IT transceivers. [See the [Hardware Compatibility Tool \(HCT\)](#) for details.]

High Availability (HA) and Resiliency

- **NSR support for IS-IS with SR (ACX Series, MX Series)**—Starting in Junos OS Release 20.4R1, ACX Series devices support NSR for IS-IS with segment routing (SR). To use NSR, you must first enable GRES on your device.

[See [Nonstop Active Routing Concepts](#)]

Junos Telemetry Interface

- **JTI support for persistent active gRPC sessions between collector and server during an SSL certificate update (ACX Series, MX Series, and PTX Series)**—Junos OS Release 20.4R1 supports persistent active remote procedure call (gRPC) sessions between the collector (client) and server during an SSL certificate update.

For secure channel authentication, the TLS protocol is used to maintain a secure channel between the collector and the server. TLS uses the server certificate and the client certificate to authenticate each other and send encrypted messages over the network. When an SSL certificate is updated, existing gRPC sessions are abruptly terminated, forcing the collector to initiate a new gRPC connection and subscribe to sensors again.

To avoid this problem, you can enable persistent active gRPC sessions by configuring **hot-reloading** at the `[edit system services extension-service request-response grpc ssl]` hierarchy level. After you enable this feature, gRPC sessions will remain active even when authentication certificates are updated.

After the certificate is updated, any new gRPC session will use the updated certificate.

[See [gRPC Services for Junos Telemetry Interface](#) and [ssl](#).]

- **Juniper Resiliency Interface for exception reporting and null route detection (ACX Series, PTX Series, and MX Series)**—Starting in Junos OS Release 20.4R1, you can use Juniper Resiliency Interface to detect and reduce Mean Time to Repair (MTTR) first-order network issues. Juniper Resiliency Interface uses a push model for data reporting from the entities in the system which encounter packet drops. This automates the workflow for detecting, reporting, and mitigating adverse exceptions.

To collect kernel routing table and routing protocol process exceptions, configure the **set system resiliency exceptions** statement at the `[edit]` hierarchy level to specify exception reporting based on kernel exceptions, and routing exceptions.

You can display exceptions from a remote collector by means of remote procedure call (gRPC) services or gRPC network management interface (gNMI) services. Display on-box exceptions by accessing the `/var/log` file or the database at `/var/db/ResiliencyExceptions.db`. No Junos operational mode commands display these exceptions.

Routing Protocols

- **Support for multiple single-hop EBGp sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGp session. You can now enable single-hop EBGp sessions on different links over multiple directly connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGp Sessions on Different Links Using the Same Link-Local Address \(IPv6\)](#).]

Timing and Synchronization

- **Support for PTP G.8275.2 profile (ACX710)**—Starting in Junos OS Release 20.4R1, we support the Precision Time Protocol (PTP) G.8275.2 profile with node type T-BC-P (BC).

You can use the [edit protocols ptp profile-type g.8275.2] hierarchy level to configure the G.8275.2 profile.

[See [Understanding the Time Management Administration Guide](#) and [profile-type](#).]

SEE ALSO

What's Changed 23
Known Limitations 25
Open Issues 27
Resolved Issues 29
Documentation Updates 32
Migration, Upgrade, and Downgrade Instructions 33

What's Changed

IN THIS SECTION

- General Routing | 23
- MPLS | 24
- Network Management and Monitoring | 24
- Routing Protocols | 24
- User Interface and Configuration | 24

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.4R1 for the ACX Series routers.

General Routing

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

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

MPLS

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

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to "deviate not-supported" nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.

Routing Protocols

- **Inet6 is disabled in VT interface (ACX5448)**—Starting in this release, the `inet6` statement at the `edit interfaces vt-interface-number unit unit-number family` hierarchy level is disabled.

User Interface and Configuration

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

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

SEE ALSO

What's New 17
Known Limitations 25
Open Issues 27
Resolved Issues 29
Documentation Updates 32

Known Limitations

IN THIS SECTION

- General Routing | 25
- Timing and Synchronization | 25

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

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

General Routing

- On the ACX710 router, Servo moves to the **Holdover-in/Holdover-out/Acq** state from the **Phase-aligned** state with impairment. [PR1550367](#)
- On the ACX710 router, PTP with **Vlan-id-range** does not work for specific VLANs. [PR1550482](#)
- On the ACX710 router, the holdover error **HOLDOVER OUT OF SPEC** does not reset during the Servo state change. [PR1556798](#)

Timing and Synchronization

- On the ACX5448 router, the two-way time error and CTE for 1 PPS does not meet the class A metrics. [PR1535434](#)
- On the ACX5448-M router, the 1 PPS CTE does not meet the class A performance in 1-Gigabits interface. [PR1542744](#)
- On the ACX5448 router, due to BRCM KBP issue route lookup might fail. [PR1533557](#)
- On the ACX5448 router, ping stops working even though the ARP entry is present during continuous script executions. [PR1533513](#)
- On the ACX710 router, T1 or T4 cTE should be tuned closer to two-way CTE. [PR1527347](#)
- On the ACX710 router, huge offset is observed initially with ACQ and holdover inspec and outspec conditions. [PR1534470](#)

- On the ACX710 router, the incremental PTP FPGA upgrades do not bundle along with the regular image upgrades. [PR1540799](#)
- On the ACX710 router, changing the PTP profile type from g.8275.1 to g.8275.2 requires the Packet Forwarding Engine to reboot and the clksyncd process to restart. As a workaround, you must reboot the Packet Forwarding Engine and restart the clocking process before you change the profile. [PR1546614](#)
- On the ACX710 router, the Servo transition is incorrect after chassis restart. [PR1550270](#)
- On the ACX710 router, the delay-asymmetry compensation update does not work at CLI with the G.8275.2 profile. [PR1550441](#)
- On the ACX710 router, the PTP Servo status shows holdover during transition between virtual port and PTP. [PR1510880](#)
- On the ACX710 router, if the client clock candidate is configured with a virtual port, the clock class is on T-BC. [PR1520204](#)
- On the ACX710 router, the SyncE to 1PPS transient test results do not meet G.8273.2 SyncE to 1PPS transient metric. [PR1522796](#)
- On the ACX710 router, the clock parameters are incorrect in certain scenarios when the Servo is in the **FREERUN** state. [PR1548192](#)
- On the ACX710 router, the PTP Servo takes longer time to lock after the clksyncd process restarts. [PR1549952](#)
- On the ACX710 router, the **show ptp global-information** command does not display correct Clock Class or ESMC QL details when the Servo goes to the **Holdover-in** state. [PR1553213](#)
- On the ACX710 router, the Servo transition is incorrect during the T-GM switchover scenario. [PR1553439](#)

SEE ALSO

[What's New | 17](#)

[What's Changed | 23](#)

[Open Issues | 27](#)

[Resolved Issues | 29](#)

[Documentation Updates | 32](#)

[Migration, Upgrade, and Downgrade Instructions | 33](#)

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 27](#)
- [General Routing | 27](#)
- [Platform and Infrastructure | 28](#)
- [VPNs | 28](#)

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

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

Class of Service (CoS)

- Unexpected behavior of Class of Service is observed with the wildcard classifier. [PR1559516](#)

General Routing

- On the ACX5448 router, latency is observed for the host-generated ICMP traffic. [PR1380145](#)
- Tx power cannot be configured using the + sign. [PR1383980](#)
- On the ACX710 router, alarm is not raised when booting the system with recovery snapshot. [PR1517221](#)
- On the ACX5448 router, the BGPV6LU traffic drop is observed when the node is deployed in ingress. [PR1538819](#)
- On the ACX500-I router, the **show services session count** does not work as expected. [PR1520305](#)
- 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 ACX710 router, the global configuration of **IPv4-dscp** naming convention must be corrected as per the stream level **dscp**, which is more meaningful for both the IPv6 and IPv4 services. [PR1557262](#)
- 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](#)

- On the ACX5448 router, the following syslog message is reported every 30 seconds;
ACX_DFW_CFG_FAILED: ACX Error (dfw):dnx_dfw_dyn_entry_counter_get: Entry is invalid. [PR1562323](#)
- On the ACX5448 router, the transit DHCPv4 and DHCPv6 packets drop in a Layer 2 domain. [PR1517420](#)
- On the ACX5448 router, the ISSU upgrade fails due to the Packet Forwarding Engine restart issue. [PR1554915](#)
- On the ACX5048 router, all the OAM sessions are not established. [PR1561751](#)
- 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 ACX5448 device as TWAMP server delays the start session acknowledgment by 10 seconds. [PR1556829](#)
- On the ACX2100 device, **laser-output-power** is seen after the interface is disabled and rebooted. [PR1560501](#)
- Inline BFD stays down with IS-IS or Static clients. [PR1561590](#)

Platform and Infrastructure

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

VPNs

- On the ACX5448 router, the MC-AE Layer 2 circuit states are not updated instantly and for some time after disabling the core interface on the MC-LAG active node, double hit in traffic is observed. [PR1543408](#)

SEE ALSO

[What's New | 17](#)

[What's Changed | 23](#)

[Known Limitations | 25](#)

[Resolved Issues | 29](#)

[Documentation Updates | 32](#)

[Migration, Upgrade, and Downgrade Instructions | 33](#)

Resolved Issues

IN THIS SECTION

- Forwarding and Sampling | 29
- General Routing | 29
- Interfaces and Chassis | 32
- Layer 2 Features | 32
- Routing Protocols | 32

This section lists the issues fixed in Junos OS Release 20.4R1 for the ACX Series.

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

Forwarding and Sampling

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

General Routing

- The **gether-options** command is enabled again under the interface hierarchy. [PR1430009](#)
- Repeated powering-off or powering-on of the device, the SMBUS transactions timeout occurs. [PR1463745](#)
- On the ACX5048 router, the egress queue statistics do not work for the aggregated Ethernet interfaces. [PR1472467](#)
- On the ACX5048 router, traffic loss is observed during the unified ISSU upgrade. [PR1483959](#)
- The following syslog error message is observed: **ACX_DFW_CFG_FAILED**. [PR1490940](#)
- On the ACX5048 and ACX5096 routers, the LACP control packets might be dropped due to high CPU utilization. [PR1493518](#)
- On the ACX710 router, high convergence is observed with the EVPN-ELAN service in a scaled scenario during FRR switchover. [PR1497251](#)
- On the ACX5448 router, the EXP rewrite for the Layer 3 VPN sends all traffic with incorrect EXP. [PR1500928](#)
- The following error message is observed during MPLS route add, change, and delete operation: **mpls_extra NULL**. [PR1502385](#)

- The ACX1100, ACX2100, ACX2200, ACX2000, and ACX4000 routers might stop forwarding transit and control traffic. [PR1508534](#)
- On the ACX710 router, the Packet Forwarding Engine might crash and the fpc process might remain down. [PR1509402](#)
- The loopback filter cannot take more than 2 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](#)
- Tagged traffic matching the vlan-list configuration in the vlan-circuit cross-connect logical interface gets dropped in the ingress interface. [PR1519568](#)
- The **Incompatible Media** alarm is not raised when the Synchronous Ethernet source is configured over the copper SFP. [PR1519615](#)
- On the ACX710 router, the alarm port configuration is not cleared after deleting the alarm-port. [PR1520326](#)
- PTP to 1PPS noise transfer test fails for frequency 1.985 Hz. [PR1522666](#)
- The **show class-of-service interface** command does not show the classifier information. [PR1522941](#)
- Interface does not come up with the auto-negotiation setting between the ACX1100 router and the other ACX Series routers, MX Series routers and QFX Series switches as the other end. [PR1523418](#)
- 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](#)
- On the ACX5448 and ACX710 routers, the **vlan-id-list** statement might not work as expected. [PR1527085](#)
- The FEC field is not displayed when the interface is down. [PR1530755](#)
- The **show class-of-service routing-instance** does not show the configured classifier. [PR1531413](#)
- Memory leak in Local OutLif in VPLS/CCC topology is observed. [PR1532995](#)
- The clksyncd process generates core file on Junos OS Release 20.3R1.3 image. [PR1537107](#)
- The rpd process generates core file at `l2ckt_vc_adv_recv, l2ckt_adv_rt_flash (taskptr=0x4363b80, rtt=0x4418100, rtl=< optimized out>, data=< optimized out>, opcode=< optimized out>)` at `../../../../../../../../src/junos/usr/sbin/rpd/l2vpn/l2ckt.c:7982`. [PR1537546](#)
- The **Management Ethernet link down** alarm is observed while verifying the system alarms in the Virtual Chassis setup. [PR1538674](#)

- On the ACX5448 router, unexpected behavior of the **show chassis network-services** command is observed. [PR1538869](#)
- The following error message is observed while deleting the remote stream 0 0 0 0 0 along with feb core file at 0x00ae6484 in `bcmdnx_queue_assert (queue=0xc599b60)` at `../../../../src/pfe/common/drivers/bcmdnx/bcmdnx_sdk_ukern_layer.c: Err] clksync_mimic_delete_clock_entry Unexpected error`. [PR1539953](#)
- The announcement or synchronization interval rate range is not as expected. [PR1542516](#)
- Synchronization Ethernet goes in the **Holdover** state and comes back to the **Locked** state when the PTP configuration is deleted. [PR1546681](#)
- The ACX5448 router as transit for the BGP labeled unicast drops traffic. [PR1547713](#)
- Multicast traffic is stopped when HQoS with multicast configurations are applied. [PR1551248](#)
- With the **no-local-switching** command, traffic between the local and remote CE devices are affected. [PR1527231](#)
- On the ACX710 router, the T-BC-P switch-over performance fails beyond the standard mask and servo moving to multiple **Holdover-in** state, **Acquiring** state, **Holdover-in** state, **Holdover-out** state, and **Acquiring** state. [PR1556087](#)
- Running SNMP MIB walk and executing the **show interfaces** command might cause the picd process to crash. [PR1533766](#)
- On the ACX5448 router, you cannot downgrade to Junos OS Release 18.4 code-base. [PR1556377](#)
- BIND does not sufficiently limit the number of fetches while processing referrals. [PR1512212](#)
- The clksyncd process generates core file during the stability test with traffic and scale. [PR1518253](#)
- The fxpc process generates core file during EEPROM read when SFP is removed. [PR1518480](#)
- On the ACX5448 routers, multicast traffic loop over ICL might be observed. [PR1521113](#)
- On the ACX710 router, PIR/CIR HQoS behavior is inconsistent. [PR1525789](#)
- Error messages are displayed while attaching tcp on physical interfaces. [PR1527541](#)
- The l2cpd memory leak might be observed with the aggregated Ethernet interface flap. [PR1527853](#)
- Upon classifying the Layer 3 packets, DSCP is not preserved and is lost at the egress due to the limitations of a chipset. [PR1535876](#)
- Other than IPv4 and IPV6, other IPs should not be forwarded. Only IP header with version 4 and 6 can pass through. [PR1550748](#)
- Profile switch between G.8275.1 and G.8275.2 works as expected. [PR1533263](#)

Interfaces and Chassis

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

Layer 2 Features

- On the ACX5448 routers, the VPLS traffic statistics are not displayed when the **show vpls statistics** command is executed. [PR1506981](#)
- The rpd might crash on the new primary Routing Engine after GRES in the VPLS or Layer 2 circuit scenario. [PR1507772](#)

Routing Protocols

- The rpd process might report 100 percent CPU usage with the BGP route damping enabled. [PR1514635](#)
- On the ACX5448 routers, the **family inet6** configuration under the vt- interface is disabled. [PR1514595](#)

SEE ALSO

What's New 17
What's Changed 23
Known Limitations 25
Open Issues 27
Documentation Updates 32
Migration, Upgrade, and Downgrade Instructions 33

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for ACX Series routers.

SEE ALSO

What's New 17
What's Changed 23
Open Issues 27

[Known Limitations | 25](#)

[Resolved Issues | 29](#)

[Migration, Upgrade, and Downgrade Instructions | 33](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 33](#)

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

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

SEE ALSO

What's New 17
What's Changed 23
Known Limitations 25
Open Issues 27
Resolved Issues 29
Documentation Updates 32

Junos OS Release Notes for cRPD

IN THIS SECTION

- What's New | 34
- What's Changed | 35
- Known Limitations | 36
- Open Issues | 36
- Resolved Issues | 36

These release notes accompany Junos OS Release 20.4R1 for the containerized routing protocol process (cRPD) container. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Platform and Infrastructure | 35

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

Platform and Infrastructure

- **Support for eventd (cRPD)**—Starting in Junos OS Release 20.4R1, we support only external event policies. You can enable these policies in container RPD. In cRPD, eventd and rsyslogd run as two independent processes. The eventd process provides event interface to processes such as rpd/auditd/mgd and supports automated event policy execution.

Use the **set event-options policy *policy name* events [*events*] then** command to enable an event policy and **restart event-processing** to restart event processing.

By default, Python 3.x support is enabled along with existing on-box Python/SLAX functions in cRPD environment.

Use the **[edit system scripts language python3]** command to enable and to support python event automation.

[See [event-options](#), [events](#) and [event-policy](#).]

- **Support for Configuring cRPD through SONiC (PTX10008)**—Juniper Networks' PTX10008 router supports configuring cRPD in SONiC through the **config_db.json** configuration utility. The **config_db.json** utility is a local redis database (redis-db). You need to do a **config save** and **config load** for the configurations to take effect in cRPD.
- **Support for cRPD in SONiC (PTX10008)**—cRPD routing stack is supported on PTX10008 router running SONiC.

What's Changed

IN THIS SECTION

- [Junos Telemetry Interface](#) | 36

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

Junos Telemetry Interface

- **cRPD supports the Junos Telemetry Interface (JTI) over TLS similar to Junos OS (cRPD)**—cRPD supports local (server-side) certificate validation for gRPC and JTI similar to Junos OS. cRPD doesn't support bidirectional authentication for gRPC and JTI. See [Configuring gRPC for the Junos Telemetry Interface](#) and [Importing SSL Certificates for Junos XML Protocol Support](#).

Known Limitations

There are no known behavior for cRPD in Junos OS Release 20.4R1.

Open Issues

There are no open issues for cRPD in Junos OS Release 20.4R1.

Resolved Issues

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

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

Routing Policy and Firewall Filters

- The show route forwarding-table or show route instance operational commands output is incomplete.
[PR1545415](#)

Junos OS Release Notes for cSRX

IN THIS SECTION

- [What's New | 37](#)
- [What's Changed | 37](#)
- [Known Limitations | 38](#)
- [Open Issues | 39](#)
- [Resolved Issues | 39](#)

These release notes accompany Junos OS Release 20.4R1 for the cSRX Container Firewall, a containerized version of the SRX Series Services Gateway. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features in Junos OS Release 20.4R1 for cSRX.

What's Changed

IN THIS SECTION

- [Platform and Infrastructure | 38](#)

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

Platform and Infrastructure

- Downloading of Signature Pack You can download the signature pack from the Signature Pack directly when the cSRX doesn't have pre-installed signature pack.
 - Configure proxy server so that IP address of proxy server is reachable from cSRX.
 - Run the following command to enter the configurational mode from CLI.

```
root@host> configure [edit]
root@host#
```
 - Configure proxy server profile on cSRX using IP address and port of proxy server.

```
root@host#set services proxy profile appid_sigpack_proxy protocol http host 4.0.0.1
root@host#set services proxy profile appid_sigpack_proxy protocol http port 3128
```
 - Attach the profile to AppID and IDP.

```
root@host#set services application-identification download proxy-profile appid_sigpack_proxy
root@host#set security idp security-package proxy-profile appid_sigpack_proxy
```
 - Commit the configuration.

```
root@host#commit and-quit
commit complete
```
 - Download IDP and APPID sigpack through proxy server.

```
root@host>request services application-identification download
root@host>request security idp security-package download
```
 - To verify if download is going through proxy server:
 Verify the logs in proxy server.

```
[root@srxdpi-lnx39 squid]# cat /var/log/squid/access.log 1593697174.470 1168 4.0.0.254
TCP_TUNNEL/200 5994 CONNECT signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
1593697175.704 1225 4.0.0.254 TCP_TUNNEL/200 11125 CONNECT signatures.juniper.net:443 -
HIER_DIRECT/66.129.242.156 - 1593697176.950 1232 4.0.0.254 TCP_TUNNEL/200 5978 CONNECT
signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 - 1593697178.195 1236 4.0.0.254
TCP_TUNNEL/200 11188 CONNECT signatures.juniper.net:443 - HIER_DIRECT/66.129.242.156 -
1593697198.337 1243 4.0.0.254 TCP_TUNNEL/200 6125 CONNECT signatures.juniper.net:443 -
HIER_DIRECT/66.129.242.156 -
```

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

Known Limitations

There are no known behavior for cSRX in Junos OS Release 20.4R1.

Open Issues

There are no open issues for cSRX in Junos OS Release 20.4R1.

Resolved Issues

There are no resolved issues for cSRX in Junos OS Release 20.4R1.

Junos OS Release Notes for EX Series

IN THIS SECTION

- What's New | 40
- What's Changed | 46
- Known Limitations | 50
- Open Issues | 51
- Resolved Issues | 55
- Documentation Updates | 58
- Migration, Upgrade, and Downgrade Instructions | 58

These release notes accompany Junos OS Release 20.4R1 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Authentication, Authorization, and Accounting | 40
- EVPN | 40
- Interfaces and Chassis | 43
- Junos OS XML, API, and Scripting | 43
- Network Management and Monitoring | 44
- Routing Protocols | 44
- Software Installation and Upgrade | 45
- Subscriber Management and Services | 46

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

NOTE: The following EX Series switches are supported in Release 20.4R1: EX2300, EX3400, EX4300, EX4600, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

Authentication, Authorization, and Accounting

- **RADIUS attributes for dynamic VLAN assignment on colorless ports (EX2300, EX2300-MP, EX3400, EX4300, and EX4300-MP)**—We now support IETF-defined RADIUS attributes that provide VLAN assignments and also indicate whether frames on the VLAN are in tagged or untagged format. This enables the network access control server to dynamically assign VLANs on colorless ports. The VLAN assignments, which are based on device profiling, can be made on either access ports or trunk ports.

[See [Dynamic VLAN Assignment on Colorless Ports](#).]

EVPN

- **MAC limit, MAC move limit, and persistent MAC learning with EVPN-VXLAN (EX4300-48MP)**—We support the following Layer 2 port security features in an EVPN-VXLAN overlay network:
 - **MAC limit**—You can limit the number of MAC addresses learned by network (local) interfaces.

NOTE: We don't support MAC limits on virtual tunnel endpoint (VTEP) interfaces.

- **MAC move limit**—You can limit the number of times a MAC address is moved to a different interface within 1 second. To configure this feature, you apply a limit to a VLAN. In an EVPN-VXLAN network, a VLAN's members can include network (local) and VTEP interfaces. We support the following MAC move use cases and actions:
 - **MAC moves between network interfaces**—By default, the configured action is applied on the interface to which the MAC address is last moved. If you configured action priority on the interfaces, the action is applied on the interface with the lesser priority.
 - **MAC moves between network and VTEP interfaces and vice-versa**—The action is applied on the network interface.

NOTE: We don't support MAC moves between the following:

- VTEP interfaces.
 - A VTEP interface and a network interface on which persistent MAC learning and static MAC addresses are configured.
- **Persistent MAC learning (sticky MAC)**—You can enable network interfaces to retain dynamically learned MAC addresses when the switch is restarted or when an interface goes down and comes back up again.

NOTE: We don't support persistent MAC learning on VTEP interfaces.

[See [Understanding MAC Limiting and MAC Move Limiting](#) and [Understanding and Using Persistent MAC Learning](#).]

- **MC-LAG emulation in an EVPN deployment (EX Series, MX Series, and vMX)**—Starting in Junos OS Release 20.4R1, you can emulate the function of an MC-LAG in active-standby mode in an EVPN configuration without having to configure an ICCP or ICL interface. In a standard EVPN configuration, logical interfaces configured on an aggregated Ethernet interface can have different designated forwarder election roles. To emulate an MC-LAG configuration, the designated forwarder (DF) takes on the role of the aggregated Ethernet interface. The provider edge (PE) that is the non-DF will send LACP out-of-sync packets to the CE. This causes LACP to go down on the CE device, and the CE device does not use the links connected to the non-DF for sending traffic. If the connection between a CE and a DF PE fails, the PE is re-elected as a DF. If the connection between a CE and a non-DF PE fails, the current DF PE is not changed.

To enable this functionality, configure the **lcp-oos-on-ndf** statement at the **[edit interfaces interface name esi df-election-granularity per-esi]** hierarchy.

- **Support for IGMP snooping and selective multicast forwarding (EX4300-MP)**—Starting in Junos OS Release 20.4R1, the EX4300-MP switch supports IGMP snooping and selective multicast forwarding in an EVPN-VXLAN centrally-routed bridging overlay network with all-active multihoming. Selective multicast Ethernet (SMET) forwarding is part of IGMP snooping. IGMP snooping and SMET forwarding reduce the volume of multicast traffic in a broadcast domain by forwarding multicast traffic only to interfaces that have IGMP listeners. SMET forwarding sends multicast packets to the leaf devices in the core that have expressed an interest in that multicast group. SMET forwarding is supported only in intra-VLAN replication. This feature supports EVPN Type 7 (IGMP Join Synch Route) and EVPN Type 8 (IGMP Leave Synch Routes). To configure IGMP snooping, include the **igmp-snooping proxy** configuration statement at the **[edit routing-instances routing-instance-name protocols]** hierarchy level.

[See [Overview of Multicast Forwarding with IGMP Snooping in an EVPN-VXLAN Environment](#) and [Overview of Selective Multicast Forwarding](#).]

- **Support for assisted replication (EX4300MP)**—Starting in Junos OS Release 20.4R1, the EX4300-MP switch supports assisted replication in an EVPN-VXLAN centrally-routed bridging overlay network with all-active multihoming. Assisted replication (AR) optimizes multicast traffic flow by offloading traffic replication to devices that can more efficiently handle replication and forwarding. You can configure the EX4300-MP only as an AR-leaf device. You can further optimize multicast traffic by configuring AR with IGMP snooping. To configure the EX4300-MP as an AR leaf, include the **assisted-replication leaf** statement at the **[edit routing-instances routing-instance-name protocols evpn]** or **[edit protocols evpn]** hierarchy level.

[See [Assisted Replication Multicast Optimization in EVPN Networks](#)

- **Support for sFlow in an EVPN-VXLAN network (EX4300-MP)**—Starting in Junos OS Release 20.4R1, sFlow monitoring is supported on EX4300-MP switches in an EVPN-VXLAN network. sFlow monitoring provides visibility into your EVPN VXLAN network by sampling VXLAN-encapsulated traffic at the ingress and egress interfaces. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously. You must enable sFlow monitoring on each interface individually. Configure sFlow monitoring at the **[edit protocols sflow]** hierarchy level. Use the **show sflow collector** command to display the collector statistics and the **clear sflow collector** command to delete the collector statistics.

[See [Overview of sFlow Technology](#).]

- **Layer 3 gateway in an EVPN-MPLS environment (EX9200 with EX9200-SF3 switch fabric module and EX9200-15C line card)**—Starting in Junos OS Release 20.4R1, an EX9200 switch with an EX9200-SF3 switch fabric module and an EX9200-15C line card can act as a default Layer 3 gateway for an EVPN instance (EVI) that can span a set of devices. In this role, the EX9200 switch 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 VLAN 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.](#)]
- 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.](#)]

Interfaces and Chassis

- **10GBASE-T SFP+ transceiver for EX4600-40F**—Starting in Junos OS Release 20.4R1, EX4600-40F switches support the 10GBASE-T SFP+ transceiver (JNP-SFPP-10GE-T), capable of working at speeds of 10 Gbps, 1Gbps, and 100Mbps, and also auto-negotiation. You can use the existing show commands such as the **show interfaces media** command to view the details of the transceivers.

[See [speed\(Ethernet\).](#)]

Junos OS XML, API, and Scripting

- **Support for Certificate Authority Chain Profile (EX2300, EX3400, EX4300, MX240, MX480, MX960, PTX-5000, VMX, vSRX and QFX5200)**—Starting in Junos OS Release 20.4R1, you can configure intermediate Certificate Authority (CA) chain profile certificate and perform https REST API request using mutual and server authentications.

To configure intermediate ca-chain certificate, configure **ca-chain ca-chain** statement at the **[edit system services rest https]** hierarchy level.

- **Start time option for interval-based internal events that trigger event policies (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.4R1, when you create an interval-based internal event for triggering event policies, you can specify the start date and time for the initial event. To specify a start time, configure the **start-time** option along with the **time-interval** option at the **[edit event-options generate-event]** hierarchy level.

[See [Generating Internal Events to Trigger Event Policies.](#)]

Network Management and Monitoring

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The **<get-configuration>** operation supports the **compare="configuration-revision"** and **configuration-revision** attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

Routing Protocols

- **BGP Prefix-Independent Convergence (PIC) Edge for MPLS VPNs (EX9200)**—You can now install a Layer 3 VPN route in the forwarding table as an alternate path, enabling fast failover when a provider edge (PE) router fails or you lose connectivity to a PE router. This already installed path is used until global convergence through the IGP is resolved.

To enable BGP PIC Edge in an MPLS VPN, include the **protect-core** statement at the **[edit routing-instances routing-instance-name routing-options]** hierarchy level. Both IS-IS LDP and OSPF LDP are supported. When BGP PIC Edge is enabled, the **show route extensive** command now displays the weight assigned to the indirect hop.

[See [Configuring BGP PIC Edge for MPLS Layer 3 VPNs](#).]

- **Support for multiple single-hop EBGp sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGp session. You can now enable single-hop EBGp sessions on different links over multiple directly connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGp Sessions on Different Links Using the Same Link-Local Address \(IPv6\)](#).]

Software Installation and Upgrade

- **Phone-home client (EX4600, EX4650, EX9200, QFX5110, QFX5200, QFX5210, QFX5120-32C, and QFX5120-48Y)**—Starting with Junos OS Release 20.4R1, you can use either the legacy DHCP-options-based ZTP or the phone-home client (PHC) to provision software for the switch. When the switch boots up, if there are DHCP options that have been received from the DHCP server for ZTP, ZTP resumes. If DHCP options are not present, PHC is attempted. PHC enables the switch to securely obtain bootstrapping data, such as a configuration or software image, with no user intervention other than having to physically connect the switch to the network. When the switch first boots up, PHC connects to a redirect server, which redirects to a phone home server to obtain the configuration or software image.

To initiate either DHCP-options-based ZTP or PHC, the switch must be in a factory-default state, or you can issue the **request system zeroize** command.

[See [Understanding the Phone-Home Client](#)

- **ZTP with DHCPv6 client support (EX3400, EX4300, PTX1000, PTX5000, PTX10002-60C, PTX10008, QFX5100, QFX5200, QFX10002, and QFX10002-60C)**—Starting in Junos OS Release 20.4R1, zero touch 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 one of the DHCPv4 bindings fails, the device continues to check for bindings until provisioning is successful. However, if there are no DHCPv4 bindings, the device checks for DHCPv6 bindings and follows 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: ZTP supports only HTTP and HTTPS transport protocols.

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

- **Phone-home client (EX4300-48MP Virtual Chassis)**—Starting in Junos OS Release 20.4R1, the phone-home client (PHC) can securely provision a Virtual Chassis consisting of all EX4300-48MP member switches without requiring user interaction. If the switches all have the factory-default configuration, you just need to:
 - Connect the switches using the Virtual Chassis ports.
 - Connect any network port or the management port to the network.

- Power on the Virtual Chassis.

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

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

Subscriber Management and Services

- **Control plane DDoS protection against DDoS attacks (EX9200 with MPC10E)**—Starting in Junos OS Release 20.4R1, control plane distributed denial of service (DDoS) protection is enabled by default on EX9200 switches with MPC10E line cards. To prevent malicious traffic from interfering with device operations, this feature uses firewall filters and policers to discard or rate-limit control plane traffic. You can disable this feature at different levels or change the default policer parameters for many protocol groups and individual packet types in the supported protocol groups.

[See [Control Plane Distributed Denial-of-Service \(DDoS\) Protection Overview.](#)]

SEE ALSO

What's Changed	 46
Known Limitations	 50
Open Issues	 51
Resolved Issues	 55
Documentation Updates	 58
Migration, Upgrade, and Downgrade Instructions	 58

What's Changed

IN THIS SECTION

- [Test](#) | [47](#)
- [MPLS](#) | [49](#)
- [Network Management and Monitoring](#) | [49](#)
- [Platform and Infrastructure](#) | [49](#)
- [User Interface and Configuration](#) | [49](#)

Learn about what changed in this release for EX Series Switches in Junos OS Release 20.4R1.

Test

Uncategorized

- **SSH session connection attempt limits and connection limits (PTX10008 and PTX10003)**—We have introduced the **connection-limit** and **rate-limit** options at the **set system services ssh** hierarchy levels. The default connection limit value is 75 connections, and the default rate limit value is 3 connections per second. Junos OS measures the rate limit value per minute but Junos OS Evolved measures the rate limit value per second.
- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

See [arp](#).

- **Change in license bandwidth command on vMX virtual routers**—Starting in Junos OS Release 21.1R1, to use the available license bandwidth, explicitly set the license bandwidth use the **set chassis license bandwidth** command.

See [Configuring Licenses on vMX Virtual Routers](#).

MPLS

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

See [show mpls lsp](#).

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to "deviate not-supported" nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.
- **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](#).

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](#).

MPLS

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

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to deviate not-supported nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.

Platform and Infrastructure

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

User Interface and Configuration

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

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

SEE ALSO

What's New	 40
Known Limitations	 50
Open Issues	 51
Resolved Issues	 55
Documentation Updates	 58
Migration, Upgrade, and Downgrade Instructions	 58

Known Limitations

IN THIS SECTION

- [EVPN](#) | [51](#)
- [Platform and Infrastructure](#) | [51](#)

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

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

Platform and Infrastructure

- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. Device can be recovered using power-cycle of the device. [PR1385970](#)
- 10G Channels shows false up even when peer end is configured with different speed. The LED on the box also shows green. [PR1530061](#)
- In a qinq configuration, xSTP should not be enabled on interface having ifls with vlan-id-list configured. If xSTP is enabled on such interface, it will only run on ifl whose vlan-id range includes native-vlan-id configured, and all other ifls of this interface will in discarding state. So, user should not enable xSTP on these kind of interfaces. Sample configuration which is not allowed: set interfaces ge-0/0/1 flexible-vlan-tagging set interfaces ge-0/0/1 native-vlan-id 3000 set interfaces ge-0/0/1 encapsulation extended-vlan-bridge set interfaces ge-0/0/1 unit 2000 vlan-id-list 1-200 set interfaces ge-0/0/1 unit 2000 input-vlan-map push set interfaces ge-0/0/1 unit 2000 output-vlan-map pop set vlans csvlan1 interface ge-0/0/1.2000 set protocols mstp interface ge-0/0/1. [PR1532992](#)

SEE ALSO

What's New 40
What's Changed 46
Open Issues 51
Resolved Issues 55
Documentation Updates 58
Migration, Upgrade, and Downgrade Instructions 58

Open Issues

IN THIS SECTION

- [Infrastructure | 52](#)
- [Juniper Extension Toolkit \(JET\) | 53](#)

- Platform and Infrastructure | 53
- Routing Policy and Firewall Filters | 54
- Routing Protocols | 54
- User Interface and Configuration | 54

Learn about open issues in Junos OS Release 20.4R1 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.

Infrastructure

- On EX Series legacy Switches, fsck is run with '-C' option, which skips the file system corruption check if the partition has been marked clean during the boot 'nand-media' check. Due to this, there have been multiple instances where the partition has had file system issues even when cleanly shut down. This change is to enforce fsck during the boot cycle to strengthen the file system check during boot time. Fixed in releases: 12.3R12-S7, 14.1X53-D46, 15.1R6 HOW TO RECOVER:* The switch will repair the corruption during the boot cycle when the file system check (fsck) is run.* In the rare instance that the file system check (fsck) is completed, and there are continued file system corruptions, then the next step is to do an 'install -format'. This will format the file system and all file system corruptions will be removed, along with the previous logs and configuration. [PR1191072](#)
- 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 HSRPv2 (Hot Standby Router Protocol for IPv6) 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 number of firewall filters on some interfaces, the FPC might crash and generate core files. [PR1434927](#)
- PROTOCOLS:SWITCHING: AI: Unable to Verify jais-7.0R3-THIN.0.tgz in EX4600 box due to space issue. [PR1548668](#)
- On EX3400 Virtual Chassis, traffic destined to IRB interface would be dropped after mac-persistence-timer was expired. [PR1557229](#)

Juniper Extension Toolkit (JET)

- gRPC stack uses IPV4 mapped IPV6 address internally, so that gRPC server can work with pure/mapped IPV4/IPV6 addresses. However, a recent change in kernel IPv4/v6 handling causes a problem when a management IP is configured. Workaround: Changing address to 0.0.0.0 solves the issue set system services extension-service request-response grpc clear-text address 0.0.0.0. [PR1559064](#)

Platform and Infrastructure

- On EX, OCX or QFX based platforms using Broadcom chipset, with SFP+ implemented, interface on the platforms might be in active status when TX or RX connector is removed. When this issue happens, traffic could be dropped. [PR1495564](#)
- Do not renumber the Virtual Chassis in non consecutive fashion , for SNMP POE MIB walk to work correctly. [PR1503985](#)
- 35 seconds delay is added in reboot time from Junos OS Release 20.2R1 release compared to Release 19.4R2. [PR1514364](#)
- The **request chassis fpc slot <slot_num> restart** command is unsupported in EX series platforms, so avoid using that command. [PR1536997](#)
- OSPF and OSPF3 adjacency uptime is more than expected after NSSU upgrade and Outage is higher than the expected. [PR1551925](#)
- Traffic drop is seen after I2 gres switchover with Layer 2 forwarding database. [PR1561344](#)
- Limited images are not supported for EX92XX on this release. [PR1561741](#)
- Client authentication is failing after performing graceful switchover. [PR1563431](#)
- On certain Junos platforms with Dual-REs (platforms capable of installing Junos packages with name format as "junos*install"), BGP replication may fail to start under GRES/NSR setup after a crash on backup Routing Engine. NSR starts un-replicating the socket since backup Routing Engine is no longer present. Massive unreplicated request leads to memory buffer getting full with multiple BGP sessions (e.g., 20 BGP peers). Hence BGP unreplicated request returned with an error. Besides, the kernel is left with stale data. It does not allow the JSR (Juniper Socket Replication, BGP in this case) when backup RE comes up due to the stale data. BGP-NSR (Nonstop Routing) is broke under the conditions. Traffic outage will be observed after performing GRES. [PR1552603](#)

Routing Policy and Firewall Filters

- On all Junos OS platforms with "set policy-options rtf-prefix-list" configured, if you upgrade to a specific version, the device might fail to validate its configuration, which eventually causes rpd to crash unexpectedly due to a software fault. [PR1538172](#)

Routing Protocols

- When I2 and I3 ifls are configured on the same ifd and vport scale is enabled on QFX 5110 and QFX 5120 and the I2 ifl is part of a vxlan, then SVP is derived from source_trunk_map table. In this case, the packet will not match with the SOURCE_FIELDS in my_station_tcam table due to which the entry is not getting hit. OSPF unicast pkts will be dropped due to this and it will be stuck in ExStart State. [PR1519244](#)
- On Trio based Virtual Chassis (VC) platform, when there are multicast tunneled packets being received, which come into the Virtual Chassis Ports (VCP) and then pop out of the tunnel, if the VCP ports and the interfaces where multicast packets enter/leave the router are located on the same Packet Forwarding Engine (PFE), it might fail in sending multicast traffic to downstream receiver due to this issue. [PR1555518](#)

User Interface and Configuration

- In Junos OS 20.4R1 release, if your switch is not connected to the Internet, then J-Web UI cannot download and install the J-Web application package automatically. [PR1563588](#)

SEE ALSO

What's New	40
What's Changed	46
Known Limitations	50
Resolved Issues	55
Documentation Updates	58
Migration, Upgrade, and Downgrade Instructions	58

Resolved Issues

IN THIS SECTION

- [Authentication and Access Control | 55](#)
- [EVPN | 55](#)
- [Infrastructure | 55](#)
- [Layer 2 Features | 56](#)
- [Network Management and Monitoring | 56](#)
- [Platform and Infrastructure | 56](#)
- [Routing Protocols | 57](#)
- [User Interface and Configuration | 57](#)
- [Virtual Chassis | 57](#)

This section lists the issues fixed in Junos OS Release 20.4R1 for EX Series switches.

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

Authentication and Access Control

- The 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](#)

Infrastructure

- qmon-sw sensor is not supported in EX3400. [PR1506710](#)
- The IP communication between directly connected interfaces on EX4600 would fail. [PR1515689](#)
- The VC system might get hanged after committing the VSTP configurations. [PR1520351](#)
- OID ifOutDiscards reports zero and sometimes shows valid value. [PR1522561](#)
- Firewall policer with discard action might fail on EX4300. [PR1532670](#)

- Errors might be seen when dumping vmcore on EX2300 and EX3400 switches. [PR1537696](#)
- The LLDP neighborship with the VoIP phones can't be established. [PR1538482](#)

Layer 2 Features

- The dcpfe/FPC might crash due to the memory leak during the vlan add/delete operation. [PR1505239](#)
- 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 member in Virtual Chassis after the MAC flaps. [PR1521324](#)

Network Management and Monitoring

- EX4300: SNMP OID 1.3.6.1.2.1.25.3.3.1.2.0 (hrProcessorLoad) always returns 0 irrespective of the real CPU utilization. [PR1508364](#)

Platform and Infrastructure

- IPv6 neighbor solicitation packets might be dropped in a transit device. [PR1493212](#)
- DHCP Binding is not happening after Graceful switchover. [PR1515234](#)
- LLDP adjacency might fail for non-AE interfaces on EX4300 platform. [PR1538401](#)
- uRPF in the Strict mode does not work. [PR1417546](#)
- Virtual Chassis split after network topology changed. [PR1427075](#)
- IRB MAC will not be programmed in hardware when MAC persistence timer expires. [PR1484440](#)
- Authentication session might be terminated if PEAP request is retransmitted by authenticator. [PR1494712](#)
- In some cases, if we have an OSPF session on the IRB over LAG interface with 40-Gigabit Ethernet port as member, the session gets stuck in restart. [PR1498903](#)
- On the 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](#)
- 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](#)

- Last commit line in configuration is updated after the configuration backup has been done. [PR1513499](#)
- The 100M SFP-FX is not supported on satellite device in a Junos Fusion setup. [PR1514146](#)
- ARP learning issue might be seen on EX4300-MP platform when configuring Layer 3 gateway interfaces. [PR1514729](#)
- "dot1x" memory leak is seen. [PR1515972](#)
- The dcpfe (PFE) process might crash due to memory leak. [PR1517030](#)
- MPPE-Send/Recv-key attribute is not extracted correctly by dot1xd. [PR1522469](#)
- "Drops" and "Dropped packets" counters in the output by "show interface extensive" are double counting. [PR1525373](#)
- EX4300-48MP device might go out of service during a software upgrade operation. [PR1526493](#)
- PoE messages "poe_get_dev_class: Failed to get PD class info" seen on EX2300. [PR1536408](#)
- EX3400, EX2300 : Upgrade failure do to lack of available storage. [PR1539293](#)
- Slaac-Snoopd child process core is observed upon multiple switchovers on Routing Engine. [PR1543181](#)
- EX9200 SF3 Fabric OIR Issues with Junos 23.1R1.8. [PR1555727](#)

Routing Protocols

- The rpd process might report 100 percent CPU usage with the BGP route damping enabled. [PR1514635](#)
- Packet loss might be observed while verifying traffic from access to core network for IPv4 and IPv6 interfaces. [PR1520059](#)
- OSPFv3 adjacency should not be established when IPsec authentication is enabled. [PR1525870](#)

User Interface and Configuration

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

Virtual Chassis

- On the EX4650 device, the following error message is observed during booting: kldload: an error occurred while loading the module. [PR1527170](#)

SEE ALSO

What's Changed	46
Known Limitations	50
Open Issues	51
Documentation Updates	58
Migration, Upgrade, and Downgrade Instructions	58

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for EX Series switches.

SEE ALSO

What's New	40
What's Changed	46
Known Limitations	50
Open Issues	51
Resolved Issues	55
Migration, Upgrade, and Downgrade Instructions	58

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases](#) | [59](#)

This section contains the upgrade and downgrade support policy for Junos OS for EX Series switches. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network. For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

SEE ALSO

What's New 40
What's Changed 46
Known Limitations 50
Open Issues 51
Resolved Issues 55
Documentation Updates 58

Junos OS Release Notes for JRR Series

IN THIS SECTION

- What's New | 60
- What's Changed | 61
- Known Limitations | 62
- Open Issues | 63

- Resolved Issues | 63
- Documentation Updates | 64
- Migration, Upgrade, and Downgrade Instructions | 64

These release notes accompany Junos OS Release 20.4R1 for the JRR Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Routing Protocols | 61


Learn about new features introduced in Junos OS Release 20.4R1 for JRR Series Route Reflectors.

Routing Protocols

- **Support for BGP Sharding (JRR200)**—Starting in Junos OS Release 20.4R1, we support BGP sharding. BGP sharding splits a BGP RIB into several sub RIBs and each sub RIB handles a subset of BGP routes. Each sub RIB is served by a separate RPD thread to achieve parallel processing. This results in reduced convergence time and faster performance. BGP sharding is disabled by default.

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

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

**NOTE:** BGP sharding is supported for IPv4, IPv6, L3VPN and BGP-LU. All the other RIBs are processed without sharding.

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

SEE ALSO

What's Changed 61
Known Limitations 62
Open Issues 63
Resolved Issues 63
Documentation Updates 64
Migration, Upgrade, and Downgrade Instructions 64

What's Changed

There are no changes in behavior and syntax in Junos OS Release 20.4R1 for JRR Series Route Reflectors.

SEE ALSO

What's New 60
Known Limitations 62
Open Issues 63
Resolved Issues 63
Documentation Updates 64
Migration, Upgrade, and Downgrade Instructions 64

Known Limitations

IN THIS SECTION

- [Routing Protocols | 62](#)

Learn about known limitations in this release for JRR200 Route Reflectors.

Routing Protocols

- These features are not supported in Junos OS 20.4R1 release for BGP Sharding:
 - routing-options validations with rib sharding
 - inet4/6 unicast rib-group along with rib sharding
 - outbound route-filter with bgp sharding.

SEE ALSO

What's New 60
What's Changed 61
Open Issues 63
Resolved Issues 63
Documentation Updates 64
Migration, Upgrade, and Downgrade Instructions 64

Open Issues

There are no open issues in Junos OS 20.4R1 Release for JRR Series Route Reflectors.

SEE ALSO

[What's New | 60](#)[What's Changed | 61](#)[Known Limitations | 62](#)[Resolved Issues | 63](#)[Documentation Updates | 64](#)[Migration, Upgrade, and Downgrade Instructions | 64](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.4R1 Release | 63](#)

This section lists the issues fixed in Junos OS Release 20.4R1 for JRR 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: 20.4R1 Release

General Routing

- On the JRR200 routers, the firewall filter with non-zero TTL value might cause a commit error. [PR1531034](#)
- tcp_timer_keep logs flood on JRR200. [PR1533168](#)
- Optics info of physical interfaces is not available for JRR200 on Junos OS. [PR1537261](#)
- The CLI "request system power-off" and "request system halt" commands do not work as expected on JRR200. [PR1534795](#)

SEE ALSO

What's New 60
What's Changed 61
Known Limitations 62
Open Issues 63
Documentation Updates 64
Migration, Upgrade, and Downgrade Instructions 64

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for JRR200 Route Reflectors.

SEE ALSO

What's New 60
What's Changed 61
Known Limitations 62
Open Issues 63
Resolved Issues 63
Migration, Upgrade, and Downgrade Instructions 64

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

SEE ALSO

What's New	 	60
What's Changed	 	61
Known Limitations	 	62
Open Issues	 	63
Resolved Issues	 	63
Documentation Updates	 	64

Junos OS Release Notes for Juniper Secure Connect

IN THIS SECTION

- What's New | 66
- What's Changed | 66
- Known Limitations | 66
- Open Issues | 67
- Resolved Issues | 67

These release notes accompany Junos OS Release 20.4R1 for Juniper Secure Connect. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features in Junos OS Release 20.4R1 for Juniper Secure Connect.

What's Changed

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 20.4R1.

Known Limitations

There are no known behavior or limitation for Juniper Secure Connect in Junos OS Release 20.4R1.

Open Issues

IN THIS SECTION

- [Juniper Secure Connect Client | 67](#)

Learn about open issues in this release for Juniper Secure Connect.

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

Juniper Secure Connect Client

- IKE configure mode payload is not pushing secondary DNS and secondary WINS attributes to Xauth module with IKEv1. Hence client is not getting assigned with secondary DNS and secondary WINS with IKEv1. [PR1558831](#)

Resolved Issues

There are no resolved issues for Juniper Secure Connect in Junos OS Release 20.4R1.

Junos OS Release Notes for Junos Fusion for Enterprise

IN THIS SECTION

- [What's New | 68](#)
- [What's Changed | 68](#)
- [Known Limitations | 69](#)
- [Open Issues | 69](#)
- [Resolved Issues | 70](#)

- Documentation Updates | 71
- Migration, Upgrade, and Downgrade Instructions | 71

These release notes accompany Junos OS Release 20.4R1 for the Junos fusion for enterprise. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

There are no new features or enhancements to existing features in Junos OS Release 20.4R1 for Junos fusion for enterprise.

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

SEE ALSO

What's Changed 68
Known Limitations 69
Open Issues 69
Resolved Issues 70
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

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 20.4R1 for Junos fusion for enterprise.

SEE ALSO

What's New 68
Known Limitations 69
Open Issues 69
Resolved Issues 70
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.4R1 for Junos fusion for enterprise.

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

SEE ALSO

What's New 68
What's Changed 68
Open Issues 69
Resolved Issues 70
Documentation Updates 71
Migration, Upgrade, and Downgrade Instructions 71

Open Issues

There are no known issues in hardware and software in Junos OS Release for 20.4R1 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	 68
What's Changed	 68
Known Limitations	 69
Resolved Issues	 70
Documentation Updates	 71
Migration, Upgrade, and Downgrade Instructions	 71

Resolved Issues

IN THIS SECTION

- [Resolved Issues: Release 20.4R1](#) | [70](#)

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 20.4R1

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

SEE ALSO

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

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 for documentation for Junos fusion for enterprise.

SEE ALSO

[What's New | 68](#)

[What's Changed | 68](#)

[Known Limitations | 69](#)

[Open Issues | 69](#)

[Resolved Issues | 70](#)

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

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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

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

Basic Procedure for Upgrading Junos OS on an Aggregation Device

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

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

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

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

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

To download and install Junos OS:

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

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

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

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

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

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

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Switch

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Downgrading Junos OS

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

NOTE: You cannot downgrade more than three releases.

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

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

SEE ALSO

[What's New | 68](#)

[What's Changed | 68](#)

[Known Limitations | 69](#)

[Open Issues | 69](#)

[Resolved Issues | 70](#)

[Documentation Updates | 71](#)

Junos OS Release Notes for Junos Fusion for Provider Edge

IN THIS SECTION

- [What's New | 77](#)
- [What's Changed | 78](#)
- [Known Limitations | 79](#)
- [Open Issues | 79](#)
- [Resolved Issues | 80](#)
- [Documentation Updates | 80](#)
- [Migration, Upgrade, and Downgrade Instructions | 81](#)

These release notes accompany Junos OS Release 20.4R1 for Junos fusion for provider edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [Hardware | 78](#)

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

Hardware

- **Support for QFX5110 as a satellite device in a Junos fusion for provider edge environment on a GNF (MX480, MX960, MX2010, and MX2020)**—With Junos node slicing, you can create guest network functions (GNFs), which are 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 20.4R1, you can configure QFX5110 switches as satellite devices in a Junos fusion for provider edge environment on a GNF.

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

SEE ALSO

What's Changed	 78
Known Limitations	 79
Open Issues	 79
Resolved Issues	 80
Documentation Updates	 80
Migration, Upgrade, and Downgrade Instructions	 81

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	 77
Known Limitations	 79
Open Issues	 79
Resolved Issues	 80
Documentation Updates	 80
Migration, Upgrade, and Downgrade Instructions	 81

Known Limitations

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 20.4R1 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	 	77
What's Changed	 	78
Open Issues	 	79
Resolved Issues	 	80
Documentation Updates	 	80
Migration, Upgrade, and Downgrade Instructions	 	81

Open Issues

There are no open issues in the Junos OS Release 20.4R1 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	 	77
What's Changed	 	78
Known Limitations	 	79
Resolved Issues	 	80
Documentation Updates	 	80
Migration, Upgrade, and Downgrade Instructions	 	81

Resolved Issues

There are no fixed issues in the Junos OS Release 20.4R1 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	77
What's Changed	78
Known Limitations	79
Open Issues	79
Documentation Updates	80
Migration, Upgrade, and Downgrade Instructions	81

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for Junos fusion for provider edge.

SEE ALSO

What's New	77
What's Changed	78
Known Limitations	79
Open Issues	79
Resolved Issues	80
Migration, Upgrade, and Downgrade Instructions	81

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading an Aggregation Device | 81
- Upgrading an Aggregation Device with Redundant Routing Engines | 84
- Preparing the Switch for Satellite Device Conversion | 84
- Converting a Satellite Device to a Standalone Device | 86
- Upgrading an Aggregation Device | 88
- Upgrade and Downgrade Support Policy for Junos OS Releases | 88
- Downgrading from Junos OS Release 20.1 | 89

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

Basic Procedure for Upgrading an Aggregation Device

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

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

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

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

The download and installation process for Junos OS Release 20.4R1 is different from that for earlier Junos OS releases.

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

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

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

- For 64-bit software:

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

```
user@host> request system software add validate reboot  
source/jinstall64-20.4R1.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.4R1.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

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

```
user@host> request system software add validate reboot
source/jinstall64-20.4R1.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-20.4R1.SPIN-export-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 20.4R1 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

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

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

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

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

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

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

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

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

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

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

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

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

9. Commit the configuration.

To commit the configuration to both Routing Engines:

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

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

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

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

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

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

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

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 20.4R1, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

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


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

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

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

Downgrading from Junos OS Release 20.1

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

**NOTE:** You cannot downgrade more than three releases.

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

SEE ALSO

What's New	 77
What's Changed	 78
Known Limitations	 79
Open Issues	 79
Resolved Issues	 80
Documentation Updates	 80

Junos OS Release Notes for MX Series

IN THIS SECTION

- [What's New](#) | 90
- [What's Changed](#) | 106
- [Known Limitations](#) | 112
- [Open Issues](#) | 113
- [Resolved Issues](#) | 120
- [Documentation Updates](#) | 135
- [Migration, Upgrade, and Downgrade Instructions](#) | 135

These release notes accompany Junos OS Release 20.4R1 for the MX Series 5G Universal Routing Platforms. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Hardware | 91
- EVPN | 93
- High Availability (HA) and Resiliency | 94
- Interfaces and Chassis | 95
- Juniper Extension Toolkit (JET) | 95
- Junos OS, XML, API, and Scripting | 96
- Junos Telemetry Interface | 96
- MPLS | 98
- Network Management and Monitoring | 99
- Routing Policy and Firewall Filters | 100
- Routing Protocols | 100
- Services Applications | 102
- Software Defined Networking | 102
- Software Installation and Upgrade | 104
- Software Licensing | 104
- Subscriber Management and Services | 104
- System Management | 105
- System Logging | 106

This section describes the new features and enhancements to existing features in Junos OS Release 20.4R1 for the MX Series routers.

Hardware

- We've added the following features to the MX Series routers in Junos OS Release 20.4R1.

Table 2: Features Supported by MPC10E and MPC11E Line Cards on MX Series Routers

Feature	Description
EVPN	<ul style="list-style-type: none"> • Support for configuring an Ethernet VPN Ethernet Tree (E-Tree) service on MX240, MX480, and MX960 routers using MPC10E-15C-MRATE line cards. [See EVPN-ETREE Overview.] • Support for configuring an EVPN point-to-multipoint (P2MP) label switch path (LSP) as a provider tunnel on a bud router. The bud router functions both as an egress router and a transit router. [See Configuring Bud Node Support.] • Support for configuring and signalling a P2MP LSP for the EVPN Inclusive Provider Tunnel for BUM traffic. [See Understanding P2MPs LSP for the EVPN Inclusive Provider Tunnel.]
Interfaces and chassis	<ul style="list-style-type: none"> • Support for configuring VLAN rewrite operations on CCC interfaces. [See Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview and Stacking and Rewriting Gigabit Ethernet VLAN Tags.] • Support for 100GE AOC optics on MPC10E-15C-MRATE and MPC10E-10C-MRATE (with SCBE3-MX) in the MX240, MX480, and MX960 routers. [See Hardware Compatibility Tool.] • Support for 4X100G FR transceivers and the channelization option on the 400G-DR4 transceiver on MPC10E-15C-MRATE and MPC10E-10C-MRATE (with SCBE3) in the MX240, MX480, and MX960 routers. [See Hardware Compatibility Tool.] • Support for configuring dynamic learning of the source and destination MAC addresses on aggregated Ethernet interfaces on the MPC10E-15C-MRATE, MPC10E-10C-MRATE, and MX2K-MPC11E line cards. [See MAC Address Accounting for Dynamically Learned Addresses.] • Support for monitoring link degradation of the 25GbE interfaces and 400GbE interfaces on the MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line cards. [See Link Degrade Monitoring Overview.] • Support for Layer 2 address learning process (ALD). [See Understanding Layer 2 Learning and Forwarding.] • Support for a bandwidth of 500 Gbps per Packet Forwarding Engine with four fabric planes on MPC10E-10C-MRATE and MPC10E-15C-MRATE (with the Packet Forwarding Engine 2 powered off) line cards. [See MPC10E-10C-MRATE and MPC10E-15C-MRATE.]

Table 2: Features Supported by MPC10E and MPC11E Line Cards on MX Series Routers (*continued*)

Feature	Description
General routing	<ul style="list-style-type: none"> Support for configuring the TCP maximum segment size (MSS). [See Configure TCP Options.] Support for configuring the GRE key to identify the traffic flows in a GRE tunnel on the MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards. [See dynamic-tunnel-gre-key.]
Layer 2 features	<ul style="list-style-type: none"> Support for packet mirroring with Layer 2 headers for Layer 3 forwarded traffic. [See Firewall Filter Nonterminating Actions.] Support for Layer2 Ethernet services over GRE tunnel interfaces. [See Configuring Layer 2 Ethernet Services over GRE Tunnel Interfaces.] Support for Spanning-Tree Protocol (STP), Rapid Spanning-Tree Protocol (RSTP), Multiple Spanning-Tree Protocol (MSTP), and VLAN Spanning-Tree Protocol (VSTP). [See Configuring STP Protocol.] Support for the base bridging feature commands. <p>NOTE: You can configure propagate option under the mac-flush command.</p> <p>[See clear bridge mac-table, global-mac-move, global-no-mac-learning, mac-flush, global-no-control-mac-aging, and global-no-hw-mac-learning .]</p>
Multicast	<ul style="list-style-type: none"> Support for redundant virtual tunnels (RVTs) and fast re-route (FRR) for both active/backup and active/active redundancy models (MX240, MX480, MX960, MX2010, and MX2020). RVT interfaces are used in Multicast Layer 3 VPNs (MVPN) to facilitate virtual routing and forwarding (VRF) table lookup based on MPLS labels and to provide resiliency. [See Resiliency in Multicast L3 VPNs with Redundant Virtual Tunnels.] Support for verifying the global table multicast (GTM) with IPv6 and Type-7 on MPC10 and MPC11 line cards. [See Multicast Overview.]
Network management and monitoring	<ul style="list-style-type: none"> Support for configuring ITU-T Y.1731 standard-compliant Ethernet synthetic loss measurement (ETH-SLM) and Ethernet delay measurement (ETH- DM) capabilities on MPC10E-10C-MRATE, MPC10E-15C-MRATE, and MX2K-MPC11E line cards. [See ITU-T Y.1731 Ethernet Service OAM Overview.]

Table 2: Features Supported by MPC10E and MPC11E Line Cards on MX Series Routers (*continued*)

Feature	Description
Services Applications	<ul style="list-style-type: none"> Support for inline monitoring services to provide the flexibility to monitor different streams of traffic at different sampling rates on the same interface. [See Inline Monitoring Services Configuration.] Support for Aggregated Multiservices Interfaces (AMS) on the MPC10E-10C-MRATE, MPC10E-15C-MRATE, 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 or MS-MIC only. [See Understanding Aggregated Multiservices Interfaces.]

- **Support for QSFP-100G-FR, QSFP-100G-DR, and QSFP-100G-LR transceivers (MX2010 and MX2020 with MX2K-MPC11E)**—Starting in Junos OS Release 20.4R1, the MX2K-MPC11E MPCs in the MX2010 and MX2020 routers support the QSFP-100G-FR, QSFP-100G-DR, and QSFP-100G-LR transceivers.

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

EVPN

- **MAC-VRF with EVPN-VXLAN (MX Series and vMX routers; QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002-60C, QFX10008, and QFX10016 switches)**—Data center service providers must support multiple customers with their own routing and bridging policies in the same physical network. To accommodate this requirement, you can now configure multiple customer-specific EVPN instances (EVIs) of type **mac-vrf**, each of which can support a different EVPN service type. This configuration results in customer-specific virtual routing and forwarding (VRF) tables with MAC addresses on each Juniper Networks device that serves as a virtual tunnel endpoint (VTEP) in the EVPN-VXLAN network.

NOTE: We support MAC-VRF routing instances for EVPN unicast routes only.

To support this feature, we introduce a uniform routing instance configuration, which complies with RFC 7432, *BGP MPLS-Based Ethernet VPN*. The uniform configuration eliminates hardware restrictions that limit the number of EVIs and combinations of EVIs with their respective policies that can simultaneously exist. The common configuration includes the following new CLI elements:

- The **mac-vrf** keyword at the **[edit routing-instances *name* instance-type]** hierarchy level.
- The **service-type** configuration statement at the **[edit routing-instances *name*]** hierarchy level. We support VLAN-based, VLAN-aware, and VLAN-bundle service types.
- (QFX10000 line of switches only) The **forwarding-instance** configuration statement at the **[edit routing-instances *name*]** hierarchy level. With this optional configuration statement, you can map

multiple routing instances to a single forwarding instance. If you don't include this configuration statement, the default forwarding instance is used.

We continue to support the existing method of routing instance configuration along with the new uniform routing instance configuration.

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

- **MC-LAG emulation in an EVPN deployment (EX-Series, MX-Series, and vMX)**—Starting in Junos OS Release 20.4R1, you can emulate the function of an MC-LAG in active-standby mode in an EVPN configuration without having to configure an ICCP or ICL interface. In a standard EVPN configuration, logical interfaces configured on an aggregated Ethernet interface can have different designated forwarder election roles. To emulate an MC-LAG configuration, the designated forwarder (DF) takes on the role of the aggregated Ethernet interface. The provider edge (PE) that is the non-DF will send LACP out-of-sync packets to the CE. This will cause LACP to go down on the CE device, and the CE device will not use the links connected to the non-DF for sending traffic. If the connection between a CE and a DF PE fails, the PE is re-elected as a DF. If the connection between a CE and a non-DF PE fails, the current DF PE is not changed.

To achieve this functionality, configure the **lACP-oos-on-ndf** statement at the **[edit interfaces interface name esi df-election-granularity per-esi]** hierarchy.

- **Support for EVPN E-Tree service (MX240, MX480, and MX960)**—Starting in Junos OS 20.4R1, on MX240, MX480, and MX960 routers using MPC10E-15C-MRATE line cards you can configure an Ethernet VPN Ethernet-Tree (E-Tree) service.

[See [EVPN-ETREE Overview](#).]

High Availability (HA) and Resiliency

- **Support for pause and resume options with unified ISSU (MX Series)**—Starting in Junos OS Release 20.4R1, MX Series routers support pausing and resuming unified ISSU operations. Use the **pause** and **resume** options with the **request system software in-service-upgrade** command to control when to pause and resume unified ISSU.

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

- **NSR support for IS-IS with SR (ACX Series, MX Series)**—Starting in Junos OS Release 20.4R1, MX Series routers support NSR for IS-IS with segment routing (SR). To use NSR, you must first enable GRES on your device.

[See [Nonstop Active Routing Concepts](#)]

Interfaces and Chassis

- **464XLAT support for mobility on MS-MPC (MX Series)**—Starting in Junos OS Release 20.4R1, you can specify the IPv6 prefix length for the CLAT source address using the new command **clat-ipv6-prefix-length**. When you configure this command, NAT rules apply 464XLAT based on destination-address of the traffic, and source-address and source-prefix are no longer required. The **clat-ipv6-prefix-length** command is available at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level.

[See [translated](#) and [clat-ipv6-prefix-length](#).]

Juniper Extension Toolkit (JET)

- **Juniper Extension Toolkit (JET) support for 64-bit applications (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX ELM, JunosV Firefly, cSRX, SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM, SRX650, SRX720E, SRX750E, SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SRX4400, SRX4600, SRX4800, SRX5400, SRX5600, SRX5800, SRX7X0E, SRX-ES7, SRX-ES8, VMX, and VSRX)**—Starting in Junos OS Release 20.4R1, JET supports 64-bit applications. Use the following commands to compile 64-bit applications for use with the AMD64 or ARM64 64-bit processor architecture.
 - **mk-amd64**: Compiles the application for use with AMD64 and Junos OS with FreeBSD.
 - **mk-amd64,bsd**: Compiles the application for use with AMD64 and Junos OS with upgraded FreeBSD.
 - **mk-arm64,bsd**: Compiles the application for use with ARM64 and Junos OS with upgraded FreeBSD.

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

- **Configure inner source MAC address for flexible VXLAN tunnels (MX Series and vMX with MPC1-MPC9E or LC2101)**—Starting in Junos OS Release 20.4R1, you can use the Juniper Extension Toolkit (JET) RIB Service API to configure the source MAC address used in IPv4 and IPv6 flexible VXLAN tunnel encapsulation profiles. The source MAC addresses is stored in the inner Ethernet header of VXLAN encapsulation. If you don't specify a source MAC address, the default source MAC address 00:00:5e:00:52:01 is used to encapsulate IPv4 and IPv6 flexible VXLAN tunnels.

Use the **show route detail**, **show route extensive**, and **show flexible-tunnels profiles** CLI commands or the **get-route-information** and **get-flexible-tunnels-profiles** RPC/NETCONF commands to view the source MAC address that is specified in the flexible tunnel profile.

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

Junos OS, XML, API, and Scripting

- **Support for Certificate Authority Chain Profile (EX2300, EX3400, EX4300, MX240, MX480, MX960, PTX-5000, VMX, vSRX and QFX5200)**—Starting in Junos OS Release 20.4R1, you can configure intermediate Certificate Authority (CA) chain profile certificate and perform https REST API request using mutual and server authentications.

To configure intermediate ca-chain certificate, configure **ca-chain ca-chain** statement at the [edit system services rest https] hierarchy level.

- **Start time option for interval-based internal events that trigger event policies (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.4R1, when you create an interval-based internal event for triggering event policies, you can specify the start date and time for the initial event. To specify a start time, configure the **start-time** option along with the **time-interval** option at the [edit event-options generate-event] hierarchy level.

[See [Generating Internal Events to Trigger Event Policies.](#)]

Junos Telemetry Interface

- **JTI support for inline Junos Traffic Vision sensors with gRPC services (MX Series and PTX Series)**—Junos OS Release 20.4R1 supports inline Jflow sensors for FPC3 and MPC 1 through 9. This feature enables you to monitor inline Junos Traffic Vision (previously known as Jflow) service statistics on a router and to export statistics to an outside collector at configurable intervals using remote procedure call (gRPC) services.

Use the resource path **/junos/system/linecard/services/inline-jflow/** in a subscription to export statistics.

You can view statistics in the collector output under **/components/**. The collector component ID in the statistics output will include the FPC slot number for which inline Junos Traffic Vision statistics are exported. For example, inline Jflow statistics for FPC 0 will be under **component id 0**, and inline Jflow statistics for FPC 1 will be under **component id 1**.

Inline Junos Traffic Vision statistics are slightly different, depending on the routing platform.

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

- **JTI support for persistent active gRPC sessions between collector and server during an SSL certificate update (ACX Series, MX Series, and PTX Series)**—Junos OS Release 20.4R1 supports persistent active remote procedure call (gRPC) sessions between the collector (client) and server during an SSL certificate update.

For secure channel authentication, the TLS protocol is used to maintain a secure channel between the collector and the server. TLS uses the server certificate and the client certificate to authenticate each other and send encrypted messages over the network. When an SSL certificate is updated, existing gRPC sessions are abruptly terminated, forcing the collector to initiate a new gRPC connection and subscribe to sensors again.

To avoid this problem, you can enable persistent active gRPC sessions by configuring **hot-reloading** at the `[edit system services extension-service request-response grpc ssl]` hierarchy level. After you enable this feature, gRPC sessions will remain active even when authentication certificates are updated.

After the certificate is updated, any new gRPC session will use the updated certificate.

[See [gRPC Services for Junos Telemetry Interface](#) and [ssl](#).]

- **BGP neighbor telemetry with sharding (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, BGP neighbor telemetry with sharding (multi-threading) is supported.

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

- **LACP sensors for actor partner states on JTI (MX Series and PTX Series)**—Starting in Junos OS Release 20.4R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export LACP actor partner states (also known as LACP port states). When a subscription is configured, ON_CHANGE or periodic streaming statistics are sent from devices to an outside collector.

You can subscribe to `/lacpd/` to collect all statistics or include the following resource paths individually in a subscription:

- `/lacpd/ae/member/partner_collecting`
- `/lacpd/ae/member/partner_synchronization`
- `/lacpd/ae/member/partner_timeout`
- `/lacpd/ae/member/partner_aggregatable`
- `/lacpd/ae/member/partner_distributing`
- `/junos/system/linecard/interface/traffic/`

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

- **Juniper Resiliency Interface for exception reporting and null route detection (ACX Series, PTX Series and MX Series)**—Starting in Junos OS Release 20.4R1, you can use Juniper Resiliency Interface to detect and reduce Mean Time to Repair (MTTR) first-order network issues. Juniper Resiliency Interface uses a push model for data reporting from the entities in the system which encounter packet drops. This automates the workflow for detecting, reporting, and mitigating adverse exceptions.

To collect kernel routing table and routing protocol process exceptions, configure the **set system resiliency exceptions** statement at the `[edit]` hierarchy level to specify exception reporting based on kernel exceptions, and routing exceptions.

You can display exceptions from a remote collector by means of remote procedure call (gRPC) services or gRPC network management interface (gNMI) services. Display on-box exceptions by accessing the `/var/log` file or the database at `/var/db/ResiliencyExceptions.db`. No Junos operational mode commands display these exceptions.

MPLS

- **Re-engineering of SR-TE (MX Series, PTX Series)**—Starting with Junos OS Release 20.4R1, you can incorporate the following features to enhance the debugging capability of segment routing traffic-engineering (SR-TE):
 - rib-group import functionality.
 - Display of SR-TE routes installed from various tunnel sources using the `show spring-traffic-engineering` command.
 - Template map for BGP SR-TE tunnels.
 - Compute profile in template with distributed Constrained Shortest Path First (CSPF) for dynamic SR-TE tunnels.
 - 6PE (IPv6 over IPv4 SR-TE tunnel)
 - no-chained-composite-next-hop option

[See [source-packet-routing](#) and [show spring-traffic-engineering](#).]

- **Support for optimizing auto-bandwidth adjustments for MPLS LSPs (MX Series and PTX Series)**—Starting in Junos OS Release 20.4R1, you can configure faster auto-bandwidth adjustment for MPLS LSPs under overflow or underflow conditions. This feature decreases the minimum allowed **adjust-threshold-overflow-limit** and **adjust-interval** to 150 seconds when **adjust-threshold-overflow-limit** and **adjust-threshold-underflow-limit** cross the configured threshold values. In releases earlier than Junos OS Evolved Release 20.4R1, the **adjust-interval** is 300 seconds under overflow or underflow conditions.

You can configure faster in-place LSP bandwidth update that avoids signaling of a new LSP instance as part of make-before-break. To configure faster in-place LSP bandwidth update, include the **in-place-lsp-bandwidth-update** configuration statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

You can also configure RSVP interfaces to support subscription percentage per priority. To configure subscription percentage per priority, include the **subscription priority *priority* percent *value*** configuration statement at the **[edit protocols rsvp interface *interface-name*]** hierarchy level.

[See [Configuring Optimized Auto-bandwidth Adjustments for MPLS LSPs](#).]

- **Support for express segments to establish end-to-end segment routing path (MX Series and PTX Series)**—Starting in Junos OS Release 20.4R1, express segments can be used to establish end-to-end TE paths between interconnected TE networks. Express segments (also known as virtual TE links) are generated dynamically through policies matching the underlay LSPs. Express segments and the corresponding abstracted topology (required by RFC7926) is generated with policies.

To apply a policy, include the **policy *policy-name*** statement at the **[edit protocols express-segment traffic-engineering]** hierarchy level.

To configure express segment, include the **express-segment** statement under the **[edit protocols]** hierarchy level.

[See *How to Establish End-to-End Segment Routing Paths Using Express Segments*.]

Network Management and Monitoring

- **Configuration support to prevent drifting of accounting records (MX Series routers, vMX)**—You can configure accounting records to record data in accounting files and archive the accounting files to analyze the information collected. Drifting of the accounting records happens if the time at which the records are written to the accounting file spills beyond the transfer window of the file. Starting in Junos OS Release 20.4R1, to prevent drifting of accounting records:

- Use the **start-time** statement with the accounting profiles (class-usage-profile, filter-profile, flat-file-profile, interface-profile, mib-profile, and routing-engine-profile) to have a predictable start time of the profiles.
- Use the **timestamp** statement with the **request accounting add records** command to record the timestamp externally instead of epoch timestamp when the command is executed.

[See [routing-engine-profile](#), [class-usage-profile](#), [interface-profile](#), [filter-profile](#), [mib-profile](#), [flat-file-profile](#).]

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The **<get-configuration>** operation supports the **compare="configuration-revision"** and **configuration-revision** attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

- **Support for an extension to the rpm-tracked static routes (MX Series, PTX Series, and vMX)**—Starting in Junos OS Release 20.4R1, you can configure route preference and tag values for each destination-prefix. This feature supports both IPv4 and IPv6 rpm-tracked static routes.

[See [show route rpm-tracking](#).]

Limitations

Qualified next hop is not supported with rpm-tracked static routes. Hence, the setting of preference, metric, and tags applies only to the **rpm-tracking** static route and not to the related next hops.

Routing Policy and Firewall Filters

- **Support for route's next-hop weight in policy match condition (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, a route with multiple next-hop paths can use the weight associated with a path to identify primary and backup paths. The path with the lowest weight is used as the primary path, and any paths with higher weights are treated as backup paths. You can use the next-hop weight as a match condition in export policies to redistribute IGP and BGP routes based on whether the primary or backup paths are active.

Configure this match condition using the **[edit policy-options policy-statement *policy-name* term *term-name* from]** statement.

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

Routing Protocols

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC 6286 (MX204, NFX Series, PTX5000, QFX Series, and vRR)**—Starting in Junos OS Release 20.4R1, you can establish a BGP connection using a BGP identifier that is a 4-octet, unsigned, nonzero integer and it needs to be unique only within the autonomous system (AS) per RFC 6286. In earlier releases, the BGP ID of a BGP speaker was required to be a valid IPv4 host address assigned to the BGP speaker.

To enable this feature, use the **bgp-identifier *identifier* group *bgp group name* bgp-identifier *identifier* neighbor *peer address* bgp-identifier *identifier*** configuration statement at the **[edit protocols bgp]** hierarchy level.

[See [router-id](#)]

- **Support for multiple single-hop EBGP sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGP session. You can now enable single-hop EBGP sessions on different links over multiple directly-connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGP Sessions on Different Links Using the Same Link-Local Address \(IPv6\)](#).]

- **Support for IPv6 L3VPN over IPv6 SR-TE and IPv6 Underlay (MX Series)**—Starting in Junos OS Release 20.4R1, You can configure an IPv6 Layer3 VPN connection with an IPv6 local address and an IPv6

neighbor address. We have extended BGP support for IPv6 Layer 3 VPN over BGP IPv6 SR-TE in IS-IS networks. You can connect an IPv6 provider edge device with a colored or non-colored IPv6 penultimate nexthop (PNH) address mapped to IPv6 SR-TE tunnels.

To configure an IPv6 address for Layer 3 VPN connection, include the **family inet6-vpn** configuration statement at the **[edit protocols bgp group name]** hierarchy level.

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

- **Support for BGP Labeled Unicast prefix SID (MX Series and PTX Series)**—Starting in Junos OS 20.4R1, BGP labeled unicast can carry segment routing global block label range and index information through the prefix segment attribute. With this feature we support segment routing using the BGP labeled unicast prefix segments and the MPLS data plane in medium to large scaled data centers. The controller directs the server to assign a stack- of labels to an incoming packet based on the available network state information. The assigned label stack avoids congested paths and steers the packet through a best available path.

To configure and advertise the SRGB label range specifically for BGP include the **source-packet-routing srgb start-label start-label index-range index-range** and **advertise-srgb** configuration statements at the **[edit protocols bgp]** hierarchy level.

To advertise prefix SIDs to external BGP peers, include the **advertise-prefix-sid** configuration statement at the **[edit protocols bgp]** hierarchy level. You can configure this statement globally or for specific BGP groups or BGP neighbors.

[See [srgb.](#)]

- **Support for SRv6 network programming and Layer 3 Services over SRv6 in BGP (MX Series)**—Starting in Junos OS Release 20.4R1, you can configure BGP based Layer 3 service over SRv6 core. You can enable Layer 3 overlay services with BGP as control plane and SRv6 as dataplane. SRv6 network programming provides flexibility to leverage segment routing without deploying MPLS. Such networks depend only on the IPv6 headers and header extensions for transmitting data.

To configure IPv4 and IPv6 transport over SRv6 core, include the **end-dt4-sid sid** and the **end-dt6-sid sid** statements at the **[edit protocols bgp source-packet-routing srv6 locator name]** hierarchy level.

To configure IPv4 VPN and IPv6 VPN service over SRv6 core, include the **end-dt4-sid sid** and the **end-dt6-sid sid** statements at the **[edit routing-instances routing-instance name protocols bgp source-packet-routing srv6 locator name]** hierarchy level.

[See [Understanding SRv6 Network Programming and Layer 3 Services over SRv6 in BGP.](#)]

- **Support for unicast ARP request on table entry expiration (MX Series)**—Starting in Junos OS Release 20.4R1, you can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces overall ARP broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and to

instead translate ARP broadcasts to unicast requests. You can verify whether this feature is configured by using the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

- **IPv6 support in TED (MX Series, PTX Series)**—Starting in Junos OS Release 20.4R1, you can configure IS-IS traffic engineering to store IPv6 information in the traffic engineering database (TED) in addition to IPv4 addresses. BGP-LS distributes this information as routes from the TED to the Isdist.0 routing table using the TED import policies. These routes are advertised to BGP-TE peers as network layer reachability information (NLRI) with IPv6 router ID type, length, and value (TLV).

With this enhancement, you can benefit from obtaining the complete network topology in the TED.

[See [Link-State Distribution Using BGP Overview](#).]

Services Applications

- **Support for passive flow monitoring (MX Series)**—Starting in Junos OS Release 20.4R1, you can configure passive flow monitoring on these routers:
 - MX240/MX480/MX960/MX2008/MX2010/MX2020 routers with either the MPC7E-MRATE or MPC7E-10G line card
 - MX10008 router with the JNP10K-LC2101 line card

[See [Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers](#) and [Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers](#).]

Software Defined Networking

- **PCEP support for color (MX480, QFX5200)**—Starting in Junos OS Release 20.4R1, the Path Computation Element Protocol (PCEP) supports color for colored segment routing LSPs. This includes Path Computation Element (PCE)-initiated, Path Computation Client (PCC)-controlled, and PCC-delegated segment routing LSPs. With this PCEP extension, you can configure candidate paths based on color and endpoints, where the active candidate path is the path with the highest segment routing preference, or based on source priority.

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

- **Support for ECMP on multiple flexible routes (MX80, MX104, MX204, MX10003, and vMX routers)**—Starting in Junos OS Release 20.4R1, we support load balancing of traffic over multiple flexible routes with 64-way ECMP. A flexible route is a static route with a tunnel encapsulation profile, which has the flexible tunnel interface (FTI) attribute. Flexible routes are installed on Juniper gateway devices using the Juniper Extension Toolkit (JET) APIs. Multiple flexible routes can go over the same logical interface. When a packet is received with the flexible route as the destination address, the packet is processed using the profile associated with a flexible route. Traffic across multiple flexible routes is load-balanced based on the traffic priority.

Use the **show route** and **show route extensive** CLI commands or the **get-route-information** RPC/NETCONF command to view details about a flexible route for a destination address.

[See [Understanding Programmable Flexible VXLAN Tunnels](#).]

- **Static VXLAN at VLAN or bridge domain level (MX5, MX10, MX40, MX80, MX150, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016 routers and QFX5120-32C, QFX5120-48T, and QFX5120-48Y switches)**—In Junos OS Release 20.3R1 and earlier, we supported the configuration of static VXLAN at the global level only. By including the **remote-vtep-list** configuration statement at the **[edit switch-options]** or **[edit routing-instances name]** hierarchy level, you can map all local VLANs or bridge domains to the remote virtual tunnel endpoints (VTEPs) in the list.

Starting in Junos OS Release 20.4R1, you can also configure static VXLAN at the VLAN or bridge domain level using the **static-remote-vtep-list** configuration statement at the **[edit vlans name vxlan]**, **[edit bridge-domains name vxlan]**, or **[edit routing-instances name bridge-domains name vxlan]** hierarchy level.

When specifying remote VTEPs at the VLAN level in the default switching instance, you must also specify the same VTEPs at the global level in the default switching instance. Or when specifying remote VTEPs at the bridge domain level in a routing instance, you must also specify the same VTEPs at the global level in the same routing instance. For example, if you specify a VTEP in the **static-remote-vtep-list** at the **[edit routing-instances name bridge-domains name vxlan]** hierarchy level, you must also specify the VTEP in the **remote-vtep-list** at the **[edit routing-instances name]** hierarchy level.

To replicate and flood BUM traffic, you must specify the **ingress-node-replication** configuration statement at the **[edit vlans name vxlan]**, **[edit bridge-domains name vxlan]**, or **[edit routing-instances name bridge-domains name vxlan]** hierarchy level. This configuration restricts the BUM traffic flood domain to only those VTEPs mapped to a particular bridge domain or VLAN.

[See [Static VXLAN](#) and [static-remote-vtep-list](#).]

Software Installation and Upgrade

- **Zero touch provisioning (ZTP) with IPv6 support (MX Series)**—Starting in Junos OS Release 20.4R1, 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.

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

Software Licensing

- **MX Series devices using SPC3 services card for IPsec VPN services requires a feature license (MX Series)**—Starting in Junos OS Release 20.4R1, you must have a valid license to use the IPsec VPN feature running on MX Series devices with SPC3 services card. For high availability, you must install IPsec base license on both the nodes.

This is a binary license, without an installed license your license count is 0, and your license count is 1 when a valid license is installed in the system.

IPsec VPNs tunnels doesn't establish without a valid license in the device, current active tunnels will stay up if the license expires. IPsec VPN tunnels that are brought down after the license expiry doesn't re-establish until you install a valid license.

[See [Software Features That Require Licenses on MX Series Routers Only](#), [MX FLEX Software License Model](#), and [Managing Licenses](#).]

Subscriber Management and Services

- **Support for mobility on Junos Multi-Access User Plane (MX204, MX240, MX480, MX960, MX10003)**—For Junos OS Release 19.4R1, we introduced Junos Multi-Access User Plane supporting a combined SGW-U/PGW-U (SAEGW-U) on MX Series routers in accordance with 3GPP Release 14 CUPS architecture. This provided high-throughput 4G and 5G fixed-wireless access service with support for 5G non-stand-alone (NSA) mode.

For Junos OS Release 20.4R1, we introduce support for running an MX router as either a standalone SGW-U or a standalone PGW-U or a combined SAEGW-U to provide high-throughput 4G and 5G mobility service (relocation of a UE to a new eNodeB, new SGW-U, or new SAEGW-U). This includes support for GTP-U based S5-U and S8-U interfaces, to provide links between SGW-U and PGW-U

devices, and tunnel relay functionality to forward user plane traffic between S1-U and S5-U/S8-U interfaces or between S5-U/S8-U and SGi interfaces respectively. We support the following mobility scenarios:

- Handover with eNodeB and no SGW change
- Handover with SGW change (direct forwarding)
- Handover with SGW change (indirect forwarding)

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

- **Support for 5G Junos Multi-Access User Plane (MX204, MX240, MX480, MX960, MX10003)**—Starting with Junos OS Release 20.4R1, Junos Multi-Access User Plan supports 3GPP TS 29.244 Release 15, which includes support for the 5G user plane function (UPF). Specifically, these enhancements are provided:

- PDI optimization for Sx messages
- GTP path management via heartbeats
- Support for User ID in PFCP Session Establishment Request
- Support for QoS control/enforcement at the bearer level
- Support for DDOS over Sx interface.

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

- **Support for increased number of pseudowire logical interface devices (MX2010 and MX2020)**—Starting in Junos OS Release 20.4R1, you can configure up to 18,000 pseudowire logical interface devices on the MX2010 and MX2020 routers with the MX2K-MPC9E or MX2K-MPC11E line card. Use the **device-count** statement at the **[edit chassis pseudowire-service]** hierarchy level.

[See [Configuring the Maximum Number of Pseudowire Logical Interface Devices Supported on the Router](#) and [device-count](#).]

- **IPv4 reassembly for fragmented soft GRE packets on the WAG (MX Series)**—Starting in Junos OS Release 20.4R1, you can enable a Wi-Fi Access Gateway (WAG) to reassemble fragmented GRE packets that the WAG receives from a Wi-Fi access point over a soft GRE tunnel.

[See [dynamic-profiles](#) and [Wi-Fi Access Gateways](#).]

System Management

- **1-Gbps support on all ports of MPC7E-10G line cards (MX240, MX480, and MX960)**—Starting in Junos OS Release 20.4R1, you can configure 1-Gbps speed on all 40 10-Gbps Ethernet ports of the MPC7E-10G line cards. The 1-Gbps interface supports the following features:
 - Synchronous Ethernet
 - Link aggregation group (LAG)

- G.8275.1 Precision Time Protocol (PTP) profile
- Hybrid mode

To configure an interface to operate at the 1-Gbps speed, use the **set interfaces interface-name gigether-options speed 1g/10g** command at the **[edit]** hierarchy level.

[See [Precision Time Protocol Overview](#), [Synchronous Ethernet Overview](#), and [Hardware Compatibility Tool](#).]

System Logging

- **Support for time averaged watermark (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, you can capture steady state data of routing and forwarding (RIB/FIB) table routes using the **time-averaged-watermark-interval** configuration statement at the **[edit routing-options]** hierarchy level. Time averaged watermark is calculated whenever the time averaged interval is changed from CLI. Time averaged watermark is logged in syslog if the logs are enabled in the system at **LOG_NOTICE** level. The default time averaged watermark interval is 1 day. You can see the timed averaged watermark using the existing **show route summary** command.

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

SEE ALSO

What's Changed 106
Known Limitations 112
Open Issues 113
Resolved Issues 120
Documentation Updates 135
Migration, Upgrade, and Downgrade Instructions 135

What's Changed

IN THIS SECTION

- [Class of Service \(CoS\) | 107](#)
- [EVPN | 107](#)
- [General Routing | 108](#)

- High Availability (HA) and Resiliency | 109
- Interfaces and Chassis | 109
- J-Web | 110
- MPLS | 110
- Network Management and Monitoring | 111
- Platform and Infrastructure | 111
- Routing Protocols | 111
- User Interface and Configuration | 111

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.4R1 for MX Series routers.

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>`.

EVPN

- **Updated XML output for show evpn p2mp**—Starting with this release, when you pipe the output of the **show evpn p2mp** command to the **display xml** option, Junos OS now returns an XML output with a subtree structure for each neighbor. Prior to this release, the display XML returns an XML output with all the neighbors under one tree structure.
- **New output flag for the show bridge mac-ip table command**—The Layer 2 address learning daemon does not send updated MAC and IP Address advertisements to the Routing Protocol daemon when an IRB interface is disabled in an EVPN-VXLAN network. Junos has 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](#).]

General Routing

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

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

Change in show oam ethernet connectivity-fault-management interface command (MX Series)— You can now view the counter values for continuity check messages (CCM) inline sessions sent messages for MPC10E (MPC10E-10C-MRATE and MPC10E-15C-MRATE) and MPC11E (MX2K-MPC11E) line cards only when you execute the **show oam connectivity-fault-management interfaces** command three times. If you execute the command twice, the values are incorrectly displayed.

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

- **MS-MPC and MS-MIC service package (MX240, MX480, MX960, MX2020, MX2010, and MX2008)**—PICs of the MS-MPC and MS-MIC do not support any service package other than extension-provider. If you try to configure any other service package for these PICs by using the **set chassis fpc slot-number pic pic-number adaptive-services service-package** command, 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.

[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 are supported.

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

- **Updates to ON-CHANGE and periodic dynamic subscriber interface metadata sensors (MX Series routers and EX9200 line of switches)**—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 \(Junos Telemetry Interface\)](#). 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](#) .]

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#) .]

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

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

- **New TLV types and TLV type values in output field (MX960 and vMX)**—We've introduced TLV SR policy identifier, TLV SR candidate path identifier, and TLV SR preference fields in the output for the **show path-computation-client tlv-types** command. These new output fields help you in easily fetching the TLV type values used by PCCD irrespective of whether the type values are experimental or standardized.

High Availability (HA) and Resiliency

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

Interfaces and Chassis

- **Change in <range> XML tag (MX480)**—Starting in Junos OS, we've changed the <range> string </range> XML tag to <transport-range> <transport-range-info> string </transport-range-info> <transport-range-suspect-flag> string </transport-range-suspect-flag> <transport-range-reason> string </transport-range-reason> </transport-range> under the output of the **show interfaces transport pm optics current interface | display hierarchy** command. Hence, the new XML tags that associate the values

to the range-info, range-suspect-flag, range-reason tags map the information to the given **show interfaces transport pm optics current | display entry** command.

[See [Supported OTN Options on MX Series Routers](#).]

- **Hardware assisted timestamping**—By default, hardware assistance is used for timestamping Ethernet frame delay frames on AFT based MX Series line cards, even if the hardware-assisted-timestamping is not configured.

J-Web

- **Adobe Flash Player support (MX Series)**—Adobe Flash Player support ends on December 31, 2020. As a result, starting in Junos OS Release 20.4R1, the following J-Web pages will not be supported:
 - Monitor > System View > Process Details
 - Monitor > Routing > OSPF Information

The Monitor > Interfaces page is supported. However, the Flash components are removed. In addition, these monitor pages will not load correctly for Junos OS Release 20.3R1 and earlier releases.

MPLS

- **The show mpls lsp extensivel and show mpls lsp detail commands display next-hop gateway LSPid**—When you use the **show mpls lsp extensivel** and **show mpls lsp detail** commands, you'll see next-hop gateway LSPid in the output.
- **Disable back-off behavior on PSB2 (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**— We've introduced the **cspf-backoff-time** statement globally for MPLS and LSP to delay the CSPF by configured number of seconds, on receiving bandwidth unavailable PathErr on PSB2. If the configured value is zero, then the CSPF starts immediately for PSB2, when bandwidth-unavailable PathErr is received. If the statement is not configured, the default exponential back-off occurs.

[See [cspf-backoff-time](#).]

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to "deviate not-supported" nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.

Platform and Infrastructure

- **Firewall Filters Application**—You can configure the proto family as part of the filter name.

[See [Understanding Multiple Firewall Filters Applied as a List](#).]

Routing Protocols

- **Loading of the default configurations in a RIFT package causes the following changes**—
 1. Output of the **show rift node status** command displays the node ID in hexadecimal number even though the node ID is configured in decimal, hexadecimal, or octal number.
 2. Some of the DDoS default configurations change because of the DDoS protection interferes with the RIFT BFD operation.

User Interface and Configuration

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

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

SEE ALSO

[What's New](#) | 90

[Known Limitations](#) | 112

[Open Issues](#) | 113

[Resolved Issues | 120](#)

[Documentation Updates | 135](#)

[Migration, Upgrade, and Downgrade Instructions | 135](#)

Known Limitations

IN THIS SECTION

- [General Routing | 112](#)
- [Interfaces and Chassis | 113](#)
- [MPLS | 113](#)
- [Network Management and Monitoring | 113](#)

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

General Routing

- Subscriber access facing CPU utilization of FPC remains 100 percent for 56 minutes after making changes to the service firewall filter configuration. [PR1447003](#)
- The NPC process continuously generates core file at **Trinity_Ktree::Trinity_FourWayBlock, Trinity_Ktree::walkSubTree** due to the NH memory exhaustion with the NH explosion. The rpd and srdd processes start hogging and the system becomes unstable. [PR1538029](#)
- On the MX480 router, expected probes are not observed while configuring and testing the Packet Forwarding Engine based RPM for the probe type icmp-ping. [PR1556697](#)
- On the MX960 routers, **spring-traffic-engineering lsp count** is not displayed as expected while validating 32000 inter-domain DCPSF LSPs. [PR1561947](#)
- SyncE source across multiple line cards cannot be used in PTP-Hybrid source as fallback clock due to PTP lanes limitations used for SyncE clock as well in SCBE2 and MPC1-9 types. [PR1536013](#)

Interfaces and Chassis

- For MC-LAG to work properly, the mc-ae interface should be configured on both the PE devices. A scenario where the mc-ae interface is deleted, deactivated, or not configured on one of the devices is a case of misconfiguration. Juniper Networks does not support such a scenario because it can lead to traffic loss and other unexpected behavior. [PR1536831](#)
- On the MPC10 line cards, DMRs or SLRs are not received with EVPN up mep on the aggregated Ethernet interface with normalization. [PR1543641](#)
- UP MEP CFM sessions over bridge-domains or VPLSs, which have ports hosted on FPC(s) does not support ISSU. The sessions needs to be explicitly deactivate and activate to recover post ISSU. [PR1543656](#)

MPLS

- The rpd process might crash.[PR1461468](#)

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](#)

SEE ALSO

What's New 90
What's Changed 106
Open Issues 113
Resolved Issues 120
Documentation Updates 135
Migration, Upgrade, and Downgrade Instructions 135

Open Issues

IN THIS SECTION

- [Class of Service \(CoS\) | 114](#)
- [EVPN | 114](#)

- Forwarding and Sampling | 114
- General Routing | 115
- Infrastructure | 117
- Interfaces and Chassis | 117
- Juniper Extension Toolkit (JET) | 117
- Layer 2 Ethernet Services | 118
- MPLS | 118
- Platform and Infrastructure | 118
- Routing Policy and Firewall Filters | 119
- Routing Protocols | 119
- User Interface and Configuration | 119
- VPNs | 119

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)

- Unexpected behavior of the Class of Service is observed with the wildcard classifier. [PR1559516](#)

EVPN

- The rpd process might leak memory when the EVPN configuration is changed. [PR1540788](#)

Forwarding and Sampling

- Packet length for ICMPv6 is displayed as 0 in the output of the **show firewall log detail** command. [PR1184624](#)
- After restarting routing, 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 that the Layer 2 learning daemon had prior to restart. This causes a mismatch between the Layer 2 learning and routing daemon's interpretation as to where the MAC IP address entries are learned, either local or remote, leading to the MAC IP table being out of synchronization. [PR1452990](#)

General Routing

- On the MX104 router, if you use the **snmpbulkget** or **snmpbulkwalk** (for example, used by the SNMP server) on a chassisd-related component (for example, jnxOperatingEntry), high CPU usage and slow response of the chassis process (chassisd) might be observed because of a hardware limitation, which might also lead to a query timeout on the SNMP client. This issue might not be observed while using an SNMP query. [PR1103870](#)
- On the MX104 router, CPU hog or busy state occurs with the sporadic L2C access error message and false alarms. [PR1223979](#)
- Online insertion and removal of a MIC in an MPC might lead to traffic destined to the MPC being silently dropped or discarded. [PR1350103](#)
- 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](#)
- Subscribing to **/linecard/packet/usage** and triggering the UDP decoder, the hardware statistics are exported with improper hierarchy. [PR1485739](#)
- The following critical syslog error messages at FPC3 user.crit aftd-trio are observed 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](#)
- Login or logout of high scale (around 1 million bearers) causes some sessions not to re-login. [PR1489665](#)
- The following error messages are observed: **unable to set line-side lane config (err 30).** [PR1492162](#)
- Backup Routing Engine reboots because of power cycle or failure when the offline and online operations are performed on CB1. [PR1497592](#)
- The log file to log the activities associated with the **request rift package activate** command is created with the permissions of the user. If multiple users run the command, the command might fail due to the write permission error. [PR1514046](#)
- LFM might flap during MX-VC ISSU. [PR1516744](#)
- All SFBs might go offline due to fabric failure and fabric self-ping probes performing the **disable-pfe** action. [PR1535787](#)
- On the MX2020 router, the next hops are less than a total of nhdb 4MPOST GRES. [PR1539305](#)
- Even though enhanced-ip is active, the following alarm is observed during ISSU: **RE0 network-service mode mismatch between configuration and kernel setting.** [PR1546002](#)
- Some routes might get incorrectly programmed in the forwarding table in the kernel which is no longer present in rpd. [PR1534455](#)
- The chassisd memory leak might cause traffic loss. [PR1537194](#)
- PTP to PTP noise transfer test fails. [PR1543982](#)

- On the MX480 router, the expected probes are not observed while configuring and testing the Packet Forwarding Engine based RPM for the icmp-ping probe type. [PR1556697](#)
- The 2 way average time error (cTE) test fails when primary and slave are on different line cards. [PR1557636](#)
- The SyncE transient response test fails. [PR1557999](#)
- Inconsistent **core.python2.7.mpc0** core is observed with stacktrace at **ea_wi_precl** and **ea_macsec_receive**. [PR1534568](#)
- The NGMPC2 process generates core file at **bv_entry_active_here::bv_vector_op::gmph_reevaluate_group::gmph_destroy_client_group**. [PR1537846](#)
- Validation of the OCSP certificate might not go through in case of certain CA servers. [PR1548268](#)
- On the MX480 routers, the following error message is observed at **rts_marker_rcv_timo: Slave peer did not send marker ack for last 360 secs for vks 0 slave_ack=0** during ISSU in MXVC. [PR1550492](#)
- On the MX480 router, the output of the **show interfaces transport pm otn current < interface>** command is not as expected. [PR1560533](#)
- On the MX960 router, mismatch between the Yang schema and RPC output is observed. [PR1559810](#)
- When the system has only one plane (in the process of plane offline or online), the MPC10-10c line card destination errors are observed. [PR1560053](#)
- On the MX240 router, the following error message is observed: **On R0 Overlay Ping FAILS tunnel-src 10.255.0.53 tunnel-dst 10.255.0.139 vni 1, invalid VNI: '1'**. [PR1560408](#)
- On the MPC11 line card, the SPMB alarm gets generated after GRES. [PR1560898](#)
- On the MPC11 line card, error messages are not simulated using the **show log message** command on injecting an error. [PR1560920](#)
- Core file is found at **re-MX104-ABB-0.gz.core.0.gz >clksync_geneva_delete_ptp_loc_entry> clksync_geneva_ptp_add_clock_entry> clksync_ptp_stream_proc_op> clksync_event_update> clksync_process_event**. [PR1561004](#)
- The rpd process might crash with the BGP RIB sharding enabled sometimes when the routing-instances are deleted and recreated quickly. [PR1562905](#)
- The following error message for port might be seen: **FAILED(-1) read of SFP eeprom**. [PR1529939](#)
- The CFM sessions goes down during FRU upgrade stage of ISSU. [PR1534628](#)
- On the MX480 routers, sflow log error message are seen when egress sampling is enabled on the dynamic IP-IP tunnel encapsulation scenario. [PR1538863](#)
- The DHCP discover packet might be dropped if the DHCP inform packet is received first. [PR1542400](#)
- Continuous rpd errors might be seen and new routes fails to be programmed by rpd. [PR1545463](#)
- On the MPC7E line cards, -1pps cte test fails. [PR1546219](#)

- Heap malloc(0) detected for **jnh_unilist_adaptive_add** on loading configurations. [PR1547240](#)
- Commit error is introduced during the **deactivate chassis synchronization source** and **smc-transmit all** configurations. [PR1549051](#)
- On the MX480 routers, the MPC7E 1g interfaces are down on router during restart scenarios but remote side interfaces are up and running. [PR1554406](#)
- Traffic is not forwarded over IRB to Layer 2 circuit on It interfaces. [PR1554908](#)
- Captive portal for phone-home bootstrap process is not supported. [PR1555112](#)
- Configuring HFRR link-protection on an interface might cause rpd to crash. [PR1555866](#)
- Upgrading satellite devices might lead to some SDs in the **SyncWait** state. [PR1556850](#)
- The l2cpd process might generate core files on reboot. [PR1561235](#)
- On the MX240 routers, VIA headers are not changed properly when SIP ALG is enabled. [PR1561312](#)
- Traffic drop is seen after the Layer 2 GRES switchovers with Layer 2 forwarding database. [PR1561344](#)

Infrastructure

- HSRPv2 IPv6 packets might get dropped if IGMP-snooping is enabled. [PR1232403](#)

Interfaces and Chassis

- The VCP port is marked as administratively down on the wrong MX-VC member. [PR1552588](#)
- The OAM Ethernet connectivity-fault-management interfaces detail are not as expected. [PR1559375](#)
- On the MX960 routers, the LFM sessions are flapped after applying the Action profile on the router. [PR1561044](#)
- On the MX10003 router, traffic loss issue is observed while verifying the VRRP state machine functionality. [PR1564551](#)

Juniper Extension Toolkit (JET)

- Abrupt shutdown or closure of the collector port results in the grpc connection with same client ID to fail until the devices detect the disconnect. [PR1549044](#)
- The following error message is observed while creating the default grpc channel with name **fw_channel**: **Issue in channel creation.** [PR1559064](#)

Layer 2 Ethernet Services

- OSPF and OSPF3 adjacency uptime is more than expected after NSSU upgrade and outage is higher than the expected. [PR1551925](#)

MPLS

- The rpd process generates core file at `rsvp_enh_lp_defer_backup_psb_creation,rsvp_psb_request_backup_psb,rsvp_telink_down`. [PR1560059](#)
- Traffic loss might be observed due to rpd crash in the MPLS scenario. [PR1528460](#)

Platform and Infrastructure

- The CFM REMOTE MEP does not come up after configuration or if the MEP remains in the **Start** state. [PR1460555](#)
- The following line card errors are seen: **HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invalid IRB topo/ IRB ifl zero in l2 nh 40495 add IRB**. [PR1472222](#)
- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)
- Interoperability between the MPC7 and MPC10 line cards cause traffic drop in an Layer 2 overgre scenario. [PR1558114](#)
- Interfaces statistics are not updated on an aggregated Ethernet interface as expected with the CCCOAE configurations. [PR1561304](#)
- On the MX2020 routers, LMM TXED is not same as LMR RX as packet loss with cycle time 100ms. [PR1561397](#)
- On the MX480 routers, verification of GRES and NSR functionality with VXLAN feature, the convergence is not as expected as Layer 2- DOMAIN-TO-Layer 3 VXLAN. [PR1520626](#)
- The vmxt_lnx process generates core file at **KtreeSpace::FourWayLeftAttachedNode::getNextDirty Trinity_Ktree::walkSubTree Trinity_Ktree::walkSubTree**. [PR1525594](#)
- No sessions are found on service-set ss2. [PR1549259](#)
- The BFD session goes down after ISSU switchover phase. [PR1561306](#)

Routing Policy and Firewall Filters

- Generate route goes to the **Hidden** state when the **protect core** statement is enabled. [PR1562867](#)

Routing Protocols

- The BFD session might get stuck in the **Init** or **Down** state after the BFD session flaps. [PR1474521](#)
- Some PIM join or prune packets might not be processed in the first attempt in the scale scenario where the PIM routers establish neighborhood and immediately join the multicast group. [PR1500125](#)
- Sending multicast traffic to downstream receiver on Trio based Virtual Chassis platforms might fail. [PR1555518](#)
- On the MX960 router, the next-hop entries in table inet.0 is not as expected when testing the IS-IS policy. [PR1558581](#)
- The IS-IS adjacency is not as expected when testing the BGP community feature using VRR. [PR1559079](#)
- 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](#)
- Multipath information keeps showing up for the BGP route even after disabling the interface for one path. [PR1557604](#)
- The ssh connections are allowed more than the configured ssh connection-limit. [PR1559305](#)
- Wrong SPF calculation might be observed for OSPF with **ldp-synchronization hold-time** configured after interface flap. [PR1561414](#)

User Interface and Configuration

- The **request system software validate on host** command does not validate the correct configuration file. [PR1553577](#)

VPNs

- The Layer 2 circuit states are not updated instantly and double hit in traffic is observed after disabling the core interface on MC-LAG active node. [PR1543408](#)
- 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	 90
What's Changed	 106
Known Limitations	 112
Resolved Issues	 120
Documentation Updates	 135
Migration, Upgrade, and Downgrade Instructions	 135

Resolved Issues

IN THIS SECTION

- [EVPN](#) | [121](#)
- [Forwarding and Sampling](#) | [122](#)
- [General Routing](#) | [122](#)
- [Infrastructure](#) | [128](#)
- [Interfaces and Chassis](#) | [129](#)
- [Intrusion Detection and Prevention \(IDP\)](#) | [130](#)
- [Juniper Extension Toolkit \(JET\)](#) | [130](#)
- [J-Web](#) | [130](#)
- [Layer 2 Ethernet Services](#) | [130](#)
- [Layer 2 Features](#) | [130](#)
- [MPLS](#) | [130](#)
- [Network Address Translation \(NAT\)](#) | [131](#)
- [Network Management and Monitoring](#) | [131](#)
- [Platform and Infrastructure](#) | [131](#)
- [Routing Policy and Firewall Filters](#) | [132](#)
- [Routing Protocols](#) | [132](#)
- [Services Applications](#) | [134](#)
- [Subscriber Access Management](#) | [134](#)
- [User Interface and Configuration](#) | [134](#)
- [VPNs](#) | [134](#)

This section lists the issues fixed in Junos OS Release 20.4R1 for MX Series routers.

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

EVPN

- EVPN-VXLAN core isolation does not work when the system is rebooted or the routing is restarted. [PR1461795](#)
- 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](#)
- With dynamic list next hop configured, a forwarding problem occurs after performing graceful switchover. [PR1513759](#)
- ARP might break when multicast snooping is enabled in EVPN for the VLAN-based and VLAN-bundle service scenarios. [PR1515927](#)
- no-arp-suppression is required for MAC learning across the EVPN domain on the static VTEP. [PR1517591](#)
- Unable to create a new VTEP interface. [PR1520078](#)
- The ARP resolution to the gateway IRB address fails if **decapsulate-accept-inner-vlan** or **encapsulate-inner-vlan** is configured. [PR1526618](#)
- The rpd process might crash when **auto-service-id** is configured in the EVPN-VPWS scenario. [PR1530991](#)
- All the ARP reply packets toward some address are flooded across the entire fabric. [PR1535515](#)
- The **GE LOS** alarm logs on the change in IFF_CCCDOWN are not logged in the syslog message file. [PR1539146](#)
- VLAN ID information is missed while installing the EVPN route from the BGP type 2 route after modifying a routing-instance from the instance-type EVPN to instance-type virtual-switch. [PR1547275](#)
- 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](#)
- The ARP table might not be updated after VMotion or network loop is performed. [PR1521526](#)
- The l2ald process might generate core file when changing the EVPN-VXLAN configuration. [PR1541904](#)

Forwarding and Sampling

- The DHCP subscribers might get stuck in the **Terminated** state for around 5 minutes after disabling the cascade ports. [PR1505409](#)
- Traffic might get dropped for not exceeding the configured bandwidth under policer. [PR1511041](#)
- The srrd process might crash in a high route churns scenario or if the process flaps. [PR1517646](#)
- The commit might fail if a filter enabled with **enhanced-mode to et- interface** is configured. [PR1524836](#)
- The l2ald process might crash when a device configuration flaps frequently. [PR1529706](#)
- VLAN-id based firewall match conditions might not work for the VPLS service. [PR1542092](#)
- MAC learning issue might happen when EVPN-VXLAN is enabled. [PR1546631](#)
- All traffic would be dropped on the aggregated Ethernet interface bundle without 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](#)

General Routing

- New subscribers might fail to connect due to the following error message: **Filter index space exhausted**. [PR1531580](#)
- In some MX Series deployments running Junos OS, the following random syslog messages are observed for FPCs: **fpcx ppe_img_ucode_redistribute Failed to evict needed instr to GUMEM - xxx left**. These messages might not have a service impact. These messages are addressed as INFO level messages. On a Packet Forwarding Engine, there are dedicated UMEM and shared GUMEM memory blocks. This informational message indicates some evicting events between UMEM and GUMEM and can be safely ignored. [PR1298161](#)
- The **max-drop-flows** statement is not available. [PR1375466](#)
- 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 slips might be observed upon executing the **show route resolution extensive 0.0.0.0/0 | no-more** command if the number of routes in the system is large (several millions). [PR1425515](#)
- Application and removal of 1-Gbps speed results in the channel being down. [PR1456105](#)
- Random packet drop with flow cache is disabled, when NIC is mapped to NUMA node 1. [PR1458742](#)
- In the MVPN instance, the traffic drops on multicast receivers within the range of 0.1 to 0.9 percent. [PR1460471](#)
- The following error message is observed after GRES: **[user.err aftd-trio: [Error] IF:Unable to add member to aggregate member list, member already exists, agglfName:ps1.0 memberIfName:lt-3/0/0.32767]**. [PR1466531](#)

- Dynamic SR-TE tunnels do not get automatically recreated at the new primary Routing Engine after the Routing Engine switchover. [PR1474397](#)
- Expected number of 512,000 MAC entries are not re-learned in the bridge table after clearing 512,000 MAC entries from the table. [PR1475205](#)
- The syslog reports simultaneous zone change reporting for all green, yellow, orange, and red zones for one or more service PICs. [PR1475948](#)
- 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](#)
- Traffic decreases during throughput testing. [PR1483100](#)
- Prolonged flow control might occur with MS-MPC or MS-MIC. [PR1489942](#)
- The following error message is observed on the MPC card in the manual mode:
clksync_as_evaluate_synce_ref: 362 - Failed to configure clk. [PR1490138](#)
- The MX10003 router might shut itself down automatically after the system upgrades or downgrades. [PR1492121](#)
- VPLS flood next hop might not get programmed correctly. [PR1495925](#)
- Some of the virtual services might not come up after GRES or rpd restart. [PR1499655](#)
- As prefix is not emitted, the path emitted for **te-lsp-timers/state/cleanup** contains incorrect value. [PR1500690](#)
- The following error message is observed after deactivating the demux logical interface: **configuration check-out failed.** [PR1501002](#)
- The packets from a nonexisting source on the GRE or UDP designated tunnel might be accepted. [PR1503421](#)
- On the vMX instances, configuring the statement ranges for auto-sensed VLANs (either stacked VLANs or single-tag VLANs) might not work. This is because the VLANs are not programmed on the NIC drivers. [PR1503538](#)
- The gNMI stream does not follow the frequency on the subscription from the collector. [PR1504733](#)
- After sending the Layer 4 or Layer 7 traffic, the HTTP redirect messages are not captured as expected. [PR1505438](#)
- The heap memory utilization might increase after extensive subscriber login or logout. [PR1508291](#)
- Outbound SSH connection flap or memory leak issues might be observed during push configuration to the ephemeral database with a high rate. [PR1508324](#)
- The disabled QSFP transceiver might fail to get turned on. [PR1510994](#)

- PFCP message acknowledgment or non-acknowledgment responses are not tracked without the fix. If the CPF peer drops an acknowledged UPF response message and CPF retries the request, the reattempts do not get an acknowledgment by the response cache at UPF and get silently dropped. This causes the CPF state machine to constantly retry requests with that message being dropped at UPF, which leads to the **Established** state at both CPF and UPF. [PR1511708](#)
- Static subscribers are logged out after creating a unit under the demux0 interface. [PR1511745](#)
- Memory leak on l2ald might be observed 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](#)
- Subscribers might not be able to bind again after performing back-to-back GRES followed by an FPC restart. [PR1514154](#)
- Not able to forward traffic to VCP FPC after the MX Virtual Chassis reboots, FPC reboots, or adding VCP link. [PR1514583](#)
- The MACsec session might fail to establish if 256-bit cipher suite is configured for MACsec connectivity association assigned to a logical interface. [PR1514680](#)
- Duplicate prefix in secondary tunnel table is observed. [PR1514947](#)
- 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](#)
- yin2tlv sets unsupported command as hidden deprecated. [PR1516910](#)
- 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](#)
- During an upgrade, vSRX3.0 would display the following incorrect license warnings when utilizing licensable features even if the license was present on the device: **such as warning: requires 'idp-sig' license**. [PR1519672](#)
- The PADI packets might be dropped when the interface encapsulation VPLS is set along with 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 firmware for different hardware. [PR1524338](#)
- Commit is successful while deactivating CB0 or CB1 interfaces with GNF. [PR1524766](#)

- According to the OC data model, the **openconfig-alarms.yang** subscription path must be used as system, alarms, or an alarm. [PR1525180](#)
- Addition and removal of an aggregated Ethernet interface member link might cause the PPPoE subscriber session and traffic to drop. [PR1525585](#)
- 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 MPC10E line card might crash with the sensord process generating a core file due to a timing issue. [PR1526568](#)
- WAG control route prefix length is observed. [PR1526666](#)
- On the MX150 router, IFDs stay up during vmhost halt or power-off. [PR1526855](#)
- Commit error messages come twice while validating the **physical-cores** command. [PR1527322](#)
- The cp added process might generate the core file after upgrading to Junos OS Release 19.4 and later. [PR1527602](#)
- The transit PTP packet might be unexpectedly modified when passing through MPC2E-NG, MPC3E-NG, and MPC5E line cards. [PR1527612](#)
- Commit confirmed rollback does not work. [PR1527848](#)
- The l2cpd process might crash when removing LLDP on an aggregated Ethernet interface. [PR1528856](#)
- 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](#)
- Non-impacting error message is observed in the message logs: **IFP error>/..../..../..../src/pfe/usp/control/applications/interface/ifp.c@3270:(errno=1000) tunnel session add failed**. [PR1529224](#)
- The multicast traffic might get dropped due to hash mismatch when there are aggregated Ethernet and ECMP links involved in the multicast tree. [PR1529475](#)
- In the subscriber management environment, the RADIUS interim accounting records do not get populated with the subscriber statistics. [PR1529602](#)
- The SFP-LX or SFP-SX optics on MIC-3D-20GE-SFP-E/EH might show as unsupported after ISSU. [PR1529844](#)
- After performing ISSU with a high-scale bridge-domain configuration, less than 0.0254 percent of traffic loss is observed for a single bridge-domain interface. [PR1531051](#)
- On the MX10003 router, PEM 0 always shows as **Absent** or **Empty** even if PEM 0 is present. [PR1531190](#)
- Deleting the address of the jmgmt0 interface might fail if the shortened version of the CLI command is used. [PR1532642](#)
- VRRP synchronization does not occur in the backup Routing Engine with NSR in the **Steady** state. [PR1533357](#)

- The **clear ike statistics** command with remote gateway does not work. [PR1535321](#)
- Certain BGP SR-TE segment lists cause the rpd process to generate the core file during tunnel attribute parsing. [PR1535632](#)
- Multicast traffic might be observed even through unexpected interfaces with distributed IGMP is enabled. [PR1536149](#)
- Enhancements are needed for debugging l2ald. [PR1536530](#)
- The following error message might be observed when the JAM packages for the MX204, MX10003, and MX10008 are installed: **JAM: Plugin installed for summit_xxx PIC**. [PR1537389](#)
- Version-alias gets missed for subscribers configured with dynamic profiles after ISSU. [PR1537512](#)
- On the MPC10 and MPC11 line cards, the aftd process might crash in case of composite chain next hop creation failures. [PR1538559](#)
- The accounting interim-updates for subscriber does not work after GRES and subsequent reboot of FPCs in the node-slicing setup. [PR1539474](#)
- With hold time configuration, the ge Interfaces remain down on reboot. [PR1541382](#)
- Subscriber might not come up on some dynamic VLAN ranges in a subscriber management environment. [PR1541796](#)
- After changing addresses in the source pool, if the carrier-grade NAT traffic does not stop, the source pool cannot perform the NAT translation from the new pool. [PR1542202](#)
- Port mirroring with **maximum-packet-length** configuration does not work over the GRE interface. [PR1542500](#)
- The nsd daemon crashes after configuring the inline NAT44 in the USF mode. [PR1547647](#)
- The **verbose** command unexpectedly becomes hidden after Junos OS Release 16.1 for **set system export-format json**. [PR1547693](#)
- **SENSOR APP DWORD** leak is observed during the period of churn for routes bound to the sensor group. [PR1547698](#)
- Family IPv6 does not come up for the L2TP subscriber when the additional attributes are not passed in the **Framed-IPv6-Route** VSA. [PR1526934](#)
- The **show dynamic-profile session client-id** command displays only one IPv6 framed-route information. [PR1555476](#)
- The ERO update by the controller for branch LSP might cause issues. [PR1508412](#)
- The mspmand process leaks memory in relation to the MX telemetry reporting the following error message: **RLIMIT_DATA exceed**. [PR1540538](#)
- The mspmand process might generate core file on activating or deactivating the interface. [PR1544794](#)
- In the syslog output, the **sylog-local-tag** name is truncated (as **SYSLOG_SF**) when he **sylog-local-tag** name is configured as **SYSLOG_SFW**. [PR1547505](#)

- **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 PPPoE subscribers might fail to login. [PR1551207](#)
- The fabric errors are observed and the FPC processes might get offlined with SCBE3, MPC3E-NG, or MPC3E and MPC7 or MPC10 in the increased-bandwidth fabric mode. [PR1553641](#)
- The l2ald process might crash with traffic on the scaled set-up. [PR1517074](#)
- Difference between the port count and terminated count might be observed upon login or logout of the subscribers indicated by the output of the **show subscribers summary port extensive** command. [PR1523813](#)
- False positive TSensor errors are observed on vjunos0. [PR1508580](#)
- Snmp mib walk for jnxSubscriber OIDs returns the general error. [PR1535754](#)
- Delay in disabling the Packet Forwarding Engine might be seen on MPC7, MPC8, and MPC9 line cards. [PR1481879](#)
- The **next hop learning** statement is enabled by default in MPC10 and MPC11 line cards irrespective of the configuration statement. [PR1489121](#)
- The AMS bundle might remain inactive while adding the member interface to the AMS bundle with the scaled service sets. [PR1489607](#)
- Slow response might be observed when the **show | compare** or **commit check** action is executed in a large-scale configuration environment. [PR1500988](#)
- Sensord crashes on MPC10E line cards even when telemetry is not enabled. [PR1502260](#)
- Transit IPv4 traffic forwarding over BGP SR-TE might not work. [PR1505592](#)
- The l2cpd might crash if the ERP is deleted after the switchover. [PR1517458](#)
- The fxpc process might generate core file during EEPROM read when SFP is removed. [PR1518480](#)
- Traffic loss might occur when an uncorrected (Fatal) AER error is detected. [PR1519530](#)
- The VMXs might go to the amnesiac mode if they are deployed on the OpenStack based platforms. [PR1519668](#)
- The phc daemon might crash while committing the phone-home client configuration. [PR1522862](#)
- The BFD session status remains down at the non-anchor FPC even though the BFD session is up after the anchor FPC reboots or panics. [PR1523537](#)
- The rpd process might crash while restarting the routing gracefully with MPLS LSPs configured. [PR1527172](#)
- CFM does not consider the 8021AD configuration for the rewrite and classification tables. [PR1527303](#)
- BiDi 1G SFP optics displays wrong value in JVision for **optics/laser_rx_power_*_thresholds**. [PR1530120](#)

- The unilists are incorrectly formed and the list of the forwarded next-hops are not resolved properly if the value of the ECMP is set to 128. [PR1530803](#)
- The interface with the pic-mode 10GE configuration might not come up if upgraded to Junos OS Release 18.4R3-S4 or later. [PR1534281](#)
- Deactivating or activating PTP or SyncE in the upstream router causes the 100G links on the LC2103 to flap. [PR1538122](#)
- Traffic drop might be seen when executing the **request system reboot** command. [PR1538252](#)
- Upon receipt of a specific BGP FlowSpec message, network traffic might be disrupted. [PR1539109](#)
- The KRT queue might get stuck after the Routing Engine switchovers. [PR1542280](#)
- On the MX2010 and MX2020 devices, traffic loss might be observed when the Switch Fabric Board 3 and MPC8E 3D combination is used. [PR1544953](#)
- The Broadcom chip FPC might crash during the system booting. [PR1545455](#)
- Unexpected log messages appearing related to Neighbor Solicitation (NS) messages with multicast as source address is observed. [PR1546501](#)
- SR-TE might stay UP when the routes are deleted through policy. [PR1547933](#)
- The **LCM Peer Absent** error message might be seen on all TVP platforms. [PR1551760](#)
- ISSU might be aborted on MX devices for version Junos Release 20.2R2-S1. [PR1557413](#)

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 OIDs 1.3.6.1.2.1.47.2.0.30 are referenced in the SNMP trap after upgrading to Junos OS Release 18.4R3-S3. [PR1508281](#)
- SNMP polling might return an unexpectedly high value for the ifHCOutOctets counter for a physical interface when any jnxDom OID is processed at the same time. [PR1508442](#)
- The kernel might crash if a file or directory is accessed for the first time and is not created locally. [PR1518898](#)
- The telnetd.real local privilege escalation vulnerabilities in SUID binaries is observed. [PR1525318](#)
- The output drops in the **show interfaces extensive** command ' might display 0 temporarily during a race condition when the SNMP query for JnxCos is also issued. [PR1533314](#)

Interfaces and Chassis

- The **sonet-options configuration** statement is disabled for the xe interface that works in the wan-phy mode. [PR1472439](#)
- Fail to configure proactive ARP detection. [PR1476199](#)
- The fpc process might crash in an inline mode with CFM configured. [PR1500048](#)
- The following error message is observed: **Request failed: OID not increasing: ieee8021CfmStackServiceSelectorType**. [PR1517046](#)
- Buffer overflow vulnerability in a device control daemon is observed. [PR1519334](#)
- 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 an 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](#)
- 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](#)
- The **startup-silent-period** command might not work in Junos OS Release 20.3R1 or later. [PR1548464](#)
- The dcd process might leak memory on pushing the configuration to the ephemeral database. [PR1553148](#)
- Distribution fails for few sessions when VRRP is configured in a large-scale with active-inherit scenario. [PR1505998](#)
- Backup router generates **VRRP_NEW_BACKUP** syslog during bring up. [PR1539277](#)
- The rpd memory leak might be observed on the backup Routing Engine due to flapping of the link. [PR1539601](#)

Intrusion Detection and Prevention (IDP)

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

Juniper Extension Toolkit (JET)

- The JET application configuration must be disabled before upgrading Junos OS vmhost images. [PR1488769](#)

J-Web

- Privilege escalation in J-Web is observed due to arbitrary command and code execution through information disclosure from another users active session. [PR1518212](#)

Layer 2 Ethernet Services

- 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 **DHCPLEASEUNASSIGNED**, which is a 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 DHCPv6 relay statistics. [PR1521227](#)
- Memory leak in the jdhcpd process might be seen if access-profile is configured under the **dhcp-relay** or **dhcp-local-server** statement. [PR1525052](#)

Layer 2 Features

- The rpd process might crash on the new primary Routing Engine after GRES in the VPLS or Layer 2 circuit scenario. [PR1507772](#)
- The host generated traffic might get lost as the current forwarding member nexthop is down while another member nexthop is up. [PR1516514](#)

MPLS

- Committing might trigger externally provisioned LSP MBB mechanism. [PR1546824](#)
- A same device responds twice for traceroute in case of the device going through an MPLS network under specific conditions. [PR1494665](#)

- Traffic loss might occur if ISSU is performed when P2MP is configured for an LSP. [PR1500615](#)
- The rpd scheduler might slip after the link flaps. [PR1516657](#)
- The rpd process might crash after upgrading Junos OS Release 18.1 to a later release. [PR1517018](#)
- The SNMP trap is sent with the incorrect OID jnxSpSvcSetZoneEntered. [PR1517667](#)
- The LDP session-group might throw a commit error and flap. [PR1521698](#)
- The inter-domain LSP with loose next hops path might get stuck in the **Down** state. [PR1524736](#)
- The **ping mpls rsvp** command does not take into account lower MTU in the path. [PR1530382](#)
- The rpd process might crash when the LDP route with indirect next hop is deleted on the aggregated Ethernet interface. [PR1538124](#)
- The LDP routes might be deleted from the MPLS routing table after the Routing Engine switchover. [PR1527197](#)
- The rpd process might crash during the restart routing when the MPLS LSPs are present. [PR1530213](#)

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

- Unable to poll dot1dStp objects with l2cpd registered context after l2cpd restarts. [PR1561736](#)

Platform and Infrastructure

- With multiple different fixed-sized traffic streams configured at 10,000,00 fps (40-Gbps combined rate) on aggregated Ethernet0 along with another independent aggregated Ethernet interface (aggregated Ethernet1, 50 percent line rate 4 streams bidirectional => 118-Gbps combined traffic rate), both hosted on a single Packet Forwarding Engine instruction of the MPC11E line card, small varying packet drops occur for every iteration on aggregated Ethernet1 on disabling aggregated Ethernet0. [PR1464549](#)
- Traffic to VRRP virtual IP or MAC addresses might be dropped when ingress queuing is enabled. [PR1501014](#)
- Traffic originated from another subnet is sent out with 0x8100 instead of 0x88a8. [PR1502867](#)
- The kernel might crash causing the router or the Routing Engine to reboot when performing virtual IP related change. [PR1511833](#)
- The output of the **show jnh qmon queues-sensor stats 0** command has no content. [PR1514881](#)

- The VPLS connection might be stuck in the **Primary Fail** status when a dynamic profile is used on the VPLS pseudowire logical interface. [PR1516418](#)
- Configured scheduler-map is not applied on the ms- interface if the service PIC is in the **Offline** state during commit. [PR1523881](#)
- Flow programming issue for lt- 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](#)
- There is a TWAMP interoperability issue between Junos OS releases. [PR1533025](#)
- The fpc process might crash when the next hop memory of ASIC is exhausted in the EVPN-MPLS scenario. [PR1533857](#)
- 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 the LSI interface. [PR1542211](#)
- Subscribers does are not come up on VPLS PS interface. [PR1536043](#)
- The rmopd process memory leak might be seen if TWAMP client is configured. [PR1541808](#)
- The PE and CE devices OAM CFM might have issues in the aggregated Ethernet interface. [PR1501656](#)
- The VXLAN encapsulation over IPv6 underlay might not work. [PR1532144](#)
- The ISSU might fail on Junos platforms with the LUCHIP based line cards. [PR1535745](#)
- Dynamic filter fails to match IPv6 prefix. [PR1536100](#)
- TWAMP interoperability issue are observed. [PR1536939](#)
- The ARP expired timer on the backup Routing Engine is not the same as the primary Routing Engine if the aging-timer is configured. [PR1544398](#)

Routing Policy and Firewall Filters

- For setting the IPv6 router ID, the **routing-options** statement is added. [PR1523283](#)
- The policy configuration might be mismatched between the rpd and mgd process when **deactivate policy-options prefix-list** is involved in the configuration sequence. [PR1523891](#)

Routing Protocols

- The output of the **show isis interface detail** command might be incorrect if **wide-metrics-only** is enabled for IS-IS and the ASCII representation of the metric in decimal is more than 6 characters long. [PR1482983](#)
- The BGP RPKI ROA withdrawal might lead to an unexpected BGP route flap. [PR1483097](#)

- The rpd process might report 100 percent CPU usage with the BGP route damping enabled. [PR1514635](#)
- The IS-IS SR routes might not be updated to reflect the change in the SRMS advertisements. [PR1514867](#)
- The BGP link-bw of the non-multipath routes are included in an aggregation. [PR1515264](#)
- The rpd process might crash after deleting and then adding a BGP neighbor. [PR1517498](#)
- The rpd process might crash if there is a huge number of SA messages in an MSDP scenario. [PR1517910](#)
- Tag matching in the VRF policy does not work properly when the independent-domain option is configured. [PR1518056](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- NLRI handling improvements for BGP-LS ID TLV is needed. [PR1521258](#)
- BFD with authentication for BGP flaps after GRES or NSR switchover on the NG-RE and SCBE2 setup. [PR1522261](#)
- The IS-IS LSP database synchronization issue might be observed while using the flood-group feature. [PR1526447](#)
- The rpd process generates core file at `is_srv6_delete_locator_end_sid_data isis_srv6_end_sid_local_data_delete isis_srv6_locator_config_check`. [PR1531830](#)
- Transit labels for Layer 3 VPN routes are pushed momentarily to the MPLS.0 table. [PR1532414](#)
- Configuring **then next hop** and **then reject** on a route policy for the same route might cause the rpd process to crash. [PR1538491](#)
- After moving peer out of the protection group, the path protection does not get removed from the PE router. Multipath routes are still present. [PR1538956](#)
- The rpd process generates the core file at `gp_rtargt_tsi_update,bgp_rtargt_flash_rt,bgp_rtargt_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](#)
- IS-IS does not call `ted_add_halfink` for P2P IPv6-only links for traffic engineering topology. [PR1548506](#)
- Telemetry key value for transport or remote-address field for link-local IPv6 peer is incorrect and logical interface is absent. [PR1548754](#)
- The BGP session with VRRP virtual address might not come up after a flap. [PR1523075](#)
- The VRF label is not assigned at ASBR when the inter AS is implemented. [PR1523896](#)
- The BGP session neighbor shutdown configuration does not effect the non-established peer. [PR1554569](#)
- The BGP session might not come up if **extended-nexthop** is enabled by default on the other vendor remote peer. [PR1555288](#)
- The rpd process might crash with BGP RPKI enabled in a race condition. [PR1487486](#)

- The ppmmd process might generate core file after FPC restarts. [PR1490918](#)
- The virtual-router option is not supported under a routing-instance in a lean rpd image. [PR1494029](#)
- Traffic loss might be seen in the next-hop-based dynamic tunnels of Layer 3 VPN scenario after changing the dynamic-tunnel preference. [PR1542123](#)
- Six PE device prefixes might not be removed from RIB upon reception of withdrawal from a BGP neighbor when RIB sharding is enabled. [PR1556271](#)
- BGP routes might be stuck in routing table in the **Accepted DeletePending** state when the BGP peering session goes down. [PR1562090](#)

Services Applications

- The following error message is observed: **SPD_CONN_OPEN_FAILURE: spd_pre_fetch_query: unable to open connection to si-1/0/0**. [PR1550035](#)

Subscriber Access Management

- Subscriber accounting message retransmissions exist even after configuring the accounting retry 0. [PR1405855](#)
- CCR-T does not contain the usage-monitoring information. [PR1517507](#)
- The **show network-access aaa subscribers statistics username "<>"** command fails to fetch the subscriber-specific AAA statistics information if a subscriber username contains a space. [PR1518016](#)

User Interface and Configuration

- NETCONF service over SSH might not work on the device that runs Junos OS if in-band management is used. [PR1517160](#)
- The command injection vulnerability in the **request system software** command is observed. [PR1519337](#)
- The dexp local privilege escalation vulnerabilities in SUID binaries is observed. [PR1529210](#)

VPNs

- The MPLS label manager might allow configuration of a duplicated VPLS static label. [PR1503282](#)
- The rpd process might crash after removing the last configured interface under the Layer 2 circuit neighbor. [PR1511783](#)

- MVPN multicast route entry might not be properly updated with the actual downstream interfaces list. [PR1546739](#)
- The Junos image upgrade or installation with **validate** fails with XML errors. [PR1525862](#)

SEE ALSO

What's New	 90
What's Changed	 106
Known Limitations	 112
Open Issues	 113
Documentation Updates	 135
Migration, Upgrade, and Downgrade Instructions	 135

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for MX Series routers.

SEE ALSO

What's New	 90
What's Changed	 106
Known Limitations	 112
Open Issues	 113
Resolved Issues	 120
Migration, Upgrade, and Downgrade Instructions	 135

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.4R1](#) | 136
- [Procedure to Upgrade to FreeBSD 11.x-Based Junos OS](#) | 137

- [Procedure to Upgrade to FreeBSD 6.x-Based Junos OS | 139](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 141](#)
- [Upgrading a Router with Redundant Routing Engines | 141](#)
- [Downgrading from Release 20.4R1 | 142](#)

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 20.4R1

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-x86-32-20.4R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.4R1.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-20.4R1.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-20.4R1.9-limited.tgz
```

Replace source with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE:

- You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the `junos-vmhost-install-x.tgz` image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the [Installation and Upgrade Guide](#).
- Starting in Junos OS Release 20.4R1, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following MX Series routers:
 - MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008

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

NOTE: After you install a Junos OS Release 20.4R1 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-20.4R1.9-signed.tgz
```

- Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

```
user@host> request system software add validate reboot  
source/jinstall-ppc-20.4R1.9-limited-signed.tgz
```

Replace source with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 20.4R1 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Upgrading a Router with Redundant Routing Engines

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

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

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

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

Downgrading from Release 20.4R1

To downgrade from Release 20.4R1 to another supported release, follow the procedure for upgrading, but replace the 20.4R1 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 | 90](#)
- [What's Changed | 106](#)
- [Known Limitations | 112](#)
- [Open Issues | 113](#)
- [Resolved Issues | 120](#)
- [Documentation Updates | 135](#)

Junos OS Release Notes for NFX Series

IN THIS SECTION

- [What's New | 143](#)
- [What's Changed | 146](#)
- [Known Limitations | 147](#)
- [Open Issues | 148](#)
- [Resolved Issues | 149](#)

- Documentation Updates | 150
- Migration, Upgrade, and Downgrade Instructions | 151

These release notes accompany Junos OS Release 20.4R1 for the NFX Series Network Services Platforms. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Application Security | 143
- High Availability | 145
- Flow-Based and Packet-Based Processing | 145
- Logical Systems and Tenant Systems | 145
- Routing Protocols | 145
- Security | 146

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

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

Application Security

- **AppQoE support for SaaS applications (NFX Series and SRX Series)**—Starting in Junos OS Release 20.4R1, we've extended application quality of experience (AppQoE) support for Software as a Service (SaaS) applications.

AppQoE performs service-level agreement (SLA) measurements across the available WAN -links such as underlay, GRE, IPsec or MPLS over GRE. It then sends SaaS application data over the most SLA-compliant link to provide a consistent service.

To configure AppQoE for SaaS applications:

1. Define the SLA rule type as SaaS (**set security advance-policy-based-routing sla-rule sla1 type saas**).
2. Include SaaS server details in the address book (**set security address-book global address *address-book-name* dns-name *saas-server-url* ipv4-only**).
3. Disable midstream switching to disengage advanced policy-based routing (APBR) and prevent further rule matching.
4. Attach the SLA rule to the policy-based APBR profile.

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

- **Granular control over DNS-over-HTTP and DNS-over-TLS application traffic (NFX Series and SRX Series)**—In Junos OS Release 20.4R1, we introduce a new micro-application, DNS-ENCRYPTED, to enhance the application signature package. By configuring this micro-application in a security policy, you can have granular control for DNS-over-HTTP and DNS-over-TLS application traffic.

The DNS-ENCRYPTED application is enabled by default. You can disable it using the **request services application-identification application disable DNS-ENCRYPTED** command.

You can view the details of the micro-applications using the **show services show services application-identification application detail** command.

[See [Application Identification Support for Micro-Applications](#).]

High Availability

- **High availability on NFX350 devices**—Starting in Junos OS Release 20.4R1, NFX350 devices support the Chassis Cluster feature. You can configure a cluster of two NFX350 devices in active/passive or active/active mode 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.

[See [Chassis Cluster on NFX350 Devices](#) and [Upgrading or Disabling a Chassis Cluster on NFX350 Devices](#).]

Flow-Based and Packet-Based Processing

- **Support of IPFIX formatting and Chassis Cluster for NFX J-Flow functionality (NFX150, NFX250 NextGen, and NFX350)**—Starting with Junos OS Release 20.4R1, you can configure Chassis Cluster and define an IPFIX flow record template suitable for IPv4 traffic or IPv6 traffic. IPFIX is an enhanced version of J-flow version 9 template. Using IPFIX, you can collect a set of sampled flows and send the record to a specified host.

See [[Configuring Inline J-Flow to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices](#), [show services accounting flow](#), [show services accounting errors](#), [show services accounting status](#), and [clear services accounting statistics inline-jflow](#).]

Logical Systems and Tenant Systems

- **Support for MAP-E confidentiality CLI statement (NFX150, NFX250, NFX350, and SRX1500)**—Starting in Junos OS Release 20.4R1, we've introduced a global MAP-E **confidentiality** CLI statement to hide MAP-E rule parameters in CLI show commands and logs. To enable this configuration, include the **confidentiality** statement at the **[edit security softwires map-e]** hierarchy level. You need to have administrator privileges to enable or disable this configuration. This feature is supported for all domains of MAP-E.

[See [confidentiality](#) and [show security softwires map-e confidentiality status](#).]

Routing Protocols

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC6286 (MX204, NFX Series, PTX5000, QFX Series, and vRR)**—Starting in Junos OS Release 20.4R1, you can establish a BGP connection using a BGP identifier that is a 4-octet, unsigned, nonzero integer and it needs to be unique only within the autonomous system (AS) per RFC 6286. In earlier releases, the BGP ID of a BGP speaker was required to be a valid IPv4 host address assigned to the BGP speaker.

To enable this feature, use the **bgp-identifier identifier group bgp group name bgp-identifier identifier neighbor peer address bgp-identifier identifier** configuration statement at the [edit protocols bgp] hierarchy level.

[See [router-id](#)]

Security

- **MACsec on NFX350 devices**—Starting in Junos OS Release 20.4R1, you can configure Media Access Control Security (MACsec) on NFX350 devices for secure communication for almost all types of traffic on Ethernet links. MACsec provides point-to-point security on Ethernet links between directly connected nodes and is capable of identifying and preventing most security threats, including denial of service, intrusion, man-in-the-middle, masquerading, passive wiretapping, and playback attacks.

[See [Configuring MACsec on NFX350 Devices](#) and [macsec](#).]

SEE ALSO

What's Changed 146
Known Limitations 147
Open Issues 148
Resolved Issues 149
Documentation Updates 150
Migration, Upgrade, and Downgrade Instructions 151

What's Changed

IN THIS SECTION

- [Junos OS XML API and Scripting | 147](#)

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

Junos OS XML API and Scripting

- The `<get-interface-information/>` RPC reply includes an `<error-severity>` element when execution fails (NFX Series)—If the `<get-interface-information/>` RPC fails to execute, the device's RPC reply includes the `<error-severity>` element. In earlier releases, the RPC reply does not include the `<error-severity>` element.

SEE ALSO

What's New 143
Known Limitations 147
Open Issues 148
Resolved Issues 149
Documentation Updates 150
Migration, Upgrade, and Downgrade Instructions 151

Known Limitations

IN THIS SECTION

- [Interfaces | 148](#)

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

Interfaces

- The link disable option puts the analyzer interface in an inconsistent state, with link state as DOWN and administrator state as UP. [PR1442224](#)

SEE ALSO

What's New 143
What's Changed 146
Open Issues 148
Resolved Issues 149
Documentation Updates 150
Migration, Upgrade, and Downgrade Instructions 151

Open Issues

IN THIS SECTION

- [Interfaces | 149](#)
- [Platform and Infrastructure | 149](#)
- [Virtual Network Functions \(VNFs\) | 149](#)

Learn about open issues in this release for NFX Series devices.

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

Interfaces

- When you issue a **show interface** command on NFX150 devices to check the interface details, the system does not check whether the interface name provided is valid or invalid. The system does not generate an error message if the interface name is invalid. [PR1306191](#)

Platform and Infrastructure

- Jumbo frames are not supported through OVS on an NFX250 device. [PR1420630](#)

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 | 143](#)

[What's Changed | 146](#)

[Known Limitations | 147](#)

[Resolved Issues | 149](#)

[Documentation Updates | 150](#)

[Migration, Upgrade, and Downgrade Instructions | 151](#)

Resolved Issues

IN THIS SECTION

- [High Availability | 150](#)
- [Interfaces | 150](#)
- [Platform and Infrastructure | 150](#)

Learn which issues were resolved in the Junos OS Release 20.4R1 for NFX Series devices.

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

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

- 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, after a system reboot, this functionality stops working and no packets are mirrored on the output analyzer port. [PR1480290](#)

Platform and Infrastructure

- On NFX150 devices, ZTP over LTE configuration commit fails for **operation=create** in XML operations configuration. [PR1511306](#)
- The device reads the board ID from eeprom directly using I2C upon power cycle. [PR1529667](#)

SEE ALSO

[What's New | 143](#)

[What's Changed | 146](#)

[Known Limitations | 147](#)

[Open Issues | 148](#)

[Documentation Updates | 150](#)

[Migration, Upgrade, and Downgrade Instructions | 151](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for NFX Series devices.

SEE ALSO

[What's New | 143](#)[What's Changed | 146](#)[Known Limitations | 147](#)[Open Issues | 148](#)[Resolved Issues | 149](#)[Migration, Upgrade, and Downgrade Instructions | 151](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrade and Downgrade Support Policy for Junos OS Releases | 151](#)
- [Basic Procedure for Upgrading to Release 20.4 | 152](#)

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 20.4

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

NOTE: 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 20.4R1:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<https://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the **Software** tab.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the Download Software page.
5. In the Install Package section of the Software tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the device or to your internal software distribution site.
10. Install the new package on the device.

SEE ALSO

[What's New | 143](#)

[What's Changed | 146](#)

[Known Limitations | 147](#)

[Open Issues | 148](#)

[Resolved Issues | 149](#)

[Documentation Updates | 150](#)

Junos OS Release Notes for PTX Series

IN THIS SECTION

- [What's New | 154](#)
- [What's Changed | 162](#)
- [Known Limitations | 164](#)
- [Open Issues | 165](#)
- [Resolved Issues | 168](#)
- [Documentation Updates | 170](#)
- [Migration, Upgrade, and Downgrade Instructions | 171](#)

These release notes accompany Junos OS Release 20.4R1 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Junos OS XML, API, and Scripting | 154
- Junos Telemetry Interface | 155
- MPLS | 157
- Network Management and Monitoring | 158
- Routing Policy and Firewall Filters | 158
- Routing Protocols | 159
- Software Installation and Upgrade | 161
- System Logging | 161

Learn about new features introduced in Junos OS Release 20.4R1 for the PTX Series.

Junos OS XML, API, and Scripting

- **Support for Certificate Authority Chain Profile (EX2300, EX3400, EX4300, MX240, MX480, MX960, PTX-5000, VMX, vSRX and QFX5200)**—Starting in Junos OS Release 20.4R1, you can configure intermediate Certificate Authority (CA) chain profile certificate and perform https REST API request using mutual and server authentications.

To configure intermediate ca-chain certificate, configure **ca-chain ca-chain** statement at the [edit system services rest https] hierarchy level.

- **Start time option for interval-based internal events that trigger event policies (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.4R1, when you create an interval-based internal event for triggering event policies, you can specify the start date and time for the initial event. To specify a start time, configure the **start-time** option along with the **time-interval** option at the [edit event-options generate-event] hierarchy level.

[See [Generating Internal Events to Trigger Event Policies](#).]

Junos Telemetry Interface

- **JTI support for inline Junos Traffic Vision sensors with gRPC services (MX Series and PTX Series)**—Junos OS Release 20.4R1 supports inline Jflow sensors for FPC3 and MPC 1 through 9. This feature enables you to monitor inline Junos Traffic Vision (previously known as Jflow) service statistics on a router and to export statistics to an outside collector at configurable intervals using remote procedure call (gRPC) services.

Use the resource path `/junos/system/linecard/services/inline-jflow/` in a subscription to export statistics.

You can view statistics in the collector output under `/components/`. The collector component ID in the statistics output will include the FPC slot number for which inline Junos Traffic Vision statistics are exported. For example, inline Jflow statistics for FPC 0 will be under **component id 0**, and inline Jflow statistics for FPC 1 will be under **component id 1**.

Inline Junos Traffic Vision statistics are slightly different, depending on the routing platform.

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

- **Enhancing debug information for JTI (PTX5000)**—Starting in Junos OS Release 20.4R1, debug commands supporting Junos telemetry interface (JTI) are enhanced to better support JTI.

The **show network-agent statistics (brief | detail)** command output now includes:

- Average latency values per sensor, which helps to check the latency of any given sensor on the device.
- Approximate circular buffer usage per sensor, which provides an early alert if drops are likely for any specific sensor.
- Time of subscription, which helps to correlate statistics information from the provisioning logs that are taken over a period of multiple subscriptions.

The **show extension-service request-response clients (brief | detail)** command output now includes:

- The username for which the session was authenticated in a remote procedure call (gRPC) session. If not authenticated, the **username** field displays as no authentication. This helps to identify which users have requested programmable operations.
- Login time of the gRPC client, which helps determine how long this client has been active.

[See [show network-agent statistics](#) and [show extension-service request-response clients](#).]

- **JTI support for persistent active gRPC sessions between collector and server during an SSL certificate update (ACX Series, MX Series, and PTX Series)**—Junos OS Release 20.4R1 supports persistent active remote procedure call (gRPC) sessions between the collector (client) and server during an SSL certificate update.

For secure channel authentication, the TLS protocol is used to maintain a secure channel between the collector and the server. TLS uses the server certificate and the client certificate to authenticate each other and send encrypted messages over the network. When an SSL certificate is updated, existing gRPC

sessions are abruptly terminated, forcing the collector to initiate a new gRPC connection and subscribe to sensors again.

To avoid this problem, you can enable persistent active gRPC sessions by configuring **hot-reloading** at the `[edit system services extension-service request-response grpc ssl]` hierarchy level. After you enable this feature, gRPC sessions will remain active even when authentication certificates are updated.

After the certificate is updated, any new gRPC session will use the updated certificate.

[See [gRPC Services for Junos Telemetry Interface](#) and [ssl](#).]

- **BGP neighbor telemetry with sharding (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, BGP neighbor telemetry with sharding (multi-threading) is supported.

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

- **LACP sensors for actor partner states on JTI (MX Series and PTX Series)**—Starting in Junos OS Release 20.4R1, you can use Junos telemetry interface (JTI) and gRPC Network Management Interface (gNMI) services to export LACP actor partner states (also known as LACP port states). When a subscription is configured, ON_CHANGE or periodic streaming statistics are sent from devices to an outside collector.

You can subscribe to `/lacpd/` to collect all statistics or include the following resource paths individually in a subscription:

- `/lacpd/ae/member/partner_collecting`
- `/lacpd/ae/member/partner_synchronization`
- `/lacpd/ae/member/partner_timeout`
- `/lacpd/ae/member/partner_aggregatable`
- `/lacpd/ae/member/partner_distributing`
- `/junos/system/linecard/interface/traffic/`

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

- **Juniper Resiliency Interface for exception reporting and null route detection (ACX Series, PTX Series, and MX Series)**—Starting in Junos OS Release 20.4R1, you can use Juniper Resiliency Interface to detect and reduce Mean Time to Repair (MTTR) first-order network issues. Juniper Resiliency Interface uses a push model for data reporting from the entities in the system which encounter packet drops. This automates the workflow for detecting, reporting, and mitigating adverse exceptions.

To collect kernel routing table and routing protocol process exceptions, configure the **set system resiliency exceptions** statement at the `[edit]` hierarchy level to specify exception reporting based on kernel exceptions, and routing exceptions.

You can display exceptions from a remote collector by means of remote procedure call (gRPC) services or gRPC network management interface (gNMI) services. Display on-box exceptions by accessing the `/var/log` file or the database at `/var/db/ResiliencyExceptions.db`. No Junos operational mode commands display these exceptions.

MPLS

- **Support for optimizing auto-bandwidth adjustments for MPLS LSPs (MX Series and PTX Series)**—Starting in Junos OS Release 20.4R1, you can configure faster auto-bandwidth adjustment for MPLS LSPs under overflow or underflow conditions. This feature decreases the minimum allowed **adjust-threshold-overflow-limit** and **adjust-interval** to 150 seconds when **adjust-threshold-overflow-limit** and **adjust-threshold-underflow-limit** cross the configured threshold values. In releases earlier than Junos OS Evolved Release 20.4R1, the **adjust-interval** is 300 seconds under overflow or underflow conditions.

You can configure faster in-place LSP bandwidth update that avoids signaling of a new LSP instance as part of make-before-break. To configure faster in-place LSP bandwidth update, include the **in-place-lsp-bandwidth-update** configuration statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level.

You can also configure RSVP interfaces to support subscription percentage per priority. To configure subscription percentage per priority, include the **subscription priority *priority* percent *value*** configuration statement at the **[edit protocols rsvp interface *interface-name*]** hierarchy level.

[See [Configuring Optimized Auto-bandwidth Adjustments for MPLS LSPs](#).]

- **Re-engineering of SR-TE (MX Series, PTX Series)**—Starting with Junos OS Release 20.4R1, you can incorporate the following features to enhance the debugging capability of segment routing traffic-engineering (SR-TE):
 - rib-group import functionality.
 - Display of SR-TE routes installed from various tunnel sources using the **show spring-traffic-engineering** command.
 - Template map for BGP SR-TE tunnels.
 - Compute profile in template with distributed Constrained Shortest Path First (CSPF) for dynamic SR-TE tunnels.
 - 6PE (IPv6 over IPv4 SR-TE tunnel)
 - no-chained-composite-next-hop option

[See [source-packet-routing](#) and [show spring-traffic-engineering](#).]

- **Support for express segments to establish end-to-end segment routing path (MX Series and PTX Series)**—Starting in Junos OS Release 20.4R1, express segments can be used to establish end-to-end TE paths between interconnected TE networks. Express segments (also known as virtual TE links) are generated dynamically through policies matching the underlay LSPs. Express segments and the corresponding abstracted topology (required by RFC7926) is generated with policies.

To apply a policy, include the **policy *policy-name*** statement at the **[edit protocols express-segment traffic-engineering]** hierarchy level.

To configure express segment, include the **express-segment** statement under the [edit protocols] hierarchy level.

[See *How to Establish End-to-End Segment Routing Paths Using Express Segments*.]

Network Management and Monitoring

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The **<get-configuration>** operation supports the **compare="configuration-revision"** and **configuration-revision** attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

Routing Policy and Firewall Filters

- **Support for route's next-hop weight in policy match condition (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, a route with multiple next-hop paths can use the weight associated with a path to identify primary and backup paths. The path with the lowest weight is used as the primary path, and any paths with higher weights are treated as backup paths. You can use the next-hop weight as a match condition in export policies to redistribute IGP and BGP routes based on whether the primary or backup paths are active.

Configure this match condition using the [edit policy-options policy-statement *policy-name* term *term-name* from] statement.

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

- **Unicast RPF support for IPv4 and IPv6 (PTX10004)**—Starting in Junos OS Release 20.4R1, PTX10004 devices support unicast reverse-path-forwarding (uRPF) for both IPv4 and IPv6 traffic flows. uRPF helps protect against DoS and DDoS attacks by verifying the unicast source address of packets arriving on a protected interface. Packets that are not from a valid path can be discarded. You can enable RPF checking

for a given interface from the `[edit interfaces name unit number family inet | inet6 rpf-check]` hierarchy level, and create a discard rule at the `[edit firewall filter name term default then reject]` hierarchy level.

[See [Example: Configuring Unicast Reverse-Path-Forwarding Check](#).]

Routing Protocols

- **Support for multiple single-hop EBGP sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGP session. You can now enable single-hop EBGP sessions on different links over multiple directly connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGP Sessions on Different Links Using the Same Link-Local Address \(IPv6\)](#).]

- **Support for IS-IS flood-reflector interfaces (PTX1000, QFX10002, QFX10008)**—Starting in Junos OS Release 20.4R1, we support the IS-IS flood reflector feature that offers better scalability for a Level 2 topology. Flood reflectors enable the creation of topologies where Level 1 areas provide transit forwarding for Level 2 destinations within a Level 2 topology.

The flexible tunnel interfaces (FTI) are designated as flood-reflector interfaces. To enable the flood reflector on an FTI, include the **flood-reflector** statement at the `[edit protocols isis interface interface name level level number]` hierarchy level.

You can configure the interface to be either the reflector or the client. To enable the reflector, you can use the **flood-reflector reflector *cluster-id*** statement at the `[edit protocols isis level level number]` hierarchy level.

To enable the flood reflector client, include the **flood-reflector client** statement at the `[edit protocols isis level level number]` hierarchy level.

NOTE: You can configure the flood reflector feature on FTIs at Level 2 only.

[See [How to Configure Flood-Reflector Interfaces in IS-IS Networks](#).]

- **Support for BGP Labeled Unicast prefix SID (MX Series and PTX Series)**—Starting in Junos OS 20.4R1, BGP labeled unicast can carry segment routing global block label range and index information through the prefix segment attribute. With this feature we support segment routing using the BGP labeled unicast prefix segments and the MPLS data plane in medium to large scaled data centers. The controller directs the server to assign a stack- of labels to an incoming packet based on the available network state

information. The assigned label stack avoids congested paths and steers the packet through a best available path.

To configure and advertise the SRGB label range specifically for BGP include the **source-packet-routing srgb start-label start-label index-range index-range** and **advertise-srgb** configuration statements at the **[edit protocols bgp]** hierarchy level.

To advertise prefix SIDs to external BGP peers, include the **advertise-prefix-sid** configuration statement at the **[edit protocols bgp]** hierarchy level. You can configure this statement globally or for specific BGP groups or BGP neighbors.

[See [srgb](#).]

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC6286 (MX204, NFX Series, PTX5000, QFX Series, and vRR)**—Starting in Junos OS Release 20.4R1, you can establish a BGP connection using a BGP identifier that is a 4-octet, unsigned, nonzero integer and it needs to be unique only within the autonomous system (AS) per RFC 6286. In earlier releases, the BGP ID of a BGP speaker was required to be a valid IPv4 host address assigned to the BGP speaker.

To enable this feature, use the **bgp-identifier identifier group bgp group name bgp-identifier identifier neighbor peer address bgp-identifier identifier** configuration statement at the **[edit protocols bgp]** hierarchy level.

[See [router-id](#).]

- **IPv6 support in TED (MX Series, PTX Series)**—Starting in Junos OS Release 20.4R1, you can configure IS-IS traffic engineering to store IPv6 information in the traffic engineering database (TED) in addition to IPv4 addresses. BGP-LS distributes this information as routes from the TED to the Isdist.0 routing table using the TED import policies. These routes are advertised to BGP-TE peers as network layer reachability information (NLRI) with IPv6 router ID type, length, and value (TLV).

With this enhancement, you can benefit from obtaining the complete network topology in the TED.

[See [Link-State Distribution Using BGP Overview](#).]

Software Installation and Upgrade

- **ZTP with DHCPv6 client support (EX3400, EX4300, PTX1000, PTX5000, PTX10002-60C, PTX10008, QFX5100, QFX5200, QFX10002, and QFX10002-60C)**—Starting in Junos OS Release 20.4R1, zero touch 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 one of the DHCPv4 bindings fails, the device continues to check for bindings until provisioning is successful. However, if there are no DHCPv4 bindings, the device checks for DHCPv6 bindings and follows 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: ZTP supports only HTTP and HTTPS transport protocols.

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

System Logging

- **Support for time averaged watermark (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, you can capture steady state data of routing and forwarding (RIB/FIB) table routes using the **time-averaged-watermark-interval** configuration statement at the **[edit routing-options]** hierarchy level. Time averaged watermark is calculated whenever the time averaged interval is changed from CLI. Time averaged watermark is logged in syslog if the logs are enabled in the system at **LOG_NOTICE** level. The default time averaged watermark interval is 1 day. You can see the timed averaged watermark using the existing **show route summary** command.

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

SEE ALSO

[What's Changed | 162](#)

[Known Limitations | 164](#)

[Open Issues | 165](#)

[Resolved Issues | 168](#)

[Documentation Updates | 170](#)

[Migration, Upgrade, and Downgrade Instructions | 171](#)

What's Changed

IN THIS SECTION

- Class of Service (CoS) | 162
- General Routing | 162
- MPLS | 163
- Network Management and Monitoring | 163
- User Interface and Configuration | 163

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 20.4R1 for the PTX Series.

Class of Service (CoS)

- We've corrected the output of the **show class-of-service interface | display xml** command. The output is 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

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial of service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

MPLS

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

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to deviate not-supported nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.

User Interface and Configuration

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

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

SEE ALSO

[What's New | 154](#)

[Known Limitations | 164](#)

[Open Issues | 165](#)

[Resolved Issues | 168](#)

[Documentation Updates | 170](#)

[Migration, Upgrade, and Downgrade Instructions | 171](#)

Known Limitations

IN THIS SECTION

- [General Routing | 164](#)
- [Routing Protocols | 164](#)

Learn about known limitations in Junos OS Release 20.4R1 for the PTX Series.

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

General Routing

- During reconfiguration 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](#)
- When counter sample is enabled, it attempts to fetch the physical interface statistics for sFlow-enabled interfaces using rtsock messages to kernel. This blocks call and wait for the reply of earlier request and sends a new request only after receiving the reply of first one. So, FPC is occupied when this request is made and could not reply on time and hence the scheduler slip occurs. [PR1517076](#)
- In an IP-in-IP and Layer 3 VPN topology, CE to CE IPv6 traceroute picks vrf v6 address in Junos OS 18.2X75 but global instance picks v6 loopback address in Junos OS 20.3. [PR1518978](#)
- BGP based SR-TE paths based on BGP-LU prefix SID labels is not supported on PTX1000 routers. [PR1544277](#)

Routing Protocols

- Due to race between route re-converge and BGP-PIC version up message to Packet Forwarding Engine, after a remote transit router reboot, certain BGP routes might reuse stale LDP next hops and cause packet to discard at the transit router during the route re-convergence window. [PR1495435](#)

SEE ALSO

[What's New | 154](#)

[What's Changed | 162](#)

[Open Issues | 165](#)

[Resolved Issues | 168](#)

[Documentation Updates | 170](#)

[Migration, Upgrade, and Downgrade Instructions | 171](#)

Open Issues

IN THIS SECTION

- [General Routing | 165](#)
- [Layer 2 Ethernet Services | 167](#)
- [MPLS | 167](#)
- [Platform and Infrastructure | 167](#)
- [Routing Protocols | 167](#)

Learn about open issues in the Junos OS Release 20.4R1 for the PTX Series.

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

General Routing

- PTX Series platforms with FPC-PTX-P1-A or FPC2-PTX-P1A line card might encounter a single event upset (SEU) event that can cause a linked-list corruption of the TQCHIP. The following syslog message is reported: **Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002.** The Junos OS chassis management error handling does detect such condition, and raises an alarm and performs the disable-pfe action for the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, a restart of the FPC is needed. Soft errors are transient or non-recurring. FPCs experiencing such SEU events do not have any permanent damage. Contact your Juniper support representative if the issue is seen after a FPC restart. [PR1254415](#)
- When CFP2-DCO-T-WDM-1 is plugged to the PIC of the device, after FPC restarts, the carrier frequency offset TCA is raised even when TCA is not enabled. [PR1301471](#)

- Telemetry statistics might not account correctly for the traffic on SR-TE policies (both byte count and packet count). This is a sensor-related issue. [PR1413680](#)
- The firewall counter for lo0 interface might not increase. As a workaround, set the lo0 filter family inet and family inet6 counters instead of filter family any. [PR1420560](#)
- Memory leaks are expected in this release. [PR1438358](#)
- On PTX1000 routers, the vmhost disk usage might keep increasing due to an incorrect sensor path. [PR1480217](#)
- SNMP index on PFE is 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](#)
- Flap might be observed on channelized ports during ZTP when one of the ports is disabled on supporting device. [PR1534614](#)
- sFlow reports incorrect **Extended Router Data** for traffic going over a non-default VRF. [PR1537190](#)
- On PTX platforms with 18.1 or higher release, chassisd memory leak might be caused by configuration commit. When chassisd consumes ~3.4 GB of memory it might crash. chassisd crash might cause the GRES or/and FPC restart. If the GRES is enabled, commits are being synchronized between Routing Engines, so backup Routing Engine chassisd might suffer from memory leak too. [PR1537194](#)
- The output VLAN is not reported correctly in the extended switch data for IP-IP transit traffic when you configure both dynamic tunnel and FTI as backup. [PR1537648](#)

Layer 2 Ethernet Services

- It is observed rarely that issuing a **request system zeroize** does not trigger ZTP. As a workaround is to re-initiate the ZTP. [PR1529246](#)

MPLS

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

Platform and Infrastructure

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. [PR1552603](#)

Routing Protocols

- On setup with dynamic tunnel IPoIP configured on it, if the **clear bgp neighbor** command is executed on it, then ECMP next hop might be created in wrong state and traffic loss might be seen. As a workaround, restart the RPD or FPC which creates the ECMP in correct state. [PR1514966](#)
- The SSH connections are allowed more than the configured SSH connection-limit. [PR1559305](#)

SEE ALSO

What's New 154
What's Changed 162
Known Limitations 164
Resolved Issues 168
Documentation Updates 170
Migration, Upgrade, and Downgrade Instructions 171

Resolved Issues

IN THIS SECTION

- General Routing | 168
- Infrastructure | 169
- Interfaces and Chassis | 169
- MPLS | 169
- Network Management and Monitoring | 169
- Routing Protocols | 170

Learn which issues were resolved in the Junos OS Release 20.4R1 for the PTX Series.

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

General Routing

- On PTX10008 and PTX5000 routers, the output of the **show filter index number counter** command shows value as zero. [PR1420057](#)
- On PTX10016 routers, after device reboot, the FPC takes a long time to come up and hence MKA sessions establishment is delayed. The error message **Frame 08: sp = 0x48d222b8, pc = 0x10fad3bc , blaze fpc2 SCHED: Thread 59 (PFE Manager) ran for 2177 ms without yielding** is observed. [PR1477585](#)
- On PTX10016 routers, if aggregated Ethernet member or interface flow control is in disabled state, then it does not enable its own. [PR1478715](#)
- The Layer 2 VPN might flap and the CE device facing interface cannot restore the TX optical laser power even if the Layer 2 VPN is in the **Up** status under the asynchronous-notification. [PR1486181](#)
- In IP-in-IP, end-to-end (CE device to CE device) traceroute is not working as expected. [PR1488379](#)
- Dynamic tunnels traceoptions does not offer state tracing and causes JTASK_SCHED_SLIP with single underlay route bounce. [PR1493236](#)
- FPC ukern core file is not transferred to Routing Engine in a scaled setup. [PR1500418](#)
- The error message **mpls_extra NULL** might be seen during add, change, and delete of MPLS route. [PR1502385](#)
- The packetio crashes during the initialization and this might result in a second reboot. [PR1505150](#)

- ERO update by the controller for branch LSP might cause issues. [PR1508412](#)
- BIND does not sufficiently limit the number of fetches performed when processing referrals. [PR1512212](#)
- The routes update might fail upon the HMC memory issue and traffic impact might be seen. [PR1515092](#)
- On PTX5000 and PTX3000 routers, the FPC E might get stuck when the packet is switched internally between FPC connected port towards Routing Engine connected ports. [PR1519673](#)
- Sampling with the rate limiter statement enabled crosses the sample rate of 65535. [PR1525589](#)
- Running SNMP MIB walk and executing **show interfaces** command might cause the picd to crash. [PR1533766](#)
- The error message **expr_dfw_action_topo_connect_anh:1434**
expr_dfw_action_topo_connect_anh:eda_anh_discard is FALSE for nh-id 568 - return is observed in PTX1000 routers. [PR1540064](#)
- The Packet Forwarding Engine might crash in MPLS IPv6-tunneling scenario when the next hop changes. [PR1540793](#)
- Traffic might drop silently after swapping an FPC type 3 card with an FPC type 1 card in the same slot on a PTX3000 router. [PR1547790](#)

Infrastructure

- Interface drop counters might display 0 during a race condition and voq statistics are also polled simultaneously. [PR1537960](#)

Interfaces and Chassis

- The error message **Request failed: OID not increasing: ieee8021CfmStackServiceSelectorType** is observed. [PR1517046](#)
- EOAM IEEE802.3ah link discovery state is **Down** instead of **Active Send Local** after deactivating interfaces. [PR1532979](#)
- Logs are not being written in **/var/log/messages** on certain PTX Series platforms. [PR1551374](#)

MPLS

- The SNMP trap is sent with incorrect OID **jnxSpSvcSetZoneEntered**. [PR1517667](#)

Network Management and Monitoring

- The SNMP MIB **ifInErrors** [OID 1.3.6.1.2.1.2.2.1.14] reports wrong values. [PR1534286](#)

- The syslog messages might not be sent with correct port. [PR1545829](#)

Routing Protocols

- The **show dynamic-tunnels database** command does not show the statistics for the first time and fetches the traffic statistics the second time. [PR1445705](#)
- The ppm process crashes after configuring S-BFD responder on the PTX Series routers with the RE-DUO-2600 Routing Engine. [PR1477525](#)
- The rpd process might report 100 percent CPU usage with the BGP route damping enabled. [PR1514635](#)

SEE ALSO

What's New 154
What's Changed 162
Known Limitations 164
Open Issues 165
Documentation Updates 170
Migration, Upgrade, and Downgrade Instructions 171

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for PTX Series routers.

SEE ALSO

What's New 154
What's Changed 162
Known Limitations 164
Open Issues 165
Resolved Issues 168
Migration, Upgrade, and Downgrade Instructions 171

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Basic Procedure for Upgrading to Release 20.4 | 171](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases | 174](#)
- [Upgrading a Router with Redundant Routing Engines | 174](#)

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 20.4

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

NOTE: 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 20.4R1:

1. Using a Web browser, navigate to the All Junos Platforms software download URL on the Juniper Networks webpage:
<https://support.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.
4. 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-x86-64-20.4R1.9.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-x86-64-20.4R1.9-limited.tgz
```

Replace the source with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.

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

NOTE: After you install a Junos OS Release 20.4 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 19.2, 19.3, and 19.4 are EEOL releases. You can upgrade from Junos OS Release 19.2 to Release 19.3 or from Junos OS Release 19.2 to Release 19.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Upgrading a Router with Redundant Routing Engines

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

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

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

SEE ALSO

[What's New | 154](#)

[What's Changed | 162](#)

[Known Limitations | 164](#)

Open Issues 165
Resolved Issues 168
Documentation Updates 170

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- What's New | 175
- What's Changed | 205
- Known Limitations | 207
- Open Issues | 210
- Resolved Issues | 215
- Documentation Updates | 219
- Migration, Upgrade, and Downgrade Instructions | 220

These release notes accompany Junos OS Release 20.4R1 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Hardware | 176
- Class of Service (CoS) | 190
- EVPN | 191
- Flow-Based and Packet-Based Processing | 194
- High Availability (HA) and Resiliency | 197

- Interfaces and Chassis | 197
- IP Tunneling | 197
- Juniper Extension Toolkit | 197
- Junos OS XML, API, and Scripting | 198
- Junos Telemetry Interface | 198
- Network Management and Monitoring | 198
- Platform and Infrastructure | 200
- Routing Policy and Firewall Filters | 200
- Routing Protocols | 201
- Software Defined Networking (SDN) | 202
- Software Installation and Upgrade | 203
- System Management | 204
- System Logging | 204

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 20.4R1: QFX5100, QFX5110-32Q, QFX5110-48S, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, and QFX10016.

Junos on White Box runs on Accton Edgecore AS7816-64X switches in this release. The software is based on Junos OS running on QFX5210 switches, so release-note items that apply to QFX5210 switches also apply to Junos on White Box.

Hardware

- **Support for QSA Adapter (QFX10002-60C)**—Starting in Junos OS Release 20.4R1, you can use the QSA Adapter to support 1GbE and 10GbE connections on the QFX10002-60C switches.

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

- **Support for transceivers, AOCs, and DACs (QFX5120-48T)**—Starting in Junos OS Release 20.4R1, transceivers, AOCs, and DACs are supported on the QFX5120-48T switches.

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

- **New QFX5120-48YM Switch (QFX Series)**—In Junos OS Release 20.4R1, we introduce the QFX5120-48YM, an ideal switch for data center top-of-rack, leaf-and-spine deployments, enterprise multicloud deployments, and campus distribution or core deployments. The QFX5120-48YM is a 25GbE/100GbE switch that offers 48 SFP28 ports and 8 QSFP28 ports. The 48 SFP28 ports support 1-Gbps, 10-Gbps, and 25-Gbps speeds and the 8 QSFP28 ports support 40-Gbps and 100-Gbps speeds. QFX5120-48YM switches support both manual and auto-channelization, but manual CLI channelization always takes precedence. [See [Port Settings](#).]

To install the QFX5120-48YM switch hardware and perform initial software configuration, routine maintenance, and troubleshooting, see [QFX5120 Switch Hardware Guide](#). See [Feature Explorer](#) for the complete list of features for any platform.

[Table 3 on page 177](#) summarizes the QFX5120-48YM features in Junos OS Release 20.4R1.

Table 3: Features Supported by the QFX5120-48YM

Feature	Description
Access security	<ul style="list-style-type: none"> • DHCP and DHCPv6 snooping. [See DHCP Snooping.] • IP and IPv6 source guard. [See Understanding IP Source Guard for Port Security on Switches.] • Dynamic ARP inspection (DAI). [See Understanding and Using Dynamic ARP Inspection (DAI).] • IPv6 neighbor discovery inspection. [See IPv6 Neighbor Discovery Inspection.]
Authentication and access control	<ul style="list-style-type: none"> • IEEE 802.1X authentication. [See User Access and Authentication User Guide.] • RADIUS and TACACS+ authentication and accounting. [See Authentication Order for RADIUS, TACACS+, and Local Password.] • Authentication bypass access (based on host MAC address) and fallback. [See Static MAC Bypass of 802.1X and MAC RADIUS.] • Captive portal authentication for Layer 2 and Layer 3 interfaces. [See Captive Portal Authentication.]
Class of service (CoS)	<ul style="list-style-type: none"> • QFX5120-48YM switches support all class-of-service (CoS) features except the following: shared unicast and multi-destination classifiers, forwarding classes, and output queues; CoS flexible hierarchical scheduling (ETS); virtual output queue (VOQ) architecture; and CoS command to detect the source of RED-dropped packets. [See CoS Support on QFX Series Switches, EX4600 Line of Switches, and QFabric Systems.]

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
DHCP	<ul style="list-style-type: none">• DHCP server. [See DHCP Server.]• DHCP relay agent and DHCP smart relay. [See DHCP Relay Agent.]• DHCP server and client in separate routing instances. [See DHCP Message Exchange Between DHCP Clients and DHCP Server in Different VRFs.]• DHCP relay with option 82 for Layer 2 VLANs and Layer 3 interface. [See DHCP Relay Agent Information Option (Option 82).]• DHCP option 82 with textual interface description. [See DHCP Relay Agent Information Option (Option 82).]• DHCPv6 option 79. [See relay-agent-option-79.]• DHCP static addresses. [See Configuring Static DHCP IP Addresses.]• Extended DHCP (also referred to as virtual router (VR) aware DHCP). [See Legacy DHCP and Extended DHCP.]

Table 3: Features Supported by the QFX5120-48YM (continued)

Feature	Description
EVPN-VXLAN	

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • In a centrally routed bridging overlay, the QFX5120-48YM switch can act as a leaf device (Layer 2 VXLAN gateway). [See Example: Configuring an EVPN Control Plane and VXLAN Data Plane.] • In an edge-routed bridging overlay, the QFX5120-48YM switch can act as a leaf device (Layer 2 and 3 VXLAN gateways). [See Example: Configuring a QFX5110 Switch as Layer 2 and 3 VXLAN Gateways in an EVPN-VXLAN Edge-Routed Bridging Overlay.] • In a campus EVPN multihoming environment, you can deploy two QFX5120-48YM switches as distribution switches with ESI-LAG (Layer 2 and 3 VXLAN gateways) to eliminate STP. • In a campus fabric environment, you can deploy QFX5120-48YM switches as distribution or core switches in a centrally routed bridging overlay or an edge routed bridging overlay. • In the spine and leaf roles, the QFX5120-48YM switches support the following features: <ul style="list-style-type: none"> • Firewall filtering and policing of EVPN-VXLAN traffic. [See Overview of Firewall Filters (QFX Series).] • IGMPv2 snooping in a centrally routed bridging overlay. Supported use cases include intra-VLAN, inter-VLAN with IRB interfaces and PIM, and inter-VLAN with a PIM gateway and Layer 2 connectivity. [See Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment.] • Support for IPv6 data traffic. [See Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay.] • MAC filtering, storm control, and port mirroring. [See MAC Filtering, Storm Control, and Port Mirroring Support in an EVPN-VXLAN Environment.] • EVPN pure Type 5 routes. [See EVPN Type-5 Route with VXLAN Encapsulation for EVPN-VXLAN.] • EVPN proxy ARP and ARP suppression, and proxy Neighbor Discovery Protocol (NDP) and NDP suppression. [See EVPN Proxy ARP and ARP Suppression, and Proxy NDP and NDP Suppression.] • Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces. [See Supported Protocols on an IRB Interface in EVPN-VXLAN.] • Virtual machine traffic optimization for ingress interfaces. [See Ingress Virtual Machine Traffic Optimization.] • Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface. [See Understanding Flexible Ethernet Services Support With EVPN-VXLAN.]

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • Selective multicast forwarding. [See Overview of Selective Multicast Forwarding.] • MAC mobility and duplicate MAC address detection and suppression. [See Overview of MAC Mobility.]
Firewalls and policers	<ul style="list-style-type: none"> • Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on the device from a source address to a destination address. The supported firewall filter and policer features include: <ul style="list-style-type: none"> • Policer mark down action and policing/rate limiting • Single-rate two-color marking (ingress), single-rate tricolor marking (color aware, color blind), and two-rate tricolor marking • Filter-based forwarding (FBF) and FBF with destination and source prefix list on IPv6 interfaces • Dynamic allocation of TCAM memory to firewall filters and error message displayed when TCAM is full • Enhanced filter classification of CPU-generated packets • Firewall filter actions: assign forwarding class, counters; logging, syslog, reject; mirroring to an interface; and permit, drop, police, and mark • Firewall filter flexible match conditions • Firewall filters on loopback interface and management interface • IPv6 fields for ingress port and VLAN firewall filters • Policer action for MPLS firewall filters • Port firewall filters (egress and ingress), routed firewall filters (egress and ingress) and VLAN firewall filters (egress and ingress) • TCP/UDP port ranges in classification • Loopback filter optimization • Firewall filtering and policing on EVPN-VXLAN traffic • Filter-based GRE de-encapsulation • Firewall filter support on Layer 3 interfaces. [See Firewall Filter Match Conditions and Actions (QFX5220).]
High availability	<ul style="list-style-type: none"> • Nonstop bridging (NSB), and nonstop active routing (NSR) for IPv6 and OSPFv2.

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
Interfaces and chassis	<ul style="list-style-type: none"> • Support for the following resiliency features: <ul style="list-style-type: none"> • Operating system resiliency to recover the Junos OS software by using the recovery mode option on the Grub menu, which is visible after BIOS has booted up. • Partial resiliency for DIMM errors, machine-check exception (MCE), and PCI Express advanced error reporting (AER). If required, you can take assistance from Juniper Networks Technical Assistance Center (JTAC) to manually debug these type of errors when they occur. <p>[See Channelizing Interfaces on QFX5120-48YM Switches.]</p> <ul style="list-style-type: none"> • Support for channelizing interfaces. The QFX5120-48YM contains a total of 56 ports, of which 8 are QSFP28 ports and 48 are SFP28 ports. To channelize speed, you use the channel-speed statement. For setting speed, you use the set chassis fpc 0 pic 0 port <25g 1g 10g 50g 100g 40g> command. The speeds supported are: <ul style="list-style-type: none"> • 1 Gbps, 10 Gbps, and 25 Gbps on SFP28 • 40 Gbps, 100 Gbps, 4x25 Gbps, 4x10 Gbps, 2x50 Gbps on QSFP28 • 2x50 Gbps, 4x25 Gbps, or 4x10 Gbps channelization is supported on ports 50 and 52. <p>[See Port Settings.]</p>
Junos OS XML API and scripting	<ul style="list-style-type: none"> • Python, SLAX, and XSLT scripting languages, commit scripts and macros, event policy and event scripts, op scripts, and SNMP scripts. [See Automation Scripting User Guide.]

Table 3: Features Supported by the QFX5120-48YM (continued)

Feature	Description
Layer 2 features	

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • Layer 2 protocol tunneling (L2PT) support 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.] • Support for the following Layer 2 multicast features: <ul style="list-style-type: none"> • IGMP snooping with IGMPv1, IGMPv2, and IGMPv3 • IGMP proxy • IGMP querier • Virtual router (VRF-lite) IGMP snooping [See IGMP Snooping Overview.] • Support for the following Layer 2 unicast features: <ul style="list-style-type: none"> • 802.1D • 802.1w (RSTP) • 802.1s (MST) • BPDU protection • Loop protection • Root protection • VSTP • 802.1Q VLAN trunking • 802.1p • IRB (Integrated routing and bridging Interface) • Layer 3 vlan-tagged sub-interfaces • 4096 VLANs • Multiple VLAN Registration Protocol (802.1ak) • MAC address filtering • MAC address aging configuration • Static MAC address assignment for interface • Pe-VLAN MAC learning (limit) • Pe-VLAN MAC learning (limit) • MAC Learning Disable • Persistent MAC (Sticky MAC) • Link Aggregation (Static and Dynamic) with LACP (Fast and Slow LACP) • LLDP • MC-LAG with configuration sync

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
	<ul style="list-style-type: none">• Uplink Failure Detection (UFD)• VxLAN L2 Gateway (Static, OVSDB, EVPN)• QinQ Tag manipulation• 802.1x (Access control) <p>[See Layer 2 Networking.]</p>

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
Layer 3 features	

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • Traceroute over Layer 3 VPN. • Virtual routing and forwarding (VRF) support in IRB interfaces in a Layer 3 VPN. • VRF-lite, BGP, IGMP, IS-IS, OSPF, PIM, and RIP. • Support for the following Layer 3 multicast features: <ul style="list-style-type: none"> • IGMP version 1 (IGMPv1), version 2 (IGMPv2), and version 3 (IGMPv3) • IGMP filtering • PIM sparse mode (PIM-SM) • PIM dense mode (PIM-DM) • PIM source-specific multicast (PIM-SSM) <p>Multicast Source Discovery Protocol (MSDP) IGMP and PIM are also supported on virtual routers.</p> <p>[See Multicast Overview.]</p> • Support for the following Layer 3 unicast features: <ul style="list-style-type: none"> • Virtual Router Redundancy Protocol (VRRP) • Static routing • OSPFv2 • IPv4 BGP • IPv4 MBGP • BGP 4-byte ASN support • BGP Add Path (BGP-AP) • IS-IS • BFD (for RIP, OSPF, ISIS, BGP, PIM) • Filter based forwarding (FBF) • Unicast reverse path forwarding (unicast RPF) • IP directed broadcast traffic forwarding • IPv4 over GRE • IPv6 neighbor discovery protocol • Path MTU discovery • IPv6 CoS (BA, classification and rewrite, scheduling based on TC) • IPv6 ping • IPv6 static routing • IPv6 traceroute • IPv6 stateless auto-configuration

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
	<ul style="list-style-type: none"> • IPv6 OSPFv3 • IPv6 IS-IS • IPv6 BGP • VRRPv3 • 32-way equal-cost multipath (ECMP) • VXLAN Layer 3 Gateway • MPLS over UDP • DHCP snooping • IPv6 Ready Logo certification
MPLS	<ul style="list-style-type: none"> • MPLS support for label edge routers (LERs) and label switch routers (LSR). [See MPLS Overview for Switches.] • MPLS signaling protocols LDP and RSVP. [See LDP Overview and RSVP Overview.] • Fast reroute (FRR) support (a component of MPLS local protection for both one-to-one and many-to-one local protection). • Static LSPs. [See LSP Overview.] • MPLS node protection, link protection, and statistics for static LSPs. • MPLS OAM (LSP ping). • MPLS statistics. [See statistics (Protocols MPLS).] • MPLS automatic bandwidth allocation and dynamic count sizing. • MPLS with RSVP-based LSPs. • IRB interfaces over an MPLS core network. [See Example: Configuring IRB Interfaces on QFX5100 Switches over an MPLS Core Network.] • MPLS stitching for virtual machine connections. [See Using MPLS Stitching with BGP to Connect Virtual Machines.] • MPLS over Layer 3 subinterfaces. [See MPLS Limitations on QFX Series and EX4600 Switches.] • RSVP-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 MPLS Applications User Guide.] • Equal-cost multipath (ECMP) operation on MPLS using firewall filters.

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
OVSDB-VXLAN	<ul style="list-style-type: none"> • In an OVSDB-VXLAN environment, the QFX5120-48YM switch can act as a Layer 2 VXLAN gateway. [See Understanding the OVSDB Protocol Running on Juniper Networks Devices.] • In a manual (PIM-based) VXLAN environment, the QFX5120-48YM switch can act as: <ul style="list-style-type: none"> • A Layer 2 VXLAN gateway. • A transit Layer 3 switch for downstream VTEPs. • A Layer 2 VXLAN gateway. <p>[See Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches.]</p>
Network management and monitoring	<ul style="list-style-type: none"> • Support for the following services: <ul style="list-style-type: none"> • sFlow networking monitoring technology—Collects samples of network packets and sends them in a UDP datagram to a monitoring station called a <i>collector</i>. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously. The inline-sampling configuration option is available. • Local, remote, and extended port mirroring—Copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface (local port mirroring), to a VLAN (remote port mirroring), or to the IP address of a device running an analyzer application on a remote network (extended port mirroring). When you use extended port mirroring, the mirrored packets are GRE-encapsulated. • Storm control—Causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the <i>storm control level</i>—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded. <p>[See Overview of sFlow Technology, Understanding Port Mirroring, and Understanding Storm Control.]</p> <ul style="list-style-type: none"> • Support for adding nonnative YANG modules to the Junos OS schema. [See Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS.] • Puppet for Junos OS support. [See Puppet for Junos OS Administration Guide.]
Port Security	<ul style="list-style-type: none"> • Support for Media Access Control security with 256-bit cipher suite. GCM-AES-256 has a maximum key length of 256 bits and is also available with extended packet numbering (GCM-AES-XPB-256). [See Understanding Media Access Control Security (MACsec).]

Table 3: Features Supported by the QFX5120-48YM (*continued*)

Feature	Description
Software installation and upgrade	<ul style="list-style-type: none"> Secure boot—The implementation is based on the UEFI 2.4 standard. [See Software Installation and Upgrade Guide.] You need a license to use the features on the QFX5120-48YM. To find out what features are supported on this device, see QFX Switch Support for the Juniper Flex Program. To add, delete, and manage licenses, see Managing Licenses. Zero-touch provisioning (ZTP). [See Zero Touch Provisioning Overview.] Virtualization enables the switch to support multiple instances of Junos OS and other operating systems on the same Routing Engine. One instance of Junos OS, which runs as a guest operating system, is launched by default. You need to log in to this instance for operations and management. The Routing Engines on the QFX5120-48YM switches support the Wind River Linux 9 (WRL9) kernel version. [See What Are VMHosts?.]

To view the hardware compatibility matrix for optical interfaces and transceivers supported on the QFX5120-48YM, see the [Hardware Compatibility Tool](#).

Class of Service (CoS)

- **Priority-based flow control (PFC) using Differentiated Services code points (DSCP) at Layer 3 for untagged traffic (QFX5210 switches)**—Starting in Junos OS Release 20.4R1, to support lossless traffic across Layer 3 connections to Layer 2 subnetworks on QFX5210 switches, you can configure priority-based flow control (PFC) to operate using 6-bit DSCP values from Layer 3 headers of untagged VLAN traffic. You can do this rather than use IEEE 802.1p priority values in Layer 2 VLAN-tagged packet headers. You need DSCP-based PFC to support remote direct memory access (RDMA) over converged Ethernet version 2 (RoCEv2).

To enable DSCP-based PFC:

1. Map a forwarding class to a PFC priority using the **pfc-priority** statement.
2. Define a congestion notification profile to enable PFC on traffic specified by a 6-bit DSCP value.
3. Set up a classifier for the DSCP value and the PFC-mapped forwarding class.

[See [Understanding PFC Using DSCP at Layer 3 for Untagged Traffic.](#)]

EVPN

- **Multicast with IGMPv3 in EVPN-VXLAN centrally-routed bridging overlay fabrics (QFX10000, QFX5110, and QFX5120)**—Starting in Junos OS Release 20.4R1, you can configure IGMPv3 multicast in an EVPN-VXLAN centrally-routed bridging overlay fabric with multihoming for the following IPv4 multicast traffic use cases:
 - Intra-VLAN forwarding
 - Inter-VLAN routing using:
 - IRB interfaces configured with PIM
 - A PIM gateway router (for Layer 2 or Layer 3 connectivity)
 - An external multicast router

IGMPv3 multicast works with these multicast optimizations:

- IGMP snooping
- Selective multicast (SMET) forwarding
- Assisted replication (AR)

These devices process IGMPv3 reports in one of two modes:

- As any-source multicast (ASM) (*,G) reports by default
- As source-specific multicast (SSM) (S,G) reports only (if you explicitly configure this mode)

[See [Overview of Multicast Forwarding with IGMP Snooping in an EVPN-VXLAN Environment](#), [Overview of Selective Multicast Forwarding](#), [Assisted Replication Multicast Optimization in EVPN Networks](#), and [evpn-ssm-reports-only](#).]

- **MAC-VRF with EVPN-VXLAN (MX Series and vMX routers; QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002-60C, QFX10008, and QFX10016 switches)**—Data center service providers must support multiple customers with their own routing and bridging policies in the same physical network. To accommodate this requirement, you can now configure multiple customer-specific EVPN instances (EVIs) of type **mac-vrf**, each of which can support a different EVPN service type. This configuration results in customer-specific virtual routing and forwarding (VRF) tables with MAC addresses on each Juniper Networks device that serves as a virtual tunnel endpoint (VTEP) in the EVPN-VXLAN network.

NOTE: We support MAC-VRF routing instances for EVPN unicast routes only.

To support this feature, we introduce a uniform routing instance configuration, which complies with RFC 7432, *BGP MPLS-Based Ethernet VPN*. The uniform configuration eliminates hardware restrictions that limit the number of EVIs and combinations of EVIs with their respective policies that can simultaneously exist. The common configuration includes the following new CLI elements:

- The **mac-vrf** keyword at the **[edit routing-instances *name* instance-type]** hierarchy level.
- The **service-type** configuration statement at the **[edit routing-instances *name*]** hierarchy level. We support VLAN-based, VLAN-aware, and VLAN-bundle service types.
- (QFX10000 line of switches only) The **forwarding-instance** configuration statement at the **[edit routing-instances *name*]** hierarchy level. With this optional configuration statement, you can map multiple routing instances to a single forwarding instance. If you don't include this configuration statement, the default forwarding instance is used.

We continue to support the existing method of routing instance configuration along with the new uniform routing instance configuration.

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

- **Filter-based forwarding in EVPN-VXLAN networks (QFX10002-36Q, QFX10002-72Q, QFX10002-60C, QFX10008, and QFX10016)**—Instead of using the routing functionality typically provided by routing protocols, you can now use firewall filter-based forwarding. With filter-based forwarding, you can use firewall filter match conditions and actions to better control how traffic is routed in your EVPN-VXLAN network.

We support the following filter-based forwarding use cases:

- Redirecting traffic received on IRB interfaces. To set up a firewall filter:
 - Create an input filter.

NOTE: When specifying an IP address for this filter, you can use either IPv4 or IPv6 addresses.

- Specify the following match criteria:
 - The source and destination IP addresses in the inner header after a packet is de-encapsulated.
 - The source and destination ports in the inner header.
- Specify an action that directs matching packets to one of the following:
 - A routing instance
 - A next hop
 - A next hop and a routing instance
- Apply the filter to an IRB interface with or without a virtual gateway address or an anycast address.

- Handling transit traffic that matches a VXLAN network identifier (VNI). To set up a firewall filter:
 - Create an input filter.

NOTE: When specifying an IP address for this filter, you must use an IPv4 address.

- Specify the following match criteria:
 - VNI.
 - The source and destination IP addresses in the outer header.
- For the action, specify **count** or any other firewall filter action that is supported by the QFX10000 line of switches.
- Apply the filter to a Layer 3 interface.

[See [Understanding Filter-Based Forwarding](#).]

- **Loop detection for EVPN-VXLAN fabrics (QFX5120, QFX5200, QFX5210, and QFX5220)**—You can configure loop detection on the server-facing Layer 2 interfaces of the leaf devices in an EVPN-VXLAN fabric. This feature can detect the following types of Ethernet loops:
 - A loop between two interfaces with different Ethernet segment identifiers (ESIs). This loop is typically caused by miswiring fabric components.
 - A loop between two interfaces with the same ESI. This loop is typically caused by miswiring a third-party switch to the fabric.

After you've enabled loop detection, the interfaces periodically send multicast loop-detection protocol data units (PDUs). If a loop detection-enabled interface receives a PDU, a loop is detected, which triggers the configured action to break the loop. For example, if the configured action is interface-down, the interface is brought down. After the revert-interval timer expires, the configured action is reverted, and the interface is brought back up again.

[See [loop-detect](#).]

- **IPv6 multicast with MLDv1 and MLDv2 in EVPN-VXLAN centrally-routed bridging overlay fabrics (QFX10000, QFX5110, and QFX5120)**—Starting in Junos OS Release 20.4R1, you can configure MLDv1 and MLDv2 multicast in an EVPN-VXLAN centrally-routed bridging overlay fabric with multihoming for the following IPv6 multicast traffic use cases:
 - Intra-VLAN forwarding
 - Inter-VLAN routing using:
 - IRB interfaces configured with PIM
 - A PIM gateway router (for Layer 2 or Layer 3 connectivity)
 - An external multicast router

MLD multicast works with these multicast optimizations:

- MLD snooping
- Selective multicast (SMET) forwarding
- Assisted replication (AR)

These devices process MLD reports as follows:

- MLDv1 reports as any-source multicast (ASM) (*,G) reports
- MLDv2 reports in one of two modes:
 - As any-source multicast (ASM) (*,G) reports by default
 - As source-specific multicast (SSM) (S,G) reports only (if you explicitly configure this mode)

[See [Overview of Multicast Forwarding with IGMP Snooping or MLD Snooping in an EVPN-VXLAN Environment](#), [Overview of Selective Multicast Forwarding](#), [Assisted Replication Multicast Optimization in EVPN Networks](#), and [evpn-ssm-reports-only](#).]

- **Seamless EVPN-VXLAN stitching with multicast support (QFX10002-36Q, QFX10002-72Q, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.4R1, we support the following multicast features with the seamless EVPN-VXLAN stitching data center interconnect (DCI) use case:
 - Protocol Independent Multicast (PIM)
 - Internet Group Management Protocol version 2 (IGMPv2) snooping
 - Assisted replication (AR) with the following use cases:
 - The super spine device, which interconnects the data centers, and the AR replicator are two separate devices.
 - The super spine device and the AR replicator are the same device.
 - Selective multicast Ethernet tag (SMET)

[See [interconnect](#).]

Flow-Based and Packet-Based Processing

- **Support for user-defined flex hashing for MPLS traffic flows (QFX5210; Accton Edgecore AS7816 running Junos OS on White Box)**—Starting in Junos OS Release 20.4R1, you can configure user-defined flex hashing to load-balance MPLS traffic based on TCP or UDP source and 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. You can configure the selected bytes to be directly used for hashing or to be used only when the data pattern in these bytes matches 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 num_labels hash-offset
offset1 base-offset1 offset1-value offset1_value offset1-mask offset1_mask offset2 base-offset2 offset2-value
offset2_value offset2-mask offset2_mask
```

To configure a conditional match (repeat the following command with values for offsets and match data 2-4):

```
set forwarding-options enhanced-hash-key conditional-match name offset1 base-offset1 offset1-value
offset1_value offset1-mask offset1_mask matchdata1 matchdata1 matchdata1-mask 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 the **show forwarding-options enhanced-hash-key** command.

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.
- Apply a similar set of commands to the various member interfaces. For example, if the members of a particular LAG are xe-0/0/2, xe-0/0/3, and xe-0/0/4, configure three slightly different flex-hashing rules on those individual member interfaces—the rules are identical except that they have different names for the incoming traffic:
 - **set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_4 ethtype mpls interface xe-0/0/2**
 - set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_4 ethtype mpls num-labels 2**
 - set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_4 ethtype mpls conditional-match COND_L2_V6_TCP_4**
 - set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_4 ethtype mpls hash-offset offset1 base-offset1 start-of-L3-OuterHeader**

**set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_4 ethtype mpls
hash-offset offset1 offset2 offset2-mask ffff**
[...configuration commands truncated...]

- **set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_5 ethtype mpls
interface xe-0/0/3**
[...configuration commands truncated...]

- **set forwarding-options enhanced-hash-key flex-hashing FLEX_L2_V6_TCP_6 ethtype mpls
interface xe-0/0/4**
[...configuration commands truncated...]

- Use unique flex-data profile names and unique conditional-data profile names for each member interface—for example, in the following conditional-data profile names, the port number is unique in each instance:
 - ...enhanced-hash-key conditional-match COND_L1_V6_UDP_SRC_PORT_1...
 - ...enhanced-hash-key conditional-match COND_L1_V6_UDP_SRC_PORT_2...

[See [flex-hashing](#).]

High Availability (HA) and Resiliency

- **Disk health monitoring (QFX5100 and QFX5120-48Y)**—Starting in Junos OS Release 20.4R1, the QFX5100 and QFX5120-48Y switches support disk-health monitoring. This feature detects solid-state drive (SSD) failures and reboots the device gracefully to handle those failures. With this feature in place, you need not manually intervene to recover the system from an SSD failure condition. The QFX5100 switches used in a Virtual Chassis as well support this feature.

[See [show chassis routing-engine](#)]

Interfaces and Chassis

- **Retrieve an ECMP or trunk interface hardware hash result for a given input for load balancing (QFX5120)**—Starting in Junos OS Release 20.4R1, you can view the hash parameters that are used by the hashing algorithm and the final egress interface for the traffic you are interested in. Use the CLI command **show forwarding-options load-balance ecmp|trunk** to retrieve this information. The command output provides you information to troubleshoot issues for which you need to know the packet path.

[See [Configuring the Fields in the Algorithm Used To Hash LAG Bundle and ECMP Traffic \(CLI Procedure\)](#) and [show forwarding-options load-balance ecmp|trunk](#).]

IP Tunneling

- **IPIP encapsulation for flexible tunnel interfaces (FTIs) (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210)**—We've extended flexible tunnel interfaces (FTIs) and existing forwarding constructs to support configuring static IPv4 IP-in-IP tunnels and RIB APIs. To configure an IP-in-IP tunnel on a FTI, use the **ipip** option at the **[edit interfaces interface-name unit logical-unit-number tunnel encapsulation]** hierarchy level.

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

Juniper Extension Toolkit

- **Support for static backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210)**—We've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to set up underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation. JET APIs notify the controller of active paths, interfaces, and changes to the interface state. The loop-free backup paths help quickly restore failed core transport networks built with only IP protocols.

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

- **Support for policy match condition to match programmed routes (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210)**—We've introduced a new option programmed that allows policy matches for

routes injected by JET APIs. To allow policy matches for routes injected by JET APIs, use the **programmed** option at the `[edit policy-options policy-statement policy-name term term-name from]` hierarchy level. To view details about programmed routes, use the **show route programmed (detail | extensive)** command.

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

- **RIB service API option to control route distribution (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210)**—We've added a no-advertise flag to the RIB service API per-route RouteAttributes object to limit re-advertisement of the provisioned route. You can set this flag to TRUE to prevent the route from being redistributed to routing protocols and advertised to peers.

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

Junos OS XML, API, and Scripting

- **Support for Certificate Authority Chain Profile (EX2300, EX3400, EX4300, MX240, MX480, MX960, PTX-5000, VMX, vSRX and QFX5200)**—Starting in Junos OS Release 20.4R1, you can configure intermediate Certificate Authority (CA) chain profile certificate and perform https REST API request using mutual and server authentications.

To configure intermediate ca-chain certificate, configure **ca-chain ca-chain** statement at the `[edit system services rest https]` hierarchy level.

- **Start time option for interval-based internal events that trigger event policies (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.4R1, when you create an interval-based internal event for triggering event policies, you can specify the start date and time for the initial event. To specify a start time, configure the **start-time** option along with the **time-interval** option at the `[edit event-options generate-event]` hierarchy level.

[See [Generating Internal Events to Trigger Event Policies](#).]

Junos Telemetry Interface

- **BGP neighbor telemetry with sharding (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, BGP neighbor telemetry with sharding (multi-threading) is supported.

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

Network Management and Monitoring

- **Support for sFlow technology with VXLAN (QFX5110, QFX5120-32C, QFX5120-48Y, and QFX5120-48T-6C)**—Starting in Junos OS Release 20.4R1, we support sFlow technology with VXLAN. sFlow is a monitoring technology for high-speed switched or routed networks. The sFlow agent performs packet sampling and gathers interface statistics, and then combines the information into UDP datagrams that are sent to sFlow collectors.

Keep the following points in mind:

- True egress sampling is not supported on the switches. Egress sampling is done at the end of the ingress pipeline, and so the egress samples do not contain modifications that are made to the packet in the egress pipeline. When the packet flow is from the access port to the network port, sFlow is configured on the egress network port, and thus packets are captured without VXLAN encapsulation at the collector.
- Egress sampling for BUM traffic is not supported.
- Extended router data with next-hop information is not supported on the switches.
- Sampling on ingress interfaces does not capture CPU-bound traffic.
- You cannot configure sFlow on a (LAG), but you can configure it individually on a LAG member interface.
- You must not configure sFlow for more than one logical interface on a physical interface.

[See [Overview of sFlow Technology](#).]

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **sFlow sampling support for IP-IP traffic (QFX5100 and QFX5200)**—Starting in Junos OS Release 20.4R1, you can use sFlow technology to sample IP over IP (IP-IP) traffic at a physical port. This feature is supported for IP-IP tunnels with an IPv4 outer header that carry IPv4 or IPv6 traffic. You can use sFlow monitoring technology to randomly sample network packets from IP-IP tunnels to send the samples to a destination collector for monitoring. Devices that act as an IP-IP tunnel entry point, transit device, or tunnel endpoint support sFlow sampling.

[See [Overview of sFlow Technology](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The **<get-configuration>** operation supports the **compare="configuration-revision"** and **configuration-revision** attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

Platform and Infrastructure

- **Flooding bridge protocol data units (BPDUs) using existing ingress port-based firewall filters (QFX5100, QFX5110, QFX5120-32C, QFX5200, and QFX5210)**—Starting in Junos OS Release 20.4R1, you can configure a new firewall filter CLI action to flood BPDUs using the **set firewall family ethernet-switching filter f1 term t1 then flood** statement. The flexibility to flood BPDUs on a per port basis for the QFX5000 line of switches can be achieved by using the existing ingress port-based firewall filters.

[See [Configuring a Firewall Filter](#).]

Routing Policy and Firewall Filters

- **Support for route's next-hop weight in policy match condition (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, a route with multiple next-hop paths can use the weight associated with a path to identify primary and backup paths. The path with the lowest weight is used as the primary path, and any paths with higher weights are treated as backup paths. You can use the next-hop weight as a match condition in export policies to redistribute IGP and BGP routes based on whether the primary or backup paths are active.

Configure this match condition using the **[edit policy-options policy-statement *policy-name* term *term-name* from]** statement.

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

- **IPv6 support for firewall filtering and policing on EVPN-VXLAN traffic (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210)**—Starting with Junos OS Release 20.4R1, you can use IPv6 for firewall filters and policers on VXLAN traffic in an EVPN topology. Configure firewall filters at the **[edit firewall]** hierarchy level. For each firewall filter that you apply to a VXLAN, you can specify **family ethernet-switching** to filter Layer 2 (Ethernet) packets or **family inet** or **family inet6** to filter on IRB interfaces. You can apply firewall filters and policers on CE-facing interfaces in the ingress direction only. For IRB interfaces, you can apply filtering only at the ingress point of a non-encapsulated frame routed through the IRB interface.

[See [Understanding VXLANs](#) and [Overview of Firewall Filters](#).]

- **Support for matching IPv6 source addresses from an inet6 egress interface (QFX5110)**—Starting in Junos OS Release 20.4R1, you can configure a firewall filter on an IPv6 egress interface to match specified IPv6 source or destination addresses, for example, to protect a third-party device connected to the switch.

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

Routing Protocols

- **Support for multiple single-hop EBGp sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGp session. You can now enable single-hop EBGp sessions on different links over multiple directly connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGp Sessions on Different Links Using the Same Link-Local Address \(IPv6\)](#).]

- **Support for IS-IS flood-reflector interfaces (PTX1000, QFX10002, QFX10008)**—Starting in Junos OS Release 20.4R1, we support the IS-IS flood reflector feature, which allows creation of IS-IS flood reflection topologies. Flood reflection allows the creation of topologies where Level 1 areas provide transit forwarding for Level 2 destinations within a Level 2 topology that provides better scalability.

We designate flexible tunnel interfaces (FTI) as flood-reflector interfaces. To enable the flood reflector on an FTI, include the **flood-reflector** statement at the **[edit protocols isis interface *interface name* level *level number*** hierarchy level.

You can configure the interface to be either the reflector or the client. To enable the reflector, you can use the **flood-reflector reflector *cluster-id*** statement at the **[edit protocols isis level *level number*** hierarchy level.

To enable the flood reflector client, include the **flood-reflector client** statement at the **[edit protocols isis level *level number*** hierarchy level.

NOTE: You can configure the flood reflector feature on FTIs at Level 2 only.

[See [How to Configure Flood Reflector Interfaces in IS-IS Networks](#).]

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC6286 (MX204, NFX Series, PTX5000, QFX Series, and vRR)**—Starting in Junos OS Release 20.4R1, you can establish a BGP connection using a BGP identifier that is a 4-octet, unsigned, nonzero integer and it needs to be unique only within the autonomous system (AS) per RFC 6286. In earlier releases, the BGP ID of a BGP speaker was required to be a valid IPv4 host address assigned to the BGP speaker.

To enable this feature, use the **bgp-identifier *identifier* group *bgp group name* bgp-identifier *identifier* neighbor *peer address* bgp-identifier *identifier*** configuration statement at the **[edit protocols bgp]** hierarchy level.

[See [router-id](#)]

- **Support for IPv4 VPN unicast and IPv6 VPN unicast address families in BGP (QFX10002-60C, QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 20.4R1, on QFX Series switches, the following address families are supported to enable advertisement or reception, or both, of multiple paths to a destination to and from the same BGP peer, instead of advertising and receiving only one path to and from the same BGP peer, under the `[edit protocols bgp group group-name]` hierarchy. You can configure the **add-path** statement at the BGP global, group level, and peer level.

- IPv4 VPN unicast (family inet-vpn)
- IPv6 VPN unicast (family inet6-vpn)

[See [Understanding the Advertisement of Multiple Paths to a Single Destination in BGP.](#)]

Software Defined Networking (SDN)

- **PCEP support for color (MX480, QFX5200)**—Starting in Junos OS Release 20.4R1, the Path Computation Element Protocol (PCEP) supports color for colored segment routing LSPs. This includes Path Computation Element (PCE)-initiated, Path Computation Client (PCC)-controlled, and PCC-delegated segment routing LSPs. With this PCEP extension, you can configure candidate paths based on color and endpoints, where the active candidate path is the path with the highest segment routing preference, or based on source priority.

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

- **Static VXLAN at VLAN or bridge domain level (MX5, MX10, MX40, MX80, MX150, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016 routers and QFX5120-32C, QFX5120-48T, and QFX5120-48Y switches)**—In Junos OS Release 20.3R1 and earlier, we supported the configuration of static VXLAN at the global level only. By including the **remote-vtep-list** configuration statement at the `[edit switch-options]` or `[edit routing-instances name]` hierarchy level, you can map all local VLANs or bridge domains to the remote virtual tunnel endpoints (VTEPs) in the list.

Starting in Junos OS Release 20.4R1, you can also configure static VXLAN at the VLAN or bridge domain level using the **static-remote-vtep-list** configuration statement at the `[edit vlans name vxlan]`, `[edit bridge-domains name vxlan]`, or `[edit routing-instances name bridge-domains name vxlan]` hierarchy level.

When specifying remote VTEPs at the VLAN level in the default switching instance, you must also specify the same VTEPs at the global level in the default switching instance. Or when specifying remote VTEPs at the bridge domain level in a routing instance, you must also specify the same VTEPs at the global level in the same routing instance. For example, if you specify a VTEP in the **static-remote-vtep-list** at the `[edit routing-instances name bridge-domains name vxlan]` hierarchy level, you must also specify the VTEP in the **remote-vtep-list** at the `[edit routing-instances name]` hierarchy level.

To replicate and flood BUM traffic, you must specify the **ingress-node-replication** configuration statement at the `[edit vlans name vxlan]`, `[edit bridge-domains name vxlan]`, or `[edit routing-instances name]`

bridge-domains name vxlan] hierarchy level. This configuration restricts the BUM traffic flood domain to only those VTEPs mapped to a particular bridge domain or VLAN.

[See [Static VXLAN](#) and [static-remote-vtep-list](#).]

Software Installation and Upgrade

- **Phone-home client (EX4600, EX4650, EX9200, QFX5110, QFX5200, QFX5210, QFX5120-32C, and QFX5120-48Y)**—Starting with Junos OS Release 20.4R1, you can use either the legacy DHCP-options-based ZTP or the phone-home client (PHC) to provision software for the switch. When the switch boots up, if there are DHCP options that have been received from the DHCP server for ZTP, ZTP resumes. If DHCP options are not present, PHC is attempted. PHC enables the switch to securely obtain bootstrapping data, such as a configuration or software image, with no user intervention other than having to physically connect the switch to the network. When the switch first boots up, PHC connects to a redirect server, which redirects to a phone home server to obtain the configuration or software image.

To initiate either DHCP-options-based ZTP or PHC, the switch must be in a factory-default state, or you can issue the **request system zeroize** command.

[See [Understanding the Phone-Home Client](#)

- **ZTP with DHCPv6 client support (EX3400, EX4300, PTX1000, PTX5000, PTX10002-60C, PTX10008, QFX5100, QFX5200, QFX10002, and QFX10002-60C)**—Starting in Junos OS Release 20.4R1, zero touch 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 one of the DHCPv4 bindings fails, the device continues to check for bindings until provisioning is successful. However, if there are no DHCPv4 bindings, the device checks for DHCPv6 bindings and follows 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: ZTP supports only HTTP and HTTPS transport protocols.

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

System Management

- **PTP transparent clock (QFX5120-32C)**—Starting in Junos OS Release 20.4R1, you can use a transparent clock to update Precision Time Protocol (PTP) packets with the residence time as the packets pass through QFX5120-32C switches. The PTP Transparent Clock (PTP TC) is defined in IEEE 1588-2008 (PTPv2). QFX5120-32C switches support end-to-end transparent clocks, which include only the residence time. To use a transparent clock, enable the **e2e-transparent** statement at the **[edit protocols ptp]** hierarchy level.

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

System Logging

- **Support for time averaged watermark (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 20.4R1, you can capture steady state data of routing and forwarding (RIB/FIB) table routes using the **time-averaged-watermark-interval** configuration statement at the **[edit routing-options]** hierarchy level. Time averaged watermark is calculated whenever the time averaged interval is changed from CLI. Time averaged watermark is logged in syslog if the logs are enabled in the system at **LOG_NOTICE** level. The default time averaged watermark interval is 1 day. You can see the timed averaged watermark using the existing **show route summary** command.

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

SEE ALSO

[What's Changed | 205](#)

[Known Limitations | 207](#)

[Open Issues | 210](#)

[Resolved Issues | 215](#)

[Documentation Updates | 219](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

What's Changed

IN THIS SECTION

- [Class of Service \(CoS\) | 205](#)
- [General Routing | 205](#)
- [MPLS | 206](#)
- [Network Management and Monitoring | 206](#)
- [User Interface and Configuration | 206](#)

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

Class of Service (CoS)

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

General Routing

- **Control plane DDoS protection packet type option for ARP traffic (PTX Series and QFX Series)**— Starting in this release, we've renamed the **arp-snoop** packet type option in the **edit system ddos-protection protocols arp** protocol group to **arp**. This packet type option enables you to change the default control plane distributed denial-of-service (DDoS) protection policer parameters for ARP traffic.

[See [protocols \(DDoS\) \(PTX Series and QFX Series\)](#).]

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether this is configured, you can issue the following command: **show configuration system arp | grep unicast-mode-on-expire**.

[See [arp](#).]

- **Change in license bandwidth command on vMX virtual routers**—Starting in Junos OS, to use the available license bandwidth, explicitly set the license bandwidth using the **set chassis license bandwidth <ln-mbps>** command.

[See [Configuring Licenses on vMX Virtual Routers](#).]

MPLS

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

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to deviate not-supported nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.

User Interface and Configuration

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

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

SEE ALSO

[What's New | 175](#)

[Known Limitations | 207](#)

[Open Issues | 210](#)

[Resolved Issues | 215](#)

[Documentation Updates | 219](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

Known Limitations

IN THIS SECTION

- [General Routing | 207](#)
- [Layer 2 Features | 209](#)
- [Routing Protocols | 209](#)

Learn about known limitations in Junos OS Release 20.4R1 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.

These EVPN-VxLAN features are not supported in some of the QFX Series platforms with Junos OS Release 20.4R1:

Feature	Unsupported Platforms	Tracking PRs
L3 Gateway functionality	QFX10002, QFX10008, QFX10016, QFX10002-60C	1561102 , 1561115 , 1522585
VPLAG (with Type-5) traffic	QFX5110, QFX5120	1545178 , 1560173
Native VLAN ID configuration	QFX5110, QFX5120	1552671 , 1559813 , 1560038
VLAN 2 (Untagged traffic routed over native vlan)	QFX5100	1560161

General Routing

- After configuring and deleting the Ethernet loopback configuration, the interface goes down and does not come up. [PR1353734](#)
- When adding or deleting routes that are pointed to by the unilist next hop, reroute counter log events might be seen. These are harmless messages and do not have any functionality impact. [PR1380350](#)
- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. The device can be recovered by power-cycling it. [PR1385970](#)

- On the QFX10000 line of devices, if an analyzer is configured to mirror traffic of an input aggregated Ethernet interface and a new member is added to the same aggregated Ethernet interface, then the analyzer might not provide sample packets that flow through a newly added child interface. [PR1417694](#)
- After changing the VLAN name on the trunk interface, while the port is receiving continuous traffic for that VLAN, local host MAC learning holds for more than 30 seconds. In case of trunk port, when the VLAN name is changed, the bridge domain entry is deleted from hardware and a new entry is installed in the hardware. When the new entry is yet to be installed in hardware, the port keeps receiving traffic for that VLAN and learning source MAC and notifies to Packet Forwarding Engine with the old bridge domain ID. When Packet Forwarding Engine software receives this MAC and drops it, as bridge domain and port mapping is not present in software, which is a must criteria for a source MAC received on a bridge domain. Once Packet Forwarding Engine drops the MAC, upper layers (L2ALD) does not get this MAC info and aging thread marks the hash index in hardware as stale. Until that hash index is not cleared in the hardware, same source MAC cannot be learned on the same hash index. Aging thread periodically scans one MAC table out of 4 tables at a time in intervals of 10 seconds and checks for stale entries and clear the hardware hash stale entry, and this time is almost 40-50 seconds based on the number of Packet Forwarding Engine chips in an FPC. In case of an access port, the default bridge domain is installed in the hardware to receive untagged traffic and does not get deleted while changing the VLAN name associated to that access port. So this issue is not seen for an access port. [PR1454274](#)
- On a fully scaled system where all the slices are utilized by different families of CLI filters, if you try to delete one family and change another family with higher number of filter terms that requires expansion of the filter, the Packet Forwarding Engine fails to add the new changed filter as out-of-sequence messages are generated, that is, change of filter is called earlier than deletion of another filter. [PR1512242](#)
- In case of fan failure, "show chassis environment" and "show chassis fan" will show "failed" and "check" status, respectively. This is expected and no discrepancy in terms of real status. [PR1527628](#)
- Each streaming telemetry sensor subscription and unsubscription results in configuration changes, which are done one by one for each sensor in the subscription. There are three threads spawned upon each subscription and these threads are cleared up when all the configurations related to the subscription are deleted. In this test case, there was aggressive subscription and unsubscription done every 20 seconds. However, the back-end infrastructure takes a little more time to clear up the configurations. Eventually this difference piles up, resulting in the na-grpcd thread count growing to a large number and eventually causing a vmcore. It is advisable to introduce at least a 60-second delay between subscriptions and unsubscriptions to enable the infrastructure to clear up the subscription. [PR1528432](#)
- On QFX Series line of switches, when you apply auto channelization disable on physically separated breakout cable, the CLI basically forces optic speed to original speed, which results in speed mismatch at both ends. [PR1531850](#)
- Runt error packets are getting dropped in the PHY itself on pyrite, as per the default runt error packets settings in the PHY. So those packets won't be sent to the switch, and so the switch won't be able to capture the runt error packets. In general CLI is implemented in such a way that it can capture only the packets that are received by the switch. [PR1533322](#)

- ECMP over GRE does not work for BGP route. Traffic is polarized to just one egress interface but not distributed to multiple egress interfaces. [PR1537924](#)
- In QFX5100 and EX4300 non-TVP platforms, the sample rate is limited by the IPC between the Packet Forwarding Engine and the sFlow process, so the supported limit is around 700 samples per second in these platforms. This is applicable to any sampled packets in these platforms and not specific to IP-IP. [PR1539815](#)
- While firmware upgrade, power interruption can land the particular module/system into unknown state. Successful re-attempt is required for the particular module upgrade. [PR1543192](#)
- When both loop-detect and CFM are configured on same VLAN, the one which is configured first will only work. Thus, depending on the order of configurations, either loop-detect or CFM will not work. [PR1553384](#)

Layer 2 Features

- With the scale of logical interfaces, if the child members for aggregated interfaces are changed, then the Packet Forwarding Engine programming will land into an erroneous state, which can cause momentary CRC error and permanent traffic impact until the device is rebooted. As a result, the Packet Forwarding Engine is restarted or the concerned interfaces are flapped. [PR1532342](#)

Routing Protocols

- Traffic silently drops and gets discarded during node failures with node protection on FTI interfaces for RSVP LSPs. [PR1456350](#)
- Chip SDK does not support variable mask for destination IP address in tunnel termination table, so firewall terms for the de-encapsulation action should always have destination address as a /32 address. Source IP address can be variable mask or optional. [PR1511893](#)
- During the deactivation or activation of an FTI, if the underlay next hop is unilist and is not programmed in hardware, the following error messages are seen: **brcm_nh_fti_ip_tunnel_ulst_install(),8412:NULL platform info for ifl 0 underlay nh 131074 [Thu Jul 2 17:23:43.470 LOG: Err] BRCM_NH-,brcm_nh_fti_ip_tunnel_create(),8945:failed to install ulst tunnel for fti_ifl 548 [Thu Jul 2 17:23:43.470 LOG: Err] BRCM_NH-,brcm_nh_fti_ip_tunnel_create(),8972:fti ip tunnel create fail. fti_ifl:548.** The Packet Forwarding Engine programs the FTI when the underlay unilist next hop is programmed in hardware. [PR1522701](#)
- During scale scenario when the hardware tunnel table is full, the following messages are seen in the Packet Forwarding Engine: **[Mon Jul 13 16:48:48.223 LOG: Err] BRCM_NH-,brcm_nh_ip_tunnel_unilist_install(),5632:IPoIP <src> NH id 2726, Tunnel id: 383 failed to create encap (0) obj vid 4088 intf 3589 of nh 131070(3) [Mon Jul 13 16:48:48.227 LOG: Debug] mem EGR_L3_INTF field TUNNEL_INDEX value does not fit.** [PR1525270](#)

SEE ALSO

What's New		175
What's Changed		205
Open Issues		210
Resolved Issues		215
Documentation Updates		219
Migration, Upgrade, and Downgrade Instructions		220

Open Issues

IN THIS SECTION

- [EVPN](#) | [211](#)
- [General Routing](#) | [211](#)
- [High Availability \(HA\) and Resiliency](#) | [213](#)
- [Layer 2 Ethernet Services](#) | [213](#)
- [Layer 2 Features](#) | [213](#)
- [Platform and Infrastructure](#) | [213](#)
- [Routing Policy and Firewall Filters](#) | [214](#)
- [Routing Protocols](#) | [214](#)
- [Virtual Chassis](#) | [214](#)

Learn about open issues in Junos OS Release 20.4R1 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

- Sometimes, the changing forwarding-instance dynamically might cause vtep leak and l2ald process crash and generate a core file. There are two workarounds for this problem : 1) After forwarding-instance is changed, restart l2-learning process 2) deactivate routing-instance, change forwarding-instance, activate routing-instance. [PR1560068](#)

General Routing

- QFX10000:Source MAC and TTL values are not updated for routed multicast packets in EVPN-VXLAN. [PR1346894](#)
- Unified ISSU is not supported on the QFX5200 switches and fails from Junos OS Release 17.2X75-D43.2 to some target versions. Also, dcpfe crash might be seen. [PR1438690](#)
- The port qualifier is supported on the QFX5000 line of switches. This installs two entries in Packet Forwarding Engine, one with source-port and the second one with destination-port with the value specified under the port stanza. [PR1440980](#)
- On QFX10K8 and QFX10K16s the chassisd process might crash and generate core files on the backup Routing Engine after commit because CHASSISD_MAIN_THREAD_STALLED for 200 seconds. [PR1481143](#)
- SNMP index on Packet Forwarding Engine is 0. This causes the sFlow records to have either input interface value (IIF) or output interface value (OIF) as 0 value in the sFlow record data at the collector. [PR1484322](#)
- On EX, OCX, or QFX Series line of devices using chipset, with SFP+ implemented, interface on the platforms might be in active status when TX or RX connector is removed. When this issue happens, traffic might drop. [PR1495564](#)
- An issue was reported for a customer with a flush cache issue on the same platform. As it was root-caused to a reliable SSD disk I/O change to be made for this platform, this caused the added delay observed in the reported issue. The previous cache mode was writethrough, which is prone to errors due to the ASYNC nature of writes. In "writethrough", the host cache is not bypassed and in case failure occurs when transferring data from the host cache to the storage device, the guest (in our case Junos OS VM) is not aware and going forward the host may return various errors causing stability issues. Many side effects can be seen. [PR1513540](#)
- Some inter-VLAN traffic flows do not converge after rebooting the QFX10002 spine device in the evpn-vxlan non-collapsed scaled scenario when the traffic is already flowing. [PR1522585](#)
- Broadcom has updated that BCMX calls are deprecated and needs to be replaced with BCM calls. [PR1541159](#)
- On QFX5000 line of switches in a Virtual Chassis setup, if the master and the backup Routing Engine members are swapped by swapping serial numbers of members then, auto channelization might fail for ports of linecard members, leaving the ports in down state. [PR1544353](#)

- On the QFX5000 line of switches, the SNMP trap of power failure might not be sent out when the power cable is removed from the PSU, and the output of the CLI command 'show chassis environment' might not display the information of the power failure. [PR1520144](#)
- On the QFX10000 line of switches, when an explicit Layer 2 classifier is applied on a Layer 3 interface, default Layer 3 classifiers are not removed. And by design the Layer 3 classifier takes precedence over the Layer 2 classifier. [PR1520570](#)
- On QFX5110 or QFX5120, when the Type 5 tunnels are destroyed, sometime error messages **brcm_virtual_tunnel_port_create() ,489:Failed NW vxlan port token(45) hw-id(7026) status(Entry not found)** might be seen. There is no functionality impact due to this. [PR1535555](#)
- On QFX5100 Virtual Chassis, the firewall counter do not get updated as expected when PACL is applied. [PR1535825](#)
- On QFX10002 devices acting as PHP, egress sFlow samples do not report MPLS explicit-null label in the raw packet header. MPLS payload can be of IPv4 or IPv6 protocol. [PR1537946](#)
- With EVPN-VXLAN configuration, when restart of l2-learning command is executed, BFD sessions on an IRB interface may not come up. [PR1538600](#)
- EVPN-VXLAN: vmcore files are seen on primary and backup Routing Engines of QFX10008 with Layer 2 and Layer 3 multicast configuration. [PR1539259](#)
- QFX10002-60C - ARP/token scale is lower than QFX10002 and QFX10008 generating dcpfe core file at high scale. Xellent multi-dimensional scale to be characterized. [PR1541686](#)
- For VLAN bundle service on the QFX10000 line of switches, there will be no entries in the mac-ip table. They are normally visible as output of the "show mac-vrf forwarding mac-ip-table" command. There is no functionality impact. [PR1548456](#)
- 100GbE port with AOC from Innolight does not come up after multiple reboots. It recovers after the interface is disabled and then enabled. [PR1548525](#)
- When ethernet-bridge is configured on an interface on the QFX10000 line of Switches, only untagged BUM/directed traffic will flow through the interface. Tagged traffic will be dropped with dlu.vlan_tag_lkup_miss trapcode. [PR1550700](#)
- Ethernet-bridge is not supported on QFX10002-60C and the user doesn't have the capability to configure the feature. [PR1551037](#)
- Loading type5 EVPN VXLAN with VLANs and VNIs greater than 1000 leads to the crashing of Packet Forwarding Engine because of the resource unavailability. [PR1556561](#)

High Availability (HA) and Resiliency

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

Layer 2 Ethernet Services

- It is observed rarely that issuing a **request system zeroize** does not trigger ZTP. As a workaround is to re-initiate the ZTP. [PR1529246](#)

Layer 2 Features

- On QFX5110 and QFX5120 platforms, changing logical interface IP address might sometimes result either in stale IP entries in the mpls_entry table or in a missing IP entry, which results in traffic drop for VXLAN traffic. [PR1472333](#)
- Observed dcpfe core file at `is_heap_block`, `ukern_free`, `eth_net_detach`, `dcbcm_eth_attach`. [PR1552798](#)

Platform and Infrastructure

- On all platforms running Junos OS that support EVPN-MPLS and EVPN-VXLAN, when an existing ESI interface flaps or is added newly to the configuration, sometimes DF (Designated Forwarder) election happens before the local bias feature is enabled and during this time, existing broadcast, unknown unicast, and multicast (BUM) traffic might be looped for a short time duration (less than several seconds). [PR1493650](#)
- The client DNS queries are unable to complete from QF director devices itself because of the stateless firewall. [PR1509383](#)

Routing Policy and Firewall Filters

- On all platforms running Junos OS with "set policy-options rtf-prefix-list" configured, when you upgrade to a specific version, the device might fail to validate its configuration, which eventually might crash the rpd unexpectedly because of the software fault. [PR1538172](#)

Routing Protocols

- On QFX5100 Virtual Chassis or Virtual Chassis Fabric, when the mini-PDT-base configuration is issued, the following error message is seen in the hardware: **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(), 128:I3 nh 6594 unintsall failed**. There is no functionality impact because of this error message. [PR1407175](#)
- On setup with dynamic tunnel IPoIP configured on it, if "clear bgp neighbor" command is executed on it then ECMP next hop might be created in an incorrect state. Due to which traffic loss can be seen. As a workaround, restart the rpd or FPC which creates the ECMP in correct state. [PR1514966](#)
- When Layer 2 and Layer 3 logical interfaces are configured on the same physical interface and the vport scale is enabled on the QFX5110 and the QFX5120 and if the Layer 2 logical interface is part of a VXLAN, then the SVP is derived from the source_trunk_map table. In this case, the packet will not match with the SOURCE_FIELDS in my_station_tcam table. As a result, the entry is not getting hit. The OSPF unicast packets will be dropped and the packets will be stuck in ExStart state. [PR1519244](#)
- Layer 2 core files might be seen because of the MAC entry list corruption in certain scenarios. It is seen only in regression scripts so far. This is not reproducible consistently. This does not impact service. [PR1547866](#)

Virtual Chassis

- QFX5110-48S reports false parity error messages like "soc_mem_array_sbusdma_read". The QFX5110-48S SDK can raise false alarms for parity error messages like "soc_mem_array_sbusdma_read". This is a false positive error message. [PR1276970](#)

SEE ALSO

[What's New | 175](#)

[What's Changed | 205](#)

[Known Limitations | 207](#)

[Resolved Issues | 215](#)

[Documentation Updates | 219](#)

[Migration, Upgrade, and Downgrade Instructions | 220](#)

Resolved Issues

IN THIS SECTION

- [Resolved Issues: 20.4R1 Release](#) | 215

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: 20.4R1 Release

EVPN

- EVPN-VXLAN core isolation do not work 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 to 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 multi-home (MH) clients. [PR1538117](#)
- Observed Layer 2 core file when system is rebooted when shared-tunnels are configured. [PR1548502](#)
- The l2ald process crashes and generates a core file l2ald_iff_rtm_delete_subintf_ifbds during the datacenter interconnect (dci) fusion run. [PR1550109](#)

General Routing

- Port LEDs do not work on the QFX5100-48T-6Q platforms. [PR1317750](#)
- On the QFX5100 switches, the interface output counter is double counted for self-generated traffic. [PR1462748](#)
- IRB MAC is not programmed in hardware when the MAC persistence timer expires. [PR1484440](#)
- Virtual Chassis is not stable with 100-Gigabit Ethernet and 40-Gigabit Ethernet interfaces. [PR1497563](#)
- BFD sessions flap after deactivating or activating the aggregated Ethernet interface or executing GRES. [PR1500798](#)
- The error message "mpls_extra NULL" might be seen when you add, change, or delete MPLS route. [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 archival function might fail in certain conditions. [PR1507044](#)
- The fxpc might crash and restart with a fxpc core file created while installing image through ZTP. [PR1508611](#)
- Traffic might be affected on QFX10002, QFX10008, and QFX10016 line of switches. [PR1509220](#)
- The output VLAN push might not work. [PR1510629](#)
- Multicast traffic loss is observed because of the few multicast routes missing in the spine node. [PR1510794](#)
- The QFX10000-36Q line card used on QFX10008 and QFX10016 line of switches might fail to detect any QSFP. [PR1511155](#)
- Display issue, Virtual Chassis environment, Configured num-65-127-prefix value is shown incorrect for the command O/P "show chassis forwarding-options" [PR1512712](#)
- 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 Junos OS Release 18.1R3-S1, 18.2R1, and later. [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 (PFE) process might crash due to memory leak. [PR1517030](#)
- The vgd process might generate a core file when the OVSDB 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](#)
- QFX5100: cprod timeout triggers high CPU. [PR1520956](#)
- The interfaces on the EX4600-EM-8F expansion module do not come up on the QFX5100-24Q with the non-QFX5E image. [PR1521523](#)
- Output interface index in SFLOW packet is zero when transit traffic is 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_sl_u_sal_update_lrnent(),1379:jprds_sl_u_sal_update_lrnent call failed. [PR1522852](#)

- The ECMP and LAG hash polarization might occur if the "hash-parameters" statement is not configured. [PR1525387](#)
- Sampling with the rate limiter command enabled, crosses the sample rate 65535. [PR1525589](#)
- Traffic loss might be observed when traffic is locally routed between the two VXLANs on the QFX5120 switch. [PR1527939](#)
- The MPLS EXP classifier might not work on QFX10000 line of switches. [PR1531095](#)
- Running SNMP MIB walk and executing 'show interfaces' command may cause the picd to crash. [PR1533766](#)
- 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 filter instance do not get removed from Packet Forwarding Engine after deactivating VLAN and IRB. [PR1537108](#)
- Interfaces are not created after channel-speed 10G is applied across ports 48 to 53 on QFX5100-48T. [PR1538340](#)
- Management Ethernet link down alarm seen while verifying system alarms in Virtual Chassis setup. [PR1538674](#)
- Traffic loss might be seen in OVSDb VxLAN scenario. [PR1540208](#)
- Inter VLAN traffic drop might be observed in EVPN-VXLAN scenario. [PR1541406](#)
- On QFX10002-60C switch, the "show pfe filter" CLI command is unavailable. [PR1545019](#)
- The neighbor solicitation might be dropped from the peer device. [PR1550632](#)
- DHCP IPv6 is not working for QFX5110-48s-4c. [PR1551710](#)
- On QFX10000 and PTX10000 line of devices with Junos OS Releases 20.1R1 and later, cannot collect RSI properly because of the authentication error. [PR1556816](#)

Infrastructure

- The kernel might crash if a file or a directory is accessed for the first time and is not created locally. [PR1518898](#)
- 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](#)
- The logical interface might flap after adding or deleting native VLAN configuration. [PR1539991](#)

Layer 2 Features

- Flow control is enabled in Packet Forwarding Engine irrespective of the interface configuration and the fix causes small amount of packet loss when a parameter related to an interface such as "interface description" on any port is changed. [PR1496766](#)
- The dcpfe or FPC might crash generate a core file because of the memory leak after the VLAN add and VLAN delete operation. [PR1505239](#)
- 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 not synchronize between the master and the member in Virtual Chassis after the MAC flap. [PR1521324](#)
- On QFX5110 switch, the EVPN-VXLAN check traffic when vxlan encap header fails. [PR1541316](#)

Platform and Infrastructure

- On QFX5110 and QFX5120 platforms, unicast RPF check in strict mode might not work properly. [PR1417546](#)

Routing Policy and Firewall Filters

- The policy configuration might be mismatched between rpd and mgd when "deactivate policy-options prefix-list" is involved in configuration sequence. [PR1523891](#)

Routing Protocols

- System upgrade or installation might fail on QFX5100-48T-6Q VC/VCF. [PR1486632](#)
- The IPv6 traffic might drop when falling back from IP-in-IP tunnel to inet.0/inet6.0. [PR1508631](#)
- Scale of filters with egress-to-ingress command is enabled. [PR1514570](#)
- The rpd process might report 100 percent CPU usage with the BGP route damping enabled. [PR1514635](#)
- The remaining BFD sessions of the aggregated Ethernet interface flap continuously if one of the BFD sessions is deleted. [PR1516556](#)
- Stale tunnel entries are seen after negative triggers. [PR1516818](#)
- The BFD sessions might flap continuously after disruptive switchover followed by GRES. [PR1518106](#)
- On QFX5210-64C, 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](#)
- Errors are seen during script run with negative triggers at scale[LOG: Err]
BRMCM_NH-,brcm_nh_ipntunnel_unilist_install(),5752:IPoIP < src: 1.1.1.1, dst: 1.1.3.1> NH id 2537, Tunnel id: 512 failed to create decap obj Table full: vrf 1 vid 4082 intf 4058 of nh 131074(3)] when the tunnel color attribute is deleted for all the tunnels at scale. [PR1526405](#)

- On QFX5000 line of switches, the IPIP firewall filter term with decapsulate action need to be duplicated for each from protocol. [PR1527755](#)
- On QFX5000 line of switches, the fxpc process might crash if the VXLAN interface flap. [PR1528490](#)

Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. [PR1490552](#)
- On QFX5120 and QFX5210 platforms unexpected storm control events might occur. [PR1519893](#)

SEE ALSO

What's New 175
What's Changed 205
Known Limitations 207
Open Issues 210
Documentation Updates 219
Migration, Upgrade, and Downgrade Instructions 220

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for the QFX Series Switches.

SEE ALSO

What's New 175
What's Changed 205
Known Limitations 207
Open Issues 210
Resolved Issues 215
Migration, Upgrade, and Downgrade Instructions 220

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrading Software on QFX Series Switches | 220
- Installing the Software on QFX10002-60C Switches | 223
- Installing the Software on QFX10002 Switches | 223
- 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 | 224
- Installing the Software on QFX10008 and QFX10016 Switches | 226
- Performing a Unified ISSU | 230
- Preparing the Switch for Software Installation | 231
- Upgrading the Software Using Unified ISSU | 231
- Upgrade and Downgrade Support Policy for Junos OS Releases | 233

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

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://www.juniper.net/support/downloads/junos.html>.

The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **20.3** in the Release pull-down list to the right of the Software tab on the Download Software page.

4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.3 release.

An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

A login screen appears.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.

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

Customers in the United States and Canada use the following command:

```
user@host> request system software add
source/jinstall-host-qfx-5-x86-64-20.3-R1.n-secure-signed.tgz reboot
```

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

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 20.3 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.

NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

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

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-20.4R1.9.tgz
```

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

For example:

```
user@switch> request vmhost software add
ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-20.4R1.9.tgz
```

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

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

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R1.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

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

For example:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz reboot
```

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

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz reboot
```

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

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

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

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

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

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

To install the software on re0:

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

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

For example:

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

To install the software on re1:

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

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

For example:

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

Reboot both Routing Engines.

For example:

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

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

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

Installing the Software on QFX10008 and QFX10016 Switches

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

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



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

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

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

2. From the command line, enter configuration mode:

```
user@switch> configure
```

3. Disable Routing Engine redundancy:

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

4. Disable nonstop-bridging:

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

5. Save the configuration change on both Routing Engines:

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

6. Exit the CLI configuration mode:

```
user@switch# exit
```

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

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

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

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

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

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz
```

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

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

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

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

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

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

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

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

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

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

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

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

12. Transfer routing control to the backup Routing Engine:

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

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

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

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

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

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

```
user@switch> request system software add validate
/var/tmp/jinstall-host-qfx-10-f-x86-64-20.4R1.n-secure-signed.tgz
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Performing a Unified ISSU

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

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

Perform the following tasks:

- [Preparing the Switch for Software Installation on page 231](#)
- [Upgrading the Software Using Unified ISSU on page 231](#)

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-x86-64-20.4R1.n-secure-signed.tgz*.

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

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

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

```

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

```

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

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

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

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

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

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

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

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

SEE ALSO

[What's New | 175](#)

[What's Changed | 205](#)

[Known Limitations | 207](#)

[Open Issues | 210](#)

[Resolved Issues | 215](#)

[Documentation Updates | 219](#)

Junos OS Release Notes for SRX Series

IN THIS SECTION

- [What's New | 235](#)
- [What's Changed | 251](#)
- [Known Limitations | 256](#)
- [Open Issues | 258](#)
- [Resolved Issues | 260](#)
- [Documentation Updates | 264](#)
- [Migration, Upgrade, and Downgrade Instructions | 265](#)

These release notes accompany Junos OS Release 20.4R1 for the SRX Series Services Gateways. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- Application Layer Gateways (ALGs) | 236
- Application Security | 236
- ATP Cloud | 237
- Authentication and Access Control | 238
- Chassis Clustering | 238
- Flow-Based and Packet-Based Processing | 240
- Interfaces and Chassis | 241
- Intrusion Detection and Prevention | 242
- Juniper Extension Toolkit (JET) | 243
- Junos OS XML and API Scripting | 244
- J-Web | 244
- Layer 2 Features | 246
- Logical Systems and Tenant Systems | 246
- Multinode High Availability | 246
- Network Management and Monitoring | 247
- Securing GTP and SCTP Traffic | 248
- Security | 249
- Unified Threat Management (UTM) | 250
- VPNs | 250

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

Application Layer Gateways (ALGs)

- **SIP ALG load-balancing enhancement (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.4R1, you can configure two new CLI commands to distribute the SIP load among all the available SPU's to better balance SIP traffic. The new commands are **set security alg sip bulk-call-mode** and **set security alg sip enable-call-distribution**.

[See [bulk-call-mode](#) and [enable-call-distribution](#).]

Application Security

- **AppQoE support for SaaS applications (NFX Series and SRX Series)**—Starting in Junos OS Release 20.4R1, we've extended application quality of experience (AppQoE) support for Software as a Service (SaaS) applications.

AppQoE performs service-level agreement (SLA) measurements across the available WAN links such as underlay, GRE, IPsec or MPLS over GRE. It then sends SaaS application data over the most SLA-compliant link to provide a consistent service.

To configure AppQoE for SaaS applications:

1. Define the SLA rule type as SaaS (**set security advance-policy-based-routing sla-rule sla1 type saas**).
2. Include SaaS server details in the address book (**set security address-book global address *address-book* dns-name *saas-server-url* ipv4-only**).
3. Attach the SLA rule to the policy-based APBR profile.

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

- **SNI-based dynamic application information for SSL proxy profile (SRX Series)**—Starting in Junos OS Release 20.4R1, we've enhanced the selection mechanism for SSL proxy profiles. An SSL proxy can now use Server Name Indication (SNI) TLS extensions to identify dynamic applications.

SSL proxy defers SSL profile selection until the dynamic application is detected in a client hello message based on the SNI. Next, the SSL proxy does a firewall rule lookup based on the identified application and selects an appropriate SSL proxy profile.

Using SNI-based dynamic application information results in more accurate SSL proxy profile selection for the session. By default, this feature is enabled on SRX Series devices.

[See [Unified Policies for SSL Proxy](#) and [global-config \(Services\)](#).]

- **Granular control over DNS-over-HTTP and DNS-over-TLS application traffic (NFX Series, SRX Series and vSRX)**—In Junos OS Release 20.4R1, we introduce a new micro-application, DNS-ENCRYPTED, to enhance the application signature package. By configuring this micro-application in a security policy, you can have granular control for DNS-over-HTTP and DNS-over-TLS application traffic.

The DNS-ENCRYPTED application is enabled by default. You can disable it using the **request services application-identification application disable DNS-ENCRYPTED** command.

You can view the details of the micro-applications using the **show services show services application-identification application detail** command.

[See [Application Identification Support for Micro-Applications.](#)]

- **Support for tunneling applications in unified policies (SRX Series and vSRX)**—In Junos OS Release 20.4R1, we've enhanced unified policy functionality on security devices to manage tunneling applications. You can now block a specific tunneling application by using a unified policy.

For example, to block tunneling applications such as QUIC or SOCKS, you can configure a unified policy with the deny or reject action for these applications.

[See [Application Identification Support for Unified Policies.](#)]

ATP Cloud

- **Support for filtering DNS requests for disallowed domains (SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 20.4R1, you can configure DNS filtering to identify DNS requests for disallowed domains. You can either:
 - Block access to the disallowed domain by sending a DNS response that contains the IP address or fully qualified domain name (FQDN) of a DNS sinkhole server. This ensures that when the client attempts to send traffic to the disallowed domain, the traffic instead goes to the sinkhole server.
 - Log the DNS request and reject access.

The DNS sinkhole must be configured only for DNS profile category.

[See [dns-filtering](#), [security-intelligence](#), [clear services security-intelligence dns-statistics](#), and [show services security-intelligence dns-statistics](#).]

Authentication and Access Control

- **Logical domain support for device identity authentication (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.4R1, we've enabled logical system and tenant system support for policy search that uses device information from JIMS. The user firewall uses the logical system name or the tenant name as a differentiator. The logical system name or tenant name must be globally unique and consistent in the JIMS server and SRX Series device. The JIMS server forwards the differentiator to be included in the device identification authentication entries. The device authentication entries are distributed into the root logical system.

[See [Understanding Integrated User Firewall support in a Logical System](#), [end-user-profile](#).]

Chassis Clustering

- **Enabling and disabling control link (SRX1500)**—Starting in Junos OS Release 20.4R1, you can enable or disable the control links on SRX1500 using the operational and configuration mode CLI commands listed below.

In earlier Junos OS releases, if you wanted to disable the control links and fabric links, you had to manually unplug the cables for control links and fabric links, which was very inconvenient.

Using these commands helps you control the status of the cluster nodes and provides protection against version mismatch during a cluster upgrade, and minimizes failovers.

Table 4: Configuration and Operational Mode Commands

Configuration Mode	Operational Mode
<p>To disable the control link, run the set chassis cluster control-interface (node0 node1) disable command on node 0 or node 1.</p> <p>NOTE: If you disable the links using the configuration command, then the links remain disabled even after system reboot.</p>	<p>To disable the control link from the local node, run the request chassis cluster control-interface (node0 node1) disable command.</p> <p>NOTE: If you have disabled control link using the operational mode CLI commands, the links will be enabled after system reboot.</p>

Table 4: Configuration and Operational Mode Commands (*continued*)

Configuration Mode	Operational Mode
<p>To enable the control link, run the delete chassis cluster control-interface (node0 node1) disable on both the nodes.</p> <hr/> <p>You need to ensure that you disable the control and fabric interfaces on both nodes using CLI configuration, to keep control and fabric links disabled when you reboot the nodes.</p> <p>Use the set interfaces (fab0 fab1) disable and delete interfaces (fab0 fab1) disable CLI commands to disable or enable the fabric interfaces.</p>	<p>To enable the control link from the local node, run the request chassis cluster control-interface (node0 node1) enable command.</p>

[See [cluster \(Chassis\)](#).]

- **In-service software upgrade (ISSU) (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.4R1, you can use the new **status** option with the **request system software in-service-upgrade** command to display ISSU status during upgrade.

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

- **Support for single PSU operation without alarms (SRX1500 and SRX4600)**—Starting in Junos OS Release 20.4R1, a new argument, **pem-absence**, is available at the **[edit chassis alarm]** hierarchy level. You can use **set chassis alarm pem-absence ignore** to ignore the power supply unit (PSU) alarm. By default, the PSU alarm is raised when any PSU is missing or not energized.

[See [Understanding Chassis Alarms](#), [show chassis alarms](#), and [pem-absence](#).]

Flow-Based and Packet-Based Processing

- **Express Path+ support for packet-based services in traditional and unified policies (SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.4R1, Express Path+ (formerly known as services offloading) is supported for packet-based services in traditional and unified policies.

[Table 5 on page 240](#) lists the features that are supported by Express Path+.

Table 5: Features Supported by Express Path+

Supported Features
ALG
APBR (application based routing)
The Express Path+ works after APBR ignores subsequent traffic.
Web Filtering
GPRS (GTP and SCTP)
IDP
Screens
See Understanding Screens Options on SRX Series Devices .
Juniper Security Intelligence (SecIntel)
Unified Policies with ApplID and URL category matching
UserFW

[See [Express Path Overview](#).]

- **New resource-manager commands (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS 20.4R1 Release, you can configure memory load and CPU resources for the SRX5400 using three new **resource-manager** commands. The new commands **service-memory**, **session-memory**, and **cpu** are available at the **[edit security resource-manager]** hierarchy level.

[See [services-memory \(resource-manager\)](#), [session-memory \(resource-manager\)](#), [cpu \(resource-manager\)](#), [show resource-manager memory](#), [show resource-manager](#), [show resource-manager cpu](#), and [show security flow session](#).]

- **Support for trace and debug of data packets (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.4R1, you can trace packet footprints.

To enable tracing of packet footprints, use the **traceoptions flag jexec** command at the **[edit security flow]** or **[edit logical-systems *logical system name* security flow]** hierarchy level.

The packet trace logs are captured in a sequential time order. The sequential trace enhances flow debuggability for packet processing with multiple logical systems and tenant systems, tunnel piping, multiple reinjection, and so on.

[See [traceoptions \(Security Flow\)](#) and [show security flow status](#).]

- **Pass-through authentication of IP-IP and GRE tunnel traffic in TAP mode (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, and vSRX)**—Starting in Junos OS Release 20.4R1, SRX Series devices perform pass-through authentication of IP-IP and GRE tunnel traffic when in TAP mode. To use TAP mode, connect the SRX Series device to the mirror port of the connected switch, which provides a copy of the traffic traversing the switch. In TAP mode, the SRX Series device processes incoming traffic from the TAP interface and generates a security log or report containing with information about threats detected, application usage, and user details.

[See [Configure User Authentication Methods](#).]

- **Enhancement in Resource Management (SRX Series Devices)**—Starting in Junos OS Release 20.4R1, when the Layer7 packets such as ALG or User Firewall create flow sessions, you can control whether to drop the packet or forward the packet if resource is busy.

Configure the **security-service** under the **edit security forwarding-options** hierarchy to implement the resource management. When you configure the **security service** as **fail-open**, the session skips the application level and forwards the packet.

By default, the **security service** is **fail-closed**, and allows the session to drop at the application level. When you use **fail-closed** option, make sure **set security forwarding-options security-service fail-open** is not configured.

[See [Traffic Processing on SRX Series Devices Overview](#)]

- **Enhancement of Flow Reroute in Multiple Routing Table (SRX Series Devices)**—Starting in Junos OS Release 20.4R1, the flow reroutes the traffic using multiple routing table. Earlier to this release, flow reroute was supported with only one routing table.

If there is one routing table involved in route lookup, there is no change in the implementation of reroute. If there are more than one routing table involved in the route lookup and when there is a route change in any one of the routing table, you can mark all the affected flows for reroute. We support 16 routing tables.

[See [Traffic Processing on SRX Series Devices Overview](#)]

Interfaces and Chassis

- **CPU load monitoring (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.4R1, the reswatch process is used to monitor Routing Engine CPU load and Junos kernel usage.

[See [resource-watch](#).]

- **Support for captive portal on Wi-Fi Mini-Physical Interface Module (SRX320, SRX340, SRX345, SRX380, and SRX550HM)**—Starting in Junos OS Release 20.4R1, we support captive portal for integrated guest access management on the Wi-Fi Mini-PIM card. You can set up captive portal authentication to do either of the following:
 - Redirect Web browser requests to a login page that requires the user to input a simple password. You can also customize the login page display settings.
 - Use username and password authentication with the RADIUS server.

To configure captive portal, you must enable a DHCP server and configure two address pools for the Wi-Fi interface.

You can enable or disable captive portal on different virtual access points (VAPs), VAPs under the same access point use the same captive portal authentication type.

[See [Wi-Fi Mini-Physical Interface Module Overview](#).]

- **Support for Annex J and G.Fast with specialized SFP (SRX380, SRX300, SRX320, SRX340, and SRX345)**—Starting in Junos OS Release 20.4R1, we support G.Fast and Annex J specification with SFP xDSL for ADSL2/ADSL2+ and all VDSL2 profiles on SRX Series devices. Annex J is a specification in ITU-T recommendations G.992.3 and G.992.5 for all digital mode ADSL with improved spectral compatibility with ADSL over ISDN. You can configure Annex J by using the **dsl-sfp-options** cli command.

[See [Configuring ADSL Interfaces](#).]

Intrusion Detection and Prevention

- **IDP utility to read PCAP and generate protocol (SRX1500, SRX4200, SRX4600, vSRX, and vSRX3.0)**—Starting in Junos OS Release 20.4R1, you can use the pcap-analysis operational command to display the generated IDP context.

Execute the operational command **request security idp pcap-analysis /var/tmp/http.pcap from-zone trust to-zone untrust** by providing the zone details for selecting the corresponding interface. The output displays as a list of IDP contexts associated with the PCAP, and the data matched for that IDP context and attacks. However, configure the rulebase, IDP active policy, interfaces with necessary configuration like IPv4 addresses, zones, and policies for pcap utility to detect contexts and attacks.

You can process upto 3MB files and save a maximum length of 8K of context data to generate a unique context.

```
Attack Detected:  Yes
Rulebase          Rule Id      Attack Name
IPS-Rulebase-1    Rule-1      HTTP:AUDIT:URL
IDP context statistics:
```

Context name	#Hits	#Data
http-url	1	/
http-get-url	1	/
http-header-host	1	7.0.0.1
http-header-user-agent	1	lwp-request/5.827 libwww-perl/5.833
http-header	2	te: deflate,gzip;q=0.3 && connection: TE, close
http-request	1	GET / HTTP/1.1
http-request-method	1	GET

[See [IDP Utility for PCAP](#)]

Juniper Extension Toolkit (JET)

- **Juniper Extension Toolkit (JET) support for 64-bit applications (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX ELM, JunosV Firefly, cSRX, SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM, SRX650, SRX720E, SRX750E, SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SRX4400, SRX4600, SRX4800, SRX5400, SRX5600, SRX5800, SRX7X0E, SRX-ES7, SRX-ES8, VMX, and VSRX)**—Starting in Junos OS Release 20.4R1, JET supports 64-bit applications. Use the following commands to compile 64-bit applications for use with the AMD64 or ARM64 64-bit processor architecture.
 - **mk-amd64:** Compiles the application for use with AMD64 and Junos OS with FreeBSD.
 - **mk-amd64,bsd:** Compiles the application for use with AMD64 and Junos OS with upgraded FreeBSD.
 - **mk-arm64,bsd:** Compiles the application for use with ARM64 and Junos OS with upgraded FreeBSD.

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

Junos OS XML and API Scripting

- **Start time option for interval-based internal events that trigger event policies (EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—Starting in Junos OS Release 20.4R1, when you create an interval-based internal event for triggering event policies, you can specify the start date and time for the initial event. To specify a start time, configure the **start-time** option along with the **time-interval** option at the **[edit event-options generate-event]** hierarchy level.

[See [Generating Internal Events to Trigger Event Policies](#).]

J-Web

- **Enhanced Setup Wizard (SRX Series)**—Starting in Junos OS Release 20.4R1, we've refreshed the Setup Wizard settings for better experience. You can:
 - Access the wizard with the changed menu by selecting Device Administration>Reset Configuration.
 - View the new look of the wizard page.
 - Create one admin user with super user permissions.
 - Add and search NTP servers inline.
 - Synchronize device time with your computer time.
 - Configure interfaces directly under Zones & Interfaces instead of using a pop-up page.
 - Configure Tap settings with reorganized options.
 - View improved statements in the commit-success page.

In Standalone mode, the Setup Wizard does not support advanced services configuration.

[See [Start J-Web](#).]

- **Support for captive portal (SRX Series)**—Starting in Junos OS Release 20.4R1, you can configure captive portal while:
 - Creating a rule on the **Security Policies & Objects > Security Policies** page.
 - Adding a logical interface on the **Network > Connectivity > Ports** page. The IP address configured for web authentication is also used for captive portal.
 - Uploading the logo on the **Security Services > Firewall Authentication > Authentication Settings** page. The selected logo is also used for captive portal.

The captive portal authenticates the user request to access an SRX Series protected resource using HTTPS browser.

[See [Add a Rule](#), [Add a Logical Interface](#), and [About the Authentication Settings Page](#).]

- **Change in the Monitor tab menus (SRX Series)**—Starting in Junos OS Release 20.4R1, we've reorganized the Monitor tab into the following menus for enhanced experience:

- Interfaces
- Logs
- Maps and Charts
- Statistics
- Reports

Additionally:

- Support for Adobe Flash Player will end on December 31, 2020. Therefore, J-Web will support only the Monitor tab submenus that don't require flash components.
- A new Traffic Map page is added under Monitor>Maps and Charts. Use this page to visualize inbound and outbound traffic between geographic regions.

[See [Monitor Interfaces](#) and [Monitor Traffic Map](#).]

- **Enhanced Source NAT feature (SRX Series)**—Starting in Junos OS Release 20.4R1, we've refreshed the Source NAT page to improve user experience. You can configure:

- Source NAT inline on the NAT Policies page (Create>Source NAT).
- Source NAT pool and destination NAT pool on the NAT Pools page (Create>Source NAT Pool or Create>Destination NAT Pool).

[See [About the NAT Policies Page](#).]

Layer 2 Features

- **LLDP on routed and reth interfaces (SRX1500)**—Starting in Junos OS Release 20.4R1, Link Layer Discovery Protocol (LLDP) is supported on routed interfaces and redundant Ethernet (reth) interfaces. LLDP is a link-layer protocol used by network devices to advertise capabilities, identity, and other information to a LAN.

[See [LLDP Overview](#).]

Logical Systems and Tenant Systems

- **Support for MAP-E confidentiality CLI statement (NFX150, NFX250, NFX350, and SRX1500)**—Starting in Junos OS Release 20.4R1, we've introduced a global MAP-E **confidentiality** CLI statement to hide MAP-E rule parameters in CLI show commands and logs. To enable this configuration, include the **confidentiality** statement at the **[edit security softwires map-e]** hierarchy level. You need administrator privileges to enable or disable this configuration. This feature is supported for all domains of MAP-E.

[See [confidentiality](#) and [show security softwires map-e confidentiality status](#).]

Multinode High Availability

- **Multinode high availability solution (SRX5400, SRX5600, and SRX5800 with SPC3 card)**—In Junos OS Release 20.4R1, we introduce the multinode high availability solution, where two SRX Series devices can be either co-located or spread across geographies. This solution also provides redundancy across service levels.

When you configure multinode high availability in the SRX Series device (in the active-backup mode), one node acts as the active device and the other acts as the backup device, ensuring failover of services to the backup device in the event of software, hardware, or path monitoring failure. Traffic is then routed toward the active node by upstream and downstream routers.

The active and backup nodes are interconnected with an IP-based link called interchassis link (ICL). The active and backup nodes synchronize data plane based session states. They also synchronize control plane states for certain services.

[Table 6 on page 247](#) lists the multinode high availability features that we support.

Table 6: Feature Support for Multinode High Availability

Feature	Description
Active and Backup modes	<p>We've introduced active and backup states on SRX Series devices that operate in the multinode HA mode.</p> <p>[See Multinode High Availability.]</p>
IPsec VPN support	<p>IPsec feature is supported on multinode HA. IPsec runs actively on one node (or active node). It can failover to the secondary node (or backup node). IKE negotiations occur from the active node and the states are synchronized with the backup node. After synchronization, the backup node takes over the primary role and continues without bringing down the tunnels after switchover.</p> <p>You can run the show command(s) on both active and backup nodes to display the status of IKE and IPsec security associations. You can delete the IKE and IPsec security associations only on the active node.</p> <p>[See Support for VPN on HA Nodes in Multinode High Availability Solution.]</p> <p>When a packet enters a flow session which is on IPsec VPN and on the backup node, the packet is dropped. When the packet enters a clear text session, irrespective of the control plane transition state the clear text session moves to the active state. When you configure the set chassis high-availability services-redundancy-group 1 process-packet-on-backup command, the IPsec VPN related packet is not dropped in the backup node.</p>

Network Management and Monitoring

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The **<get-configuration>** operation supports the **compare="configuration-revision"** and **configuration-revision** attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

Securing GTP and SCTP Traffic

- **Support for listening mode, syslog identity information, and rate-limit configuration enhancement (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.4R1, we support:
 - **GTP listening mode**—The SRX Series firewall enables GPRS tunneling protocol (GTP) listening mode on the S5, S8, or S11 interface. In listening mode, the firewall can perform sanity check and stateful inspection on GTP packets per GTP configuration. To configure, include the **listening-mode** statement at the **[edit security gtp profile profile-name]** hierarchy level.
 - **Syslog identity information**—This feature adds information such as the user equipment (UE) IP address, International Mobile Station Identity (IMSI), and access point name (APN) to existing syslog messages during GTP management. This helps in mapping identities between a user ID and an IP address.
 - **rate-limit configuration enhancement**—We've added an alarm threshold option and a drop threshold option to the existing **rate-limit** configuration. This enhancement reduces the duplicate drop logs for the destination GPRS support node (GSN).

[See [listening-mode](#), and [rate-limit \(Security GTP\)](#).]

Security

- **Unified policies support for zone-context and global-level policies (SRX Series and vSRX)**—Starting in Junos OS Release 20.4R1, unified policies support both zone-context and global-level policies at the same time. In previous releases, unified policies supported only zone-context policies.

If there is any unified policy match, either in a zone-context or in a global context, then it is added to potential match list.

If there is no match in the zone-context, policy search occurs in the global context.

[See [Global Security Policies](#).]

- **Tunnel inspection for VXLAN passthrough (SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 20.4R1, you can allow L4 or L7 services to perform an inspection against the inner ethernet frame. VXLAN is one of the supported protocols and is designed to accommodate most overlay or underlay protocols which require inner inspection. VXLAN traffic is only inspected if there is a security policy configured to perform the inspection.

[See [tunnel-inspection](#) and [show security flow session](#).]

- **Security policy support for security inspection on VXLAN tunnels (SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 20.4R1, you can perform security inspection on VXLAN tunnels by performing policy control twice. Configure an outer policy for the outer header and an inner policy for the inner header.

Configure a tunnel inspection profile to connect the outer policy and inner policy. The tunnel inspection profile is attached to the outer policy and it points to a group of inner policies (policy set). When the packet matches the outer policy, the SRX device decapsulates the packet to get the inner header. Using inner packet content along with the attached tunnel inspection profile of outer policy, the second policy lookup gets the desired inner policy applies the security services to inner packet.

[See [tunnel-inspection](#).]

- **Support for unidirectional session refreshing (SRX Series)**—Starting in Junos OS Release 20.4R1, SRX Series device support unidirectional session refreshing. You can do either of the following:
 - Refresh a session by any packet from any direction. This is an existing session-refreshing mechanism and the default behavior.
 - Refresh a session by only the packets in the initial direction (unidirectional refreshing).

By default, unidirectional session refreshing is disabled. To enable the feature, include the **unidirectional-session-refreshing** statement at the **[edit security zones security-zone zone-name]** hierarchy level.

[See [unidirectional-session-refreshing](#).]

Unified Threat Management (UTM)

- **Custom response page in UTM Web filtering profile (SRX Series and vSRX)**—Starting in Junos OS Release 20.4R1, you can configure a custom response page for a URL that is configured with the block or quarantine actions in the UTM Web filtering profile. The custom response page can include predefined page variables, your corporate branding, acceptable use policies, and links to your internal resources. You can enable the **custom-page** statement at the **[edit security utm custom-objects custom-message *name type*]** hierarchy level and configure the customized HTML file at the **[edit security utm custom-objects custom-message *name custom-page-file file-name*]** hierarchy level.

[See [custom-page](#), [custom-page-file](#), and [custom-message \(Security Web Filtering\)](#).]

- **URL pattern wildcard enhancement (SRX Series and vSRX)**—Starting in Junos OS Release 20.4R1, the URL pattern supports new regular expressions and defines new pattern matching rules for the domain name and URL path. This enhancement allows you to configure better and user-friendly URL pattern matching in the Web filtering function. You can use the asterisk (*), caret (^), and question mark (?) wildcards for a domain name match. The URL match supports the prefix match and keyword match. You can use the asterisk (*) wildcard for the URL match.

[See [url-pattern](#).]

- **Dynamic-address group rescan enhancement (SRX Series and vSRX)**—In the current dynamic-address implementation, when you add a host address to the dynamic-address group, the system does not terminate and rescan the existing sessions of the host. Starting in Junos OS Release 20.4R1, when you add a host address to the dynamic-address group, the system rescans the sessions including the existing sessions to ensure that the traffic matches the updated policy. The **session-scan** option is disabled by default. You can enable the **session-scan** option at the **[edit security dynamic-address address-name *name session-scan*]** or **[edit security dynamic-address session-scan]** hierarchy level.

[See [session-scan](#) and [hold-interval](#).]

VPNs

- **Support for load redistribution (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 20.4R1, you can redistribute or migrate a tunnel that belongs to a site-to-site or Auto VPN gateway to a new processing unit. Load distribution of tunnels tears down the tunnels associated with a peer and anchors them to the user-specified processing unit the next time a tunnel for the same peer is established. Use the **request security re-distribution ipsec-vpn** command to specify the processing unit to which you want to distribute or migrate the tunnels.

[See [request security re-distribution ipsec-vpn](#), [show security re-distribution ipsec-vpn](#), and [show security ipsec tunnel-distribution](#).]

- **PowerMode IPsec support (SRX4600)**—PowerMode IPsec (PMI) is supported on SRX5400, SRX5600, SRX5800, and vSRX. Starting in Junos OS Release 20.4R1, you can configure PMI on the SRX4600 also. PMI provides IPsec performance improvements using Vector Packet Processing (VPP) and Intel Advanced

Encryption Standard New Instructions (AES-NI). SRX4600 also supports the fat core feature that increases single tunnel performance. If one of the tunnels loaded with high traffic and other tunnels have less traffic, the resources share within the fat group and result in an even CPU utilization of the resources.

[See [Improving IPsec Performance with PowerMode IPsec](#) and [power-mode-ipsec](#).]

- **PowerMode IPsec performance improvement (SRX5400, SRX5600, and SRX5800 with SPC3 cards)**—Starting in Junos OS Release 20.4R1, we've improved PowerMode IPsec (PMI) performance by distributing load between the AES-NI instructions on the SPUs and the on-board Intel QuickAssist Technology (QAT), Hardware-based cryptographic acceleration for symmetric fat tunnels in SPC3 cards provides higher performance. Load balancing helps to provide higher throughput for IPsec. PMI uses AES-NI and QAT for encryption and FPGA for decryption of cryptographic operation. To enable QAT with AES-NI, include **power-mode-ipsec-qat** at the **[edit security flow]** hierarchy level.

[See [power-mode-ipsec-qat](#) and [inline-fpga-crypto](#).]

SEE ALSO

What's Changed 251
Known Limitations 256
Open Issues 258
Resolved Issues 260
Documentation Updates 264
Migration, Upgrade, and Downgrade Instructions 265

What's Changed

IN THIS SECTION

- [Class of Service \(CoS\) | 252](#)
- [Flow-Based and Packet-Based Processing | 252](#)
- [Intrusion Detection and Prevention \(IDP\) | 253](#)
- [Interfaces and Chassis | 253](#)
- [J-Web | 253](#)
- [Network Address Translation \(NAT\) | 254](#)
- [Network Management and Monitoring | 254](#)
- [Platform and Infrastructure | 254](#)

- [Securing GTP and SCTP Traffic | 254](#)
- [User Interface and Configuration | 255](#)
- [VPNs | 255](#)

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

Class of Service (CoS)

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

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.

[See [Understanding Session Cache](#).]

- **Default MKA transmit interval (SRX380)**—On SRX380 devices, the default MACsec Key Agreement (MKA) transmit interval is 2000 milliseconds. If you deploy an SRX380 device with other security peer device with MACsec secure link, you must change the MKA transmit interval on the peer device to 2000 milliseconds to match the new default MKA transmit interval of the SRX380 device.

[See [transmit-interval \(MACsec\)](#)`transmit-interval (MACsec)`.]

Intrusion Detection and Prevention (IDP)

- **Intelligent Offload State (SRX Series)**— We have introduced a new field in the **show security idp status** command to see the status of the IDP Intelligent offload.

[See [show security idp status](#).]

Interfaces and Chassis

- **g mode supported on radio 2.4GHz of Wi-Fi MPIM (SRX320, SRX340, SRX345, and SRX550M)**—Starting in Junos OS Release 20.4R1, radio 2 with frequency 2.4 GHz supports mode g on SRX Wi-Fi MPIM.

[See [Wi-Fi Mini Physical Interface Module \(MPIM\)](#)].

J-Web

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

- **Change in the J-Web browser tab title (SRX Series)**—The J-Web browser tab title displays the device model and hostname. These details are also displayed when you hover over the J-Web browser tab.

For example, when you access J-Web for an SRX320 device with the hostname srx320-xyz, the J-Web browser tab displays the title as *J-Web (srx320 - srx320-xyz)*.

If the hostname isn't configured, the J-Web browser tab title displays the host URL or IP address; for example, *J-Web (srx320 - <device IP address>)*.

Network Address Translation (NAT)

- **Port block allocation support (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600)**—You can configure the port block allocation size 1 through 64512. To save system memory, the recommended port block allocation size is 64. If you configure the port block allocation 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 the earlier releases, you can configure port block allocation size 1 through 64512 on SRX5400, SRX5600, and SRX5800 only.

[See [Configure Port Block Allocation Size](#).]

Network Management and Monitoring

- **Warning changed for configuration statements that correspond to deviate not-supported nodes in YANG data models (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—If you configure a statement corresponding to a YANG data model node that defines the **deviate not-supported** statement, the Junos OS configuration annotates that statement with the comment **Warning: statement ignored: unsupported platform**. In earlier releases, the warning is **Warning: 'statement' is deprecated**.

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](#).]

Securing GTP and SCTP Traffic

- **Deprecated CLI configuration statements and operational commands for GTP and SCTP (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.4R1, we've removed the term `gprs` from all the configuration statements and operational commands for GTP and SCTP. As a part of this change:
 - We've deprecated the `[set security gprs]` hierarchy-level and all the configuration options under this hierarchy-level.
 - All the configuration statements previously available under the `[set security gprs gtp]` hierarchy-level are now available under the `[set security gtp]` hierarchy-level.
 - All the configuration statements previously available under the `[set security gprs sctp]` hierarchy-level are now available under the `[set security sctp]` hierarchy-level.
 - Replace the `show security gprs gtp configuration` command by `show security gtp profile` command.

- Replace the **identifier** option by **profile-name** in the **show security gtp profile** command.
- For default applications like **junos-gprs-gtp** and **junos-gprs-sctp**, you need not remove the term **gprs**.

[See [Configuration Statements and Operational Commands](#).]

User Interface and Configuration

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

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

VPNs

- **Dynamic CA profiles loaded only on active nodes (SRX5400, SRX5600, and SRX5800)**—When you enable the multinode high availability feature, the dynamic CA profiles are loaded only on the node during the IKE negotiation. If a failover occurs, the new active node undergoes a new IKE negotiation and loads the dynamic CA certificates as part of that negotiation. When PKID restarts, dynamic CA certificates are deleted only from the node where PKID was restarted.
- **Public key infrastructure warning message (SRX5400, SRX5600, SRX5800)**—When you generate a public key infrastructure (PKI) public/private key pair for a local digital certificate, with key pair size of 4096 bits and DSA encryption, a warning message is displayed. **root@hostname> request security pki generate-key-pair certificate-id test type dsa size 4096** **Generating a key-pair with a large modulus is very time-consuming. Progress is reported to the trace log, and a log message is generated upon completion.** Because generating a local digital certificate with large key pair size is time consuming, we recommend you to check the trace log for the progress of generating a key pair.

[See [request security pki generate-key-pair \(Security\)](#).request security pki generate-key-pair (Security).]

- **Delay in VPN tunnel establishment negotiation (SRX5400, SRX5600, and SRX5800)** —In an IPsec VPN configuration, if you configure the **establish-tunnels immediately** option under the **[edit security ipsec vpn <vpn-name>]** hierarchy, it may take up to five seconds to start the negotiation for VPN tunnel establishment. In the earlier Junos OS releases, the negotiation for VPN tunnel establishment starts immediately.

[See [vpn \(Security\)](#).]

- **The junos-ike package installed by default (SRX5000 line of devices)**— For the SRX5000 line of devices with RE3 installed, the junos-ike package is installed by default. As a result, the iked and ikemd processes

run on the Routing Engine by default instead of the IPsec key management daemon (kmd). In earlier Junos OS releases, the `junos-ike` package is an optional package for SRX5000 devices with RE3, and IPsec Key Management Daemon (KMD) runs by default.

[See [Enabling IPsec VPN Feature Set on SRX5K-SPC3 Services Processing Card.](#)]

- **IKE index displayed in `show security ipsec security-associations detail` output (SRX5400, SRX5600, and SRX5800)**—When you execute the `show security ipsec security-associations detail` command, a new output field, **IKE SA Index**, corresponding to every IPsec Security Association (SA) within a tunnel is displayed under each IPsec SA information.

[See [show security ipsec security-associations.](#)]

SEE ALSO

[What's New | 235](#)

[Known Limitations | 256](#)

[Open Issues | 258](#)

[Resolved Issues | 260](#)

[Documentation Updates | 264](#)

[Migration, Upgrade, and Downgrade Instructions | 265](#)

Known Limitations

IN THIS SECTION

- [Class of Service \(CoS\) | 257](#)
- [Flow-Based and Packet-Based Processing | 257](#)
- [J-Web | 257](#)
- [VPNs | 258](#)

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.

Class of Service (CoS)

- On SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500 devices, if CoS code points with values other than 000 are configured, packet loss might be seen for certain traffic patterns due to the firewall not having enough buffer. [PR1544709](#)

Flow-Based and Packet-Based Processing

- For accelerated flows such as Express Path, the packet or byte counters in the session close log and show session output only take into account the values that accumulated while traversing the NP. [PR1546430](#)

J-Web

- The Threat Map page available under the Monitor menu does not support a redirection to Logs page from top destination countries, top source countries, and specific country threat count [PR1542392](#)
- In the NAT Policies page:
 - White space when reordering multiple rules: When you try to reorder more than one rule at a time using drag and drop, a white space is seen in the rule dropped area. It will disappear when the grid is scrolled up or down.
 - Hit count value zero when searching for a rule: The hit count field in the search results will always show zero when searching for rules. But in the normal grid view when not performing the search operation, the Hit count is updated correctly.
 - Policy grid is greyed out while adding a new rule with an existing context or ruleset expanded and scrolled down almost more than 20 rules. The page loads normally if you refresh the menu.
 - Select all checkbox and delete: When Select all checkbox is used repeatedly or when deleting multiple rules using Select all checkbox, you may receive a browser warning: A Web page is slowing down your browser.

[PR1558757](#)

- The Firefox browser displays an unsaved changes error message in the J-Web Basic Settings page if the Autofill logins and passwords option is selected under the browser Privacy and security settings. [PR1560549](#)

VPNs

- 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](#)

SEE ALSO

[What's New | 235](#)

[What's Changed | 251](#)

[Open Issues | 258](#)

[Resolved Issues | 260](#)

[Documentation Updates | 264](#)

[Migration, Upgrade, and Downgrade Instructions | 265](#)

Open Issues

IN THIS SECTION

- [Flow-Based Packet-Based Processing | 259](#)
- [Interfaces and Chassis | 259](#)
- [J-Web | 259](#)
- [Protocols | 259](#)
- [Routing Policy and Firewall Filters | 259](#)
- [VPNs | 260](#)

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 Packet-Based Processing

- The LLDP protocol can be configured on SRX4000 and SRX5000 lines of devices, which it is not actually supported on those platforms. [PR1540797](#)

Interfaces and Chassis

- On SRX4100 and SRX 4200 devices, if PEM0 is removed, the output of jnxOperatingDescr.2 command might be incomplete. [PR1547053](#)

J-Web

- J-Web GUI does not allow you to save a rule if the cumulative shared objects are more than 2500 before the policy grid is saved. When there are several shared objects, there will be a noticeable delay in opening sources and destinations of a rule, and performing rule action. [PR1540047](#)
- When the commit pending changes message is shown on the J-Web UI, the contents of other messages, landing page, or pop-ups will not be clearly visible. [PR1554024](#)
- Sometimes after a longer usage, you cannot log in to J-Web again. [PR1561930](#)

Protocols

- The LLDP protocol can be configured on SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 devices which is not supported. [PR1540797](#)

Routing Policy and Firewall Filters

- On SRX Series devices, a higher CPU utilization than normal might be observed, which might cause performance to decline rapidly if global policies are used and zones are declared explicitly in those policies. [PR1549366](#)
- On SRX5000 line of devices, the secondary node might get stuck in performing ColdSync after a reboot. [PR1558382](#)
- On SRX Series devices, when inserting one global policy (including adding, deleting or reordering a policy) above others, swapping policies will happen on the global policies after the inserted policy. At this time, the swapped global policies might not be found during the first path search. In this case, the traffic used to initiate a session creation that matched these undetected policies might be dropped, but the retransmission packets will pass successfully. [PR1558827](#)

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, during in-service software upgrade (ISSU), the IPsec tunnels flap, causing a disruption of traffic. The IPsec tunnels recover automatically after the ISSU process is completed. [PR1416334](#)
- IPsec VPN flaps if more than 500 IPsec VPN tunnels are connected for the first time. [PR1455951](#)
- 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](#)
- On SRX5000 line of devices with SPC3 card, when the encryption algorithm is not configured in IPsec proposal, the output of `show security ipsec security-associations` command might display empty space instead of keyword null for encryption algorithm. [PR1507270](#)
- In multinode high availability, if the link encryption tunnel fails to get established after some attempts, IKED process generates core files. [PR1559121](#)

SEE ALSO

[What's New | 235](#)

[What's Changed | 251](#)

[Known Limitations | 256](#)

[Resolved Issues | 260](#)

[Documentation Updates | 264](#)

[Migration, Upgrade, and Downgrade Instructions | 265](#)

Resolved Issues

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

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

Application Layer Gateways (ALGs)

- The SCCP ALG does not work on SRX Series devices running with Junos OS Release 17.3R1 and later. [PR1535356](#)

Flow-Based and Packet-Based Processing

- CLI autocomplete is now available for both SecIntel and advanced anti-malware products. [PR1487419](#)
- A condition within TCP proxy could result in downloads becoming permanently stuck or not completing. TCP proxy is used by multiple services, including Juniper ATP Cloud in block mode, ICAP, SSL proxy, antivirus, content filtering, and antispam. [PR1502977](#)
- In a dual CPE scenario, if the rule match is completed before application identification is done, AppQoE moves the session to the other node. [PR1514973](#)
- VRRP does not work on the redundant Ethernet interface with a VLAN ID greater than 1023. [PR1515046](#)
- PCAP file generated using packet capture was improper on the SRX5000 line of devices. [PR1515691](#)
- A logic issue was corrected in SSL proxy that could lead to an srpxfe or flowd core file under load. [PR1516903](#)
- The PPPoE session does not come up after return to zero on SRX Series devices. [PR1518709](#)
- FQDN-based security log stream does not dynamically update the IP address. [PR1520071](#)
- TAP mode behavior has been improved and the configuration has been greatly simplified. [PR1521066](#)
- The TCP packet might be dropped if syn-proxy protection is enabled. [PR1521325](#)
- Hide routing-instance under edit system name-server for SRX Series devices starts from Junos OS Release 20.4. [PR1521666](#)
- On SRX Series devices with chassis cluster, high CPU usage might be seen due to the llmd process. [PR1521794](#)
- Adaptive threat profiling would stop submitting new IP addresses to the feed after a limit of 10,000 has been reached. [PR1524284](#)
- 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 SRX Series devices, a node of chassis cluster might stop passing traffic. The traffic forwarding can be restored by a manual failover to node 1. [PR1528898](#)
- When no LSYS or TSYS flow trace is configured and no root-override is configured, the latest behavior is to not log any flow trace for that LSYS or TSYS, instead of dumping all to root flow trace as before. [PR1530904](#)
- 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 rst-invalidate-session command does not work if configured together with the no-sequence-check command. [PR1541954](#)
- The nsd process might crash when DNS-based allowlisting is configured under SSL proxy. [PR1542942](#)
- Application fragmented traffic might get dropped on SRX Series devices. [PR1543044](#)
- The Wi-Fi Mini-Physical Interface Module (Mini-PIM) does not support pure g mode with radio 2.4 GHz. [PR1543824](#)
- 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](#)

Interfaces and Chassis

- Fabric interface might be monitored down after chassis cluster reboot. [PR1503075](#)
- On SRX320, SRX340, SRX345, SRX380, and SRX550M devices with an LTE Mini-Physical Interface Module (Mini-PIM), the LTE connection might drop and fail to automatically recover because of firmware issue. [PR1520879](#)
- When SRX Series devices receive proxy ARP requests on VRRP interfaces, SRX Series devices send ARP replies with the underlying interface MAC address. [PR1526851](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srpxfe process might generate core files during the idpd process commit on SRX Series devices. [PR1521682](#)
- Adaptive threat profiling incorrectly classifies hosts when Server-to-Client (S2C) IDP signatures are used. [PR1533116](#)
- SOF support for partial packet plugins on traditional or unified policy. [PR1542497](#)

J-Web

- The parameters show another LSYS at J-Web in a multiple LSYS scenario. [PR1518675](#)
- Sometimes, when you edit the local gateway in the remote access VPN workflow under VPN>IPsec VPN, J-web might not display one or more drop-down values. [PR1521788](#)

- In the SRX5000 line of devices, J-Web can take up to 60 seconds to 90 seconds to load 60000 security policies. [PR1521841](#)

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

- 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](#)

Routing Policy and Firewall Filters

- Junos OS upgrade may encounter failure in certain conditions when enabling ATP. [PR1519222](#)
- Traffic might be dropped when policies are changed in SRX Series devices. [PR1527570](#)
- The show security dynamic-address feed-name command could not list secprofiling feed. [PR1537714](#)
- The flowd or srpxfe process might crash when an SRX Series or NFX Series device running Junos OS Release 18.2R1 or later supports the unified policy feature. [PR1544554](#)
- Traffic might be dropped unexpectedly when the URL category match condition is used on a security policy. [PR1546120](#)
- NSD process stops when the secprofiling feed name is 64 bytes. [PR1549676](#)

Routing Protocols

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

Subscriber Access Management

- Incorrect counter type (counter instead of gauge) specified for some values in MIB jnxUserAAAMib. [PR1533900](#)

Unified Threat Management (UTM)

- UTM causes e-mails from outside to inside to not be received. [PR1523222](#)
- Stream buffer memory leak might happen when UTM is configured under unified policies. [PR1557278](#)

VPNs

- The IKE tunnel negotiation might fail if IKE_INIT request is re-transmitted. [PR1460907](#)
- IPsec traffic may get dropped after RGO failover. [PR1522931](#)
- On all SRX Series devices using IPsec with NAT traversal, MTU size for the external interface might be changed after IPsec SA is re-established. [PR1530684](#)
- After IPsec tunnel using policy-based VPN is overwritten by another VPN client, traffic using this IPsec tunnel will be dropped. [PR1546537](#)

SEE ALSO

[What's New | 235](#)

[What's Changed | 251](#)

[Known Limitations | 256](#)

[Open Issues | 258](#)

[Documentation Updates | 264](#)

[Migration, Upgrade, and Downgrade Instructions | 265](#)

Documentation Updates

There are no errata or changes in Junos OS Release 20.4R1 documentation for the SRX Series.

SEE ALSO

[What's New | 235](#)[What's Changed | 251](#)[Known Limitations | 256](#)[Open Issues | 258](#)[Resolved Issues | 260](#)[Migration, Upgrade, and Downgrade Instructions | 265](#)

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 20.2 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 20.2.

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 | 235](#)[What's Changed | 251](#)[Known Limitations | 256](#)[Open Issues | 258](#)[Resolved Issues | 260](#)[Documentation Updates | 264](#)

Junos OS Release Notes for vMX

IN THIS SECTION

- [What's New | 267](#)
- [What's Changed | 269](#)
- [Known Limitations | 270](#)
- [Open Issues | 270](#)
- [Resolved Issues | 270](#)
- [Licensing | 271](#)
- [Upgrade Instructions | 271](#)

These release notes accompany Junos OS Release 20.4R1 for vMX. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [EVPN | 267](#)
- [Juniper Extension Toolkit \(JET\) | 267](#)
- [Junos OS XML ,API, and Scripting | 268](#)
- [Network Management and Monitoring | 268](#)
- [Routing Protocols | 269](#)

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

EVPN

- **MC-LAG emulation in an EVPN deployment (EX Series, MX Series, and vMX)**—Starting in Junos OS Release 20.4R1, you can emulate the function of an MC-LAG in active-standby mode in an EVPN configuration without having to configure an ICCP or ICL interface. In a standard EVPN configuration, logical interfaces configured on an aggregated Ethernet interface can have different designated forwarder election roles. To emulate an MC-LAG configuration, the designated forwarder (DF) takes on the role of the aggregated Ethernet interface. The provider edge (PE) that is the non-DF will send LACP out-of-sync packets to the CE. This causes LACP to go down on the CE device, and the CE device does not use the links connected to the non-DF for sending traffic. If the connection between a CE and a DF PE fails, the PE is re-elected as a DF. If the connection between a CE and a non-DF PE fails, the current DF PE is not changed.

To enable this functionality, configure the **lACP-oos-on-ndf** statement at the **[edit interfaces interface name esi df-election-granularity per-esi]** hierarchy.

Juniper Extension Toolkit (JET)

- **Juniper Extension Toolkit (JET) support for 64-bit applications (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX ELM, JunosV Firefly, cSRX, SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM, SRX650, SRX720E, SRX750E, SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SRX4400, SRX4600, SRX4800, SRX5400, SRX5600, SRX5800, SRX7X0E, SRX-ES7, SRX-ES8, VMX, and VSRX)**—Starting in Junos OS Release 20.4R1, JET supports 64-bit applications. Use the following commands to compile 64-bit applications for use with the AMD64 or ARM64 64-bit processor architecture.

- **mk-amd64:** Compiles the application for use with AMD64 and Junos OS with FreeBSD.
- **mk-amd64,bsd:** Compiles the application for use with AMD64 and Junos OS with upgraded FreeBSD.
- **mk-arm64,bsd:** Compiles the application for use with ARM64 and Junos OS with upgraded FreeBSD.

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

- **Configure inner source MAC address for flexible VXLAN tunnels (MX Series and vMX with MPC1-MPC9E or LC2101)**—Starting in Junos OS Release 20.4R1, you can use the Juniper Extension Toolkit (JET) RIB Service API to configure the source MAC address used in IPv4 and IPv6 flexible VXLAN tunnel encapsulation profiles. The source MAC addresses is stored in the inner Ethernet header of VXLAN encapsulation. If you don't specify a source MAC address, the default source MAC address 00:00:5e:00:52:01 is used to encapsulate IPv4 and IPv6 flexible VXLAN tunnels.

Use the **show route detail**, **show route extensive**, and **show flexible-tunnels profiles** CLI commands or the **get-route-information** and **get-flexible-tunnels-profiles** RPC/NETCONF commands to view the source MAC address that is specified in the flexible tunnel profile.

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

Junos OS XML ,API, and Scripting

- **Support for Certificate Authority Chain Profile (EX2300, EX3400, EX4300, MX240, MX480, MX960, PTX-5000, VMX, vSRX and QFX5200)**—Starting in Junos OS Release 20.4R1, you can configure intermediate Certificate Authority (CA) chain profile certificate and perform https REST API request using mutual and server authentications.

To configure intermediate ca-chain certificate, configure **ca-chain ca-chain** statement at the **[edit system services rest https]** hierarchy level.

Network Management and Monitoring

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with

the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The `<get-configuration>` operation supports the `compare="configuration-revision"` and `configuration-revision` attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

Routing Protocols

- **Support for multiple single-hop EBGP sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGP session. You can now enable single-hop EBGP sessions on different links over multiple directly connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGP Sessions on Different Links Using the Same Link-Local Address \(IPv6\)](#).]

What's Changed

IN THIS SECTION

- [Licensing](#) | 270

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

Licensing

- **Change in license bandwidth command on vMX virtual routers**—Starting in Junos OS, to use the available license bandwidth, explicitly set the license bandwidth use the `set chassis license bandwidth <In Mbps>` command

[See [Configuring Licenses on vMX Virtual Routers](#).]

Known Limitations

There are no known behaviors and limitations for vMX in Junos OS Release 20.4R1.

Open Issues

There are no open issues for vMX in Junos OS Release 20.4R1.

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.

Interfaces and Chassis

- Random packet drop with flow cache disabled, when NIC mapped to NUMA node 1. [PR1458742](#)
- On vMX instances, configuring the statement ranges for autosensed VLANs (either stacked VLANs or single-tag VLANs) might not work. This is because the VLANs are not programmed on the NIC drivers. [PR1503538](#)
- 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](#)

Network Management and Monitoring

- After l2cpd service is restarted, the context of registration from l2cpd to snmpd was failing due to incorrect reinitialization. Because of this, if an NMS polls the dot1dStp objects by prefixing the context might fail. As a workaround, restart snmpd or reconfigure protocols hierarchy. [PR1561736](#)

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 the 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 Release 19.2R1 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 | 272](#)
- [What's Changed | 273](#)
- [Known Limitations | 273](#)
- [Open Issues | 273](#)
- [Resolved Issues | 274](#)

These release notes accompany Junos OS Release 20.4R1 for vRR. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [Routing Protocols | 273](#)

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

To learn about common BGP or routing Junos OS features supported on vRR for Junos OS Release 20.4R1, see [What's New](#) for MX Series routers.

Routing Protocols

- **Support for relaxing BGP router ID format from /32 to a nonzero ID per RFC 6286 (MX204, NFX Series, PTX5000, QFX Series, and vRR)**—Starting in Junos OS Release 20.4R1, you can establish a BGP connection using a BGP identifier that is a 4-octet, unsigned, nonzero integer and it needs to be unique only within the autonomous system (AS) per RFC 6286. In earlier releases, the BGP ID of a BGP speaker was required to be a valid IPv4 host address assigned to the BGP speaker.

To enable this feature, use the **bgp-identifier identifier group bgp group name bgp-identifier identifier neighbor peer address bgp-identifier identifier** configuration statement at the **[edit protocols bgp]** hierarchy level.

[See [router-id](#)]

What's Changed

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

To learn more about common BGP or routing changes in behavior or syntax in Junos OS Release 20.4R1, 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 Release 20.4R1, 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

Learn about open issues in this release for vRR.

To learn more about common BGP or routing open issues in Junos OS Release 20.4R1, see [Open Issues](#) for MX Series routers.

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

Resolved Issues

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

To learn more about common BGP or routing resolved issues in Junos OS Release 20.4R1, see [Resolved Issues](#) for MX Series routers.

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

Junos OS Release Notes for vSRX

IN THIS SECTION

- [What's New | 274](#)
- [What's Changed | 278](#)
- [Known Limitations | 279](#)
- [Open Issues | 279](#)
- [Resolved Issues | 280](#)
- [Migration, Upgrade, and Downgrade Instructions | 282](#)

These release notes accompany Junos OS Release 20.4R1 for vSRX. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

What's New

IN THIS SECTION

- [ATP Cloud | 275](#)
- [Flow-Based Packet-Based Processing | 275](#)

- High Availability | 276
- Juniper Extension Toolkit (JET) | 276
- Junos OS XML ,API, and Scripting | 277
- Network Management and Monitoring | 277
- Platform and Infrastructure | 278
- Routing Protocols | 278
- VPNs | 278

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

ATP Cloud

- **Support for filtering DNS requests for disallowed domains (SRX4100, SRX4200, SRX4600, and vSRX)**—Starting in Junos OS Release 20.4R1, you can configure DNS filtering to identify DNS requests for disallowed domains. You can either:
 - Block access to the disallowed domain by sending a DNS response that contains the IP address or fully qualified domain name (FQDN) of a DNS sinkhole server. This ensures that when the client attempts to send traffic to the disallowed domain, the traffic instead goes to the sinkhole server.
 - Log the DNS request and reject access.

The DNS sinkhole must be configured only for DNS profile category.

[See [dns-filtering](#), [security-intelligence](#), [clear services security-intelligence dns-statistics](#), and [show services security-intelligence dns-statistics](#).]

Flow-Based Packet-Based Processing

- **Pass-through authentication of IP-IP and GRE tunnel traffic in TAP mode (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550 HM, SRX1500, and vSRX)**—Starting in Junos OS Release 20.4R1, SRX Series devices perform pass-through authentication of IP-IP and GRE tunnel traffic when in TAP mode. To use TAP mode, connect the SRX Series device to the mirror port of the connected switch, which provides a copy of the traffic traversing the switch. In TAP mode, the SRX Series device processes incoming traffic from the TAP interface and generates a security log or report containing with information about threats detected, application usage, and user details.

[See [Configure User Authentication Methods](#).]

- **Support for trace and debug of data packets (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 20.4R1, you can trace packet footprints. To enable tracing of packet footprints, use the **traceoptions flag jexec** command at the **[edit security flow]** or **[edit logical-systems *logical system name* security flow]** hierarchy level.

The packet trace logs are captured in a sequential time order. The sequential trace enhances flow debuggability for packet processing with multiple logical systems and tenant systems, tunnel piping, multiple reinjection, and so on.

[See [traceoptions \(Security Flow\)](#) and [show security flow status](#).]

High Availability

- **SR-IOV 10GbE high availability support (vSRX 3.0)**—Starting in Junos OS Release 20.4R1, vSRX 3.0 supports high availability (HA) single-root I/O virtualization (SR-IOV) deployment.

If you have a physical network interface card (NIC) that supports SR-IOV, you can attach SR-IOV-enabled vNICs or virtual functions to the vSRX 3.0 instance.

With this feature, you can access the hardware directly from a virtual machines environment and efficiently share the PCIe devices to optimize performance and capacity. Also, this feature allows you to create many VFs associated with a single physical function (PF) extending the capacity of a device and lowering hardware costs.

We recommend that you configure all revenue ports of vSRX 3.0 as SR-IOV. On KVM, you can configure SR-IOV high availability on management port: -fxp0/ control port- em0 / fabric port-ge-0/0/*.

SR-IOV high availability Layer 2 function is not supported. Also, SR-IOV high availability with the vSRX 3.0 on VMWare and Mellanox NICs is not supported.

[See [Configuring SR-IOV 10-Gigabit High Availability on vSRX 3.0](#).]

Juniper Extension Toolkit (JET)

- **Juniper Extension Toolkit (JET) support for 64-bit applications (MX5, MX10, MX40, MX80, MX104, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX ELM, JunosV Firefly, cSRX, SRX100, SRX110, SRX210, SRX220, SRX240, SRX300, SRX320, SRX340, SRX345, SRX550, SRX550HM, SRX650, SRX720E, SRX750E, SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SRX4400, SRX4600, SRX4800, SRX5400, SRX5600, SRX5800, SRX7X0E, SRX-ES7, SRX-ES8, VMX, and vSRX)**—Starting in Junos OS Release 20.4R1, JET supports 64-bit applications. Use the following commands to compile 64-bit applications for use with the AMD64 or ARM64 64-bit processor architecture.
 - **mk-amd64**: Compiles the application for use with AMD64 and Junos OS with FreeBSD.
 - **mk-amd64,bsd**: Compiles the application for use with AMD64 and Junos OS with upgraded FreeBSD.

- `mk-arm64,bsd`: Compiles the application for use with ARM64 and Junos OS with upgraded FreeBSD.

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

Junos OS XML ,API, and Scripting

- **Support for Certificate Authority Chain Profile (EX2300, EX3400, EX4300, MX240, MX480, MX960, PTX-5000, VMX, vSRX and QFX5200)**—Starting in Junos OS Release 20.4R1, you can configure intermediate Certificate Authority (CA) chain profile certificate and perform https REST API request using mutual and server authentications.

To configure intermediate ca-chain certificate, configure **ca-chain ca-chain** statement at the `[edit system services rest https]` hierarchy level.

Network Management and Monitoring

- **Configuration retrieval using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you can use the configuration revision identifier feature to view the configuration for a specific revision. This configuration database revision can be viewed with the CLI command **show system configuration revision**.

[See [show system configuration revision](#).]

- **Junos XML protocol operations support loading and comparing configurations using the configuration revision identifier (EX3400, EX4300, MX204, MX240, MX480, MX960, MX2020, PTX3000, PTX10008, QFX5100, QFX10002-60C, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, the Junos XML management protocol operations support loading and comparing configurations by referencing the configuration revision identifier of a committed configuration. You can execute the **<load-configuration>** operation with the **configuration-revision** attribute to load the configuration with the given revision identifier into the candidate configuration. Additionally, you can compare the candidate or active configuration to a previously committed configuration by referencing the configuration revision identifier for the comparison configuration. The **<get-configuration>** operation supports the **compare="configuration-revision"** and **configuration-revision** attributes to perform the comparison.

[See [<get-configuration>](#) and [<load-configuration>](#).]

Platform and Infrastructure

- **LiquidIO DPDK driver support (vSRX3.0)**—Starting in Junos OS Release 20.4R1, vSRX3.0 supports LiquidIO DPDK driver with KVM hypervisor. If you use the LiquidIO II smart NICs, then you can use vSRX3.0 by the virtual function of SR-IOV.

[See [Requirements for vSRX on KVM.](#)]

Routing Protocols

- **Support for multiple single-hop EBGp sessions on different links using the same IPv6 link-local address (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 20.4R1, you are no longer required to have unique peer addresses for Juniper devices for every EBGp session. You can now enable single-hop EBGp sessions on different links over multiple directly connected peers that use the same IPv6 link-local address.

In earlier Junos OS Releases, BGP peers could be configured with link-local addresses, but multiple BGP peers could not be configured to use the same link-local address on different interfaces.

[See [Configure Multiple Single-Hop EBGp Sessions on Different Links Using the Same Link-Local Address \(IPv6\).](#)]

VPNs

- **AWS Key Management Service (KMS) Integration support (vSRX 3.0)**—Starting in Junos OS Release 20.4R1, you can safeguard the private keys used by the PKI daemon and IKED using AWS Key Management Service (KMS). You can establish a PKI daemon-based VPN tunnel using the keypairs generated at the KMS. The KMS server creates, stores, and performs the needed keypair operations. After you enable KMS, all the PKI daemons keypairs previously created are deleted.

[See *Deploying vSRX 3.0 for Securing Data using AWS KMS.*]

What's Changed

IN THIS SECTION

- [Platform and Infrastructure](#) | 279

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

Platform and Infrastructure

- 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).
- **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](#).

Known Limitations

There are no known behaviors for vSRX in Junos OS Release 20.4R1.

Open Issues

IN THIS SECTION

- [J-Web](#) | 280
- [Platform and Infrastructure](#) | 280

Learn about open issues in this release for vSRX.

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

J-Web

- J-Web GUI does not allow you to save a rule if the cumulative shared objects are more than 2500 before the policy grid is saved. When there are several shared objects, there will be a noticeable delay in opening sources and destinations of a rule, and performing rule action. [PR1540047](#)
- When the commit pending changes message is shown on the J-Web UI, the contents of other messages, landing page, or pop-ups will not be clearly visible. [PR1554024](#)
- When upgrading to Junos OS Release 20.4R1 or later, any existing entries within the on-box logging database (security logs) are cleared. This is due to the high performance database design that is not forward-compatible from Junos OS Release 20.3 or earlier versions. These are the logs normally visible within J-Web under the Monitoring>Logs page. [PR1541674](#)

Platform and Infrastructure

- The IPv6 traffic redirection by NSX-T edge infrastructure is not supported. [PR1527130](#)
- Ensure the MTU on the host is large enough before setting the MTU in vSRX. [PR1537984](#)

Resolved Issues

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

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

Application Security

- The flowd or srxpfe process might crash when SSL proxy and AppSecure process traffic simultaneously. [PR1516969](#)
- During rare circumstances, if the AppID unknown packet capture functionality is enabled, the srxpfe process might crash and generate a core file. [PR1538991](#)

Chassis Clustering

- The control link might be broken when there is excessive traffic load on the control link in a vSRX cluster deployment. [PR1524243](#)

CLI

- On Microsoft Azure deployments, SSH public key authentication is not supported for vSRX 3.0 CLI and portal deployment. [PR1402028](#)
- Commit is not successful when configuration committed without active probe settings options (all options under active probe settings are optional). [PR1533420](#)
- The master-password configuration is rejected if master-encryption-password (MEK) is not set. [PR1537251](#)

Flow-Based and Packet-Based Processing

- A chassis cluster node might stop passing traffic. [PR1528898](#)

Install and Upgrade

- Upgrading to Junos OS Release 20.4R1 or later releases with a large, pre-existing security-log database might result in LLMD consuming large amounts of CPU. [PR1548423](#)

Interfaces and Chassis

- LiquidIO SR-IOV configuring ge interface as DHCP client does not work; no IP address obtained. [PR1529228](#)

Intrusion Detection and Prevention (IDP)

- The flowd or srpxfe process might generate core files during the idpd process commit on SRX Series devices. [PR1521682](#)
- When adaptive threat profiling is configured within an IDP rule base and logging is enabled, on the vSRX instances the Packet Forwarding Engine process might stop and generate a core file. [PR1532737](#)

Platform and Infrastructure

- The vSRX may restart unexpectedly. [PR1479156](#)
- In vSRX3.0 on Azure with key-vault enabled, change in MEK results in deletion of certificates. [PR1513456](#)
- With CSO SD-WAN configuration loaded, the flowd process generates core files while deleting the GRE IPsec configuration. [PR1513461](#)
- Configuration integrity mismatch error in vSRX3.0 running on Azure with key-vault integrated. [PR1551419](#)

Routing Policy and Firewall Filters

- Junos OS upgrade may encounter failure in certain conditions when enabling ATP. [PR1519222](#)

User Access and Authentication

- On vSRX 3.0 on Azure, with Microsoft Azure Hardware Security Module (HSM) enabled, keypair generation fails if you reuse the certificate ID for creating a new keypair—even if the previous keypair was deleted. [PR1490558](#)

VPNs

- The Ping-icmp test fails after configuring ECMP routes over multipoint tunnel interface VPNs. [PR1438311](#)
- The flowd process might stop in an IPsec VPN scenario. [PR1517262](#)

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- [Upgrading Software Packages | 284](#)
- [Validating the OVA Image | 289](#)

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 20.4R1 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 20.4R1 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 20.4R1 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-x86-64-20.4-2020-10-12.0_RELEASE_20.4_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-x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz

```

no-copy no-validate reboot

```

Verified junos-vsrx-x86-64-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING:      This package will load JUNOS 20.4 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-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31
junos-srx-mr-vsrx-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31
junos-srx-mr-vsrx-20.4-2020-10-12.0_RELEASE_20.4_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-20.4-2020-10-12.0_RELEASE_20.4_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-20.4-2020-10-12.0_RELEASE_20.4_THROTTLE-linux.tgz ...
upgrade_platform: Input package
/var/tmp/junos-srx-mr-vsrx-20.4-2020-10-12.0_RELEASE_20.4_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-20.4-2020-10-12.0_RELEASE_20.4_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-20.4-2020-10-12.0_RELEASE_20.4_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 20.4R1 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 20.4-2020-10-12.0_RELEASE_20.4_THROTTLE Kernel 64-bit
JNPR-11.0-20171012.170745_fbsd-
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@:~ # cli
root> show version
Model: vsrx
Junos: 20.4-2020-10-12.0_RELEASE_20.4_THROTTLE
JUNOS OS Kernel 64-bit [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs [20171012.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20171012.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20171012.170745_fbsd-builder_stable_11]
JUNOS py extensions [20171017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20171017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20171012.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20171012.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities
[20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package
[20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20171017.110007_ssd-builder_release_174_throttle]
```

```
JUNOS srx Data Plane Crypto Support
[20171017.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20171017.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20171017.110007_ssd-builder_release_174_throttle]
```

Validating the OVA Image

If you have downloaded a vSRX .ova image and need to validate it, see [Validating the vSRX .ova File for VMware](#).

Note that only .ova (VMware platform) vSRX images can be validated. The .qcow2 vSRX images for use with KVM cannot be validated the same way. File checksums for all software images are, however, available on the download page.

Upgrading Using ISSU

In-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

For additional information about using ISSU on routing and switching devices, see the [High Availability User Guide](#).

For additional information about using ISSU on security devices, see the [Chassis Cluster User Guide for SRX Series Devices](#).

For information about ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\) Web application](#).

Licensing

Starting in 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that have been developed at Juniper Networks over the past several years.

The major components of the framework are:

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

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

Compliance Advisor

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

Finding More Information

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


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

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

Documentation Feedback

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

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



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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

Revision History

31 March 2022—Revision 11, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

17 March 2022—Revision 10, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

28 January 2022—Revision 9, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

7 October 2021—Revision 8, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

23 September 2021—Revision 7, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

4 March 2021—Revision 6, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

28 January 2021—Revision 5, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

13 January 2021—Revision 4, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

7 January 2021—Revision 3, Junos OS Release 20.4R1— ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

29 December 2020—Revision 2, Junos OS Release 20.4R1—ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

28 December 2020—Revision 1, Junos OS Release 20.4R1—ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.