

Junos[®] OS 12.1 Release Notes

Release 12.1R7
06 August 2013
Revision 4

These release notes accompany Release 12.1R7 of the Junos operating system (Junos OS). They describe device documentation and known problems with the software. Junos OS runs on all Juniper Networks M Series, MX Series, and T Series routing platforms, SRX Series Services Gateways, J Series Services Routers, and the EX Series Ethernet Switches..

For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://prsearch.juniper.net>.

You can also find these release notes on the Juniper Networks Junos OS Documentation Web page, which is located at <https://www.juniper.net/techpubs/software/junos/>.

Contents

Junos OS Release Notes for EX Series Switches	8
New Features in Junos OS Release 12.1 for EX Series Switches	8
Hardware	9
Access Control and Port Security	11
Class of Service (CoS)	12
Converged Networks (LAN and SAN)	12
Ethernet Switching and Spanning Trees	12
FIPS	12
Firewall Filters	13
High Availability	13
Infrastructure	14
Interfaces	16
J-Web Interface	16
Layer 2 and Layer 3 Protocols	16
Management and Remote Monitoring	16
MPLS	16
Multicast Protocols	18
Power over Ethernet (PoE)	18
Software Installation and Upgrade	18

Virtual Chassis	18
Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches	18
Infrastructure	19
Power over Ethernet (PoE)	19
Limitations in Junos OS Release 12.1 for EX Series Switches	20
Access Control and Port Security	20
Ethernet Switching and Spanning Trees	20
Firewall Filters	20
Hardware	20
High Availability	21
Infrastructure	21
Interfaces	22
J-Web Interface	23
Layer 2 and Layer 3 Protocols	24
Management and RMON	24
Virtual Chassis	24
Outstanding Issues in Junos OS Release 12.1 for EX Series Switches	26
Access Control and Port Security	26
Converged Networks (LAN and SAN)	26
Ethernet Switching and Spanning Trees	26
Hardware	26
High Availability (HA) and Resiliency	27
Infrastructure	27
Interfaces	28
J-Web Interface	28
Layer 2 and Layer 3 Protocols	31
Management and RMON	31
Software Upgrade and Installation	31
Virtual Chassis	32
Resolved Issues in Junos OS Release 12.1 for EX Series Switches	32
Issues Resolved in Release 12.1R1	33
Issues Resolved in Release 12.1R2	42
Issues Resolved in Release 12.1R3	45
Issues Resolved in Release 12.1R4	47
Issues Resolved in Release 12.1R5	50
Issues Resolved in Release 12.1R6	51
Issues Resolved in Release 12.1R7	53
Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches	55
Changes to Junos OS for EX Series Switches Documentation	56
Errata	56
Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches	57
Upgrade and Downgrade Support Policy for Junos OS Releases	57
Upgrading to Release 12.1R2 or Later Releases, with Existing VSTP Configurations	58
Upgrading from Junos OS Release 10.4R3 or Later	58
Upgrading from Junos OS Release 10.4R2 or Earlier	60

Upgrading EX Series Switches Using NSSU	60
Junos OS Release Notes for Branch SRX Series Services Gateways and J Series Services Routers	62
Release Notes Updates	62
New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	63
Release 12.1R4 New Features	63
Release 12.1R3 New Features	65
Release 12.1R2 Software Features	67
Release 12.1 R2 Hardware Features	72
Release 12.1 R1 Software Features	74
Hardware Features—SRX550 Services Gateways	81
Software Features—SRX550 Services Gateways	83
Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	84
AppSecure	84
Chassis Cluster	84
Command-Line Interface (CLI)	84
Deprecated Items for Security Hierarchy	85
Flow and Processing	86
Hardware	87
Interfaces and Routing	87
Internet Protocol Security (IPsec)	87
Intrusion Detection and Prevention	88
J-Web	88
Remote ID Check for Site-to-Site VPNs	88
Security Policy Rules	88
System Logging	88
Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	89
ADSL Mini-PIM	89
AppSecure	89
AX411 Access Points	90
Chassis Cluster	90
Command-Line Interface (CLI)	91
DOCSIS Mini-PIM	91
Dynamic Host Configuration Protocol (DHCP)	91
Dynamic VPN	91
Firewall Filters	92
Flow and Processing	92
Group VPN Interoperability with Cisco's GET VPN for Juniper Networks Security Devices that Support Group VPN	93
Interfaces and Routing	94
Internet Key Exchange Version 2 (IKEv2)	97
Internet Protocol Security (IPsec)	97
Intrusion Detection and Prevention (IDP)	97
IPv6 IPsec	99
Layer 2 Transparent Mode	100
IPv6 Support	100

J-Web	100
Network Address Translation (NAT)	101
Power over Ethernet (PoE)	102
Security	102
Simple Network Management Protocol (SNMP)	103
Switching	103
Unified Threat Management (UTM)	104
USB	104
Upgrade and Downgrade	104
Virtual Private Networks (VPNs)	104
Unsupported CLI for Branch SRX Series Services Gateways and J Series Services Routers	104
Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	110
Outstanding Issues in Junos OS Release 12.1R7 for Branch SRX Series Services Gateways	110
Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	111
Resolved Issues in Junos OS Release 12.1R7 for Branch SRX Series Services Gateways	112
Resolved Issues in Junos OS Release 12.1R6 for Branch SRX Series Services Gateways	113
Resolved Issues in Junos OS Release 12.1R5 for Branch SRX Series Services Gateways	115
Resolved Issues in Junos OS Release 12.1R4 for Branch SRX Series Services Gateways	116
Resolved Issues in Junos OS Release 12.1 R3 for Branch SRX Series Services Gateways	119
Resolved Issues in Junos OS Release 12.1 R2 for Branch SRX Series Services Gateways	123
Resolved Issues in Junos OS Release 12.1 R1 for Branch SRX Series Services Gateways	126
Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	135
Errata for the Junos OS Software Documentation	135
Errata for the Junos OS Hardware Documentation	143
Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers	146
Upgrading and Downgrading among Junos OS Releases	146
Upgrading an AppSecure Device	148
Upgrade and Downgrade Scripts for Address Book Configuration	148
Hardware Requirements for Junos OS Release 12.1 for SRX Series Services Gateways and J Series Services Routers	151
Junos OS Release Notes for High-End SRX Series Services Gateways	153
Release Notes Updates	154
New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways	154
Release 12.1R6 New Features	154
Release 12.1R3 New Features	158

Release 12.1R2 New Features	159
Software Features	160
Changes in Default Behavior and Syntax in Junos OS Release 12.1 for	
High-End SRX Series Services Gateways	168
AppSecure Application Package Upgrade Changes	168
Chassis Cluster	169
Command-Line Interface (CLI)	169
Deprecated Items for High-End SRX Series Services Gateways	169
Flow and Processing	170
In-Service Software Upgrade (ISSU)	171
J-Web	172
Logical Systems	172
Management Information Base (MIB)	172
Remote ID Check for Site-to-Site VPNs	172
Security	173
Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services	
Gateways	174
AppSecure	174
Chassis Cluster	174
Dynamic Host Configuration Protocol (DHCP)	175
Firewall Filters	175
Flow and Processing	175
Hardware	177
Interfaces and Routing	178
Internet Key Exchange Version 2 (IKEv2)	178
Intrusion Detection and Prevention (IDP)	178
IPv6 IPsec	180
IPv6 Support	181
J-Web	182
Logical Systems	182
Network Address Translation (NAT)	183
Security	185
Simple Network Management Protocol (SNMP)	186
Virtual Private Networks (VPNs)	186
Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services	
Gateways	187
Outstanding Issues in Junos OS Release 12.1R7 for High-End SRX Series	
Services Gateways	187
Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services	
Gateways	188
Resolved Issues in Junos OS Release 12.1R7 for High-End SRX Series	
Services Gateways	188
Resolved Issues in Junos OS Release 12.1R6 for High-End SRX Series	
Services Gateways	190
Resolved Issues in Junos OS Release 12.1R5 for High-End SRX Series	
Services Gateways	192
Resolved Issues in Junos OS Release 12.1R4 for High-End SRX Series	
Services Gateways	194

Resolved Issues in Junos OS Release 12.1R3 for High-End SRX Series Services Gateways	196
Resolved Issues in Junos OS Release 12.1R2 for High-End SRX Series Services Gateways	199
Resolved Issues in Junos OS Release 12.1R1 for High-End SRX Series Services Gateways	201
Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways	210
Errata for the Junos OS Software Documentation	210
Upgrade and Downgrade Instructions for Junos OS Release 12.1 for High-End SRX Series Services Gateways	217
Upgrading and Downgrading among Junos OS Releases	217
Upgrading an AppSecure Device	219
Upgrade and Downgrade Scripts for Address Book Configuration	219
Upgrade Policy for Junos OS Extended End-Of-Life Releases	222
Hardware Requirements for Junos OS Release 12.1 for High-End SRX Series Services Gateways	222
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers	223
New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers	223
Class of Service	223
High Availability	231
Interfaces and Chassis	231
Junos OS XML API and Scripting	248
Layer 2 Ethernet Services	249
MPLS Applications	250
Multicast	253
Network Management	254
Routing Protocols	256
Subscriber Access Management	262
System Logging	273
User Interface and Configuration	280
VPNs	282
Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers	284
Changes in Default Behavior and Syntax	284
Changes Planned for Future Releases	297
Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers	298
Current Software Release	298
Previous Releases	333
Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers	404
Errata	404
Changes to the Junos OS Documentation Set	431

Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers	432
Basic Procedure for Upgrading to Release 12.1	432
Upgrade and Downgrade Support Policy for Junos OS Releases	435
Upgrading a Router with Redundant Routing Engines	435
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1	436
Upgrading the Software for a Routing Matrix	437
Upgrading Using ISSU	438
Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR	439
Downgrading from Release 12.1	440
Junos OS Documentation and Release Notes	441
Documentation Feedback	441
Requesting Technical Support	441
Revision History	443

Junos OS Release Notes for EX Series Switches

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)
- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

New Features in Junos OS Release 12.1 for EX Series Switches

This section describes new features in Release 12.1 of the Junos operating system (Junos OS) for EX Series switches.

Not all EX Series software features are supported on all EX Series switches in the current release. For a list of all EX Series software features and their platform support, see [EX Series Switch Software Features Overview](#) and [EX Series Virtual Chassis Software Features Overview](#).

New features are described on the following pages:

- [Hardware on page 9](#)
- [Access Control and Port Security on page 11](#)
- [Class of Service \(CoS\) on page 12](#)
- [Converged Networks \(LAN and SAN\) on page 12](#)
- [Ethernet Switching and Spanning Trees on page 12](#)
- [FIPS on page 12](#)
- [Firewall Filters on page 13](#)
- [High Availability on page 13](#)
- [Infrastructure on page 14](#)
- [Interfaces on page 16](#)
- [J-Web Interface on page 16](#)
- [Layer 2 and Layer 3 Protocols on page 16](#)
- [Management and Remote Monitoring on page 16](#)
- [MPLS on page 16](#)
- [Multicast Protocols on page 18](#)
- [Power over Ethernet \(PoE\) on page 18](#)

- [Software Installation and Upgrade on page 18](#)
- [Virtual Chassis on page 18](#)

Hardware

- **Enhancements for EX6210 switch line cards and SRE modules**—The EX6210 switch has 10 horizontal slots on the front of the chassis. Slots 0 through 3 and 6 through 9 accept one line card each. You can now install either a line card or a Switch Fabric and Routing (SRE) module in slots 4 and 5. You can install a maximum of nine line cards in a switch in slots 0 through 9; however, at least one SRE module must be installed in the switch. [See [Slot Numbering for an EX6210 Switch](#).]
- **Extra-scale line cards on EX8200 switches**—The following extra-scale line cards, supported on standalone EX8200 switches and on EX8200 Virtual Chassis, provide larger route table sizes than the associated non-extra-scale models to store more IPv4 and IPv6 unicast routes:
 - EX8200-8XS-ES
 - EX8200-40XS-ES
 - EX8200-48F-ES
 - EX8200-48T-ES

[See [Line Card Model and Version Compatibility in an EX8200 Switch](#).]

- **LCD panel support for the XRE200 External Routing Engine**—You can now use the LCD panel on the XRE200 External Routing Engine to configure and monitor the external Routing Engine. You can now navigate to the Maintenance menu and the Status menu in the LCD panel. You can use the Maintenance menu to perform basic maintenance tasks, such as halting or rebooting the external Routing Engine or loading a rescue or factory-default configuration. You can use the Status menu to monitor external Routing Engine status, including monitoring of the Virtual Chassis ports (VCPs), power supplies, temperatures, and the installed Junos OS version. [See [LCD Panel in an XRE200 External Routing Engine](#).]
- **New AC power supply support on EX6200 switches**—EX6200 switches now support 5000-W AC power supplies. [See [AC Power Supplies in an EX6200 Switch](#).]
- **New optical transceiver support for EX2200 switches**—EX2200 switches now support the following optical transceivers:
 - EX-SFP-GE80KCW1470
 - EX-SFP-GE80KCW1490
 - EX-SFP-GE80KCW1510
 - EX-SFP-GE80KCW1530
 - EX-SFP-GE80KCW1550
 - EX-SFP-GE80KCW1570
 - EX-SFP-GE80KCW1590
 - EX-SFP-GE80KCW1610

[See [Pluggable Transceivers Supported on EX2200 Switches.](#)]

- **New optical transceiver support for EX3200 and EX4200 switches**—EX3200 and EX4200 switches now support the following optical transceivers:
 - EX-SFP-GE80KCW1470
 - EX-SFP-GE80KCW1490
 - EX-SFP-GE80KCW1510
 - EX-SFP-GE80KCW1530
 - EX-SFP-GE80KCW1550
 - EX-SFP-GE80KCW1570
 - EX-SFP-GE80KCW1590
 - EX-SFP-GE80KCW1610
 - EX-XFP-10GE80KDWDM (on EX4200 switches only)

[See [Pluggable Transceivers Supported on EX3200 Switches](#) and [Pluggable Transceivers Supported on EX4200 Switches.](#)]

- **New optical transceiver support for EX3300 switches**—EX3300 switches now support the following optical transceivers:
 - EX-SFP-1GE-LH
 - EX-SFP-1GE-LX40K
 - EX-SFP-1GE-T
 - EX-SFP-10GE-ER
 - EX-SFP-GE10KT13R14
 - EX-SFP-GE10KT13R15
 - EX-SFP-GE10KT14R13
 - EX-SFP-GE10KT15R13
 - EX-SFP-GE40KT13R15
 - EX-SFP-GE40KT15R13

[See [Pluggable Transceivers Supported on EX3300 Switches.](#)]

- **New optical transceiver support for EX6200 switches**—EX6200 switches now support the following optical transceivers:
 - EX-SFP-1GE-LH
 - EX-SFP-1GE-LX
 - EX-SFP-1GE-LX40K
 - EX-SFP-1GE-SX
 - EX-SFP-1GE-T

[See [Optical Interface Support in EX6200 Switches.](#)]

- **New optical transceiver support for EX8200 switches**—The 40-port SFP+ and 48-port SFP line cards in EX8200 switches now support the following optical transceivers:
 - EX-SFP-FE20KT13R15
 - EX-SFP-FE20KT15R13

[See [Pluggable Transceivers Supported on EX8200 Switches.](#)]

- **RPS support on EX2200 and EX3300 switches**—Unlike other EX Series switches, which support redundant power supplies, EX2200 switches and EX3300 switches have only one power supply. If you deploy one of these switches in a critical situation, we recommend that you connect an RPS to that switch to supply backup power in case a loss of power occurs. The RPS is not a primary power supply—it provides backup power to switches only when the single dedicated power supply fails. An RPS operates in parallel with the single dedicated power supplies of the switches connected to it and provides all connected switches with either PoE or non-PoE backup power.

The RPS is supported on EX2200 switches, EX3300 switches, and EX3300 Virtual Chassis. You can connect multiple RPSs to an EX3300 Virtual Chassis.

[See [EX Series Redundant Power System \(RPS\) Documentation.](#)]

Access Control and Port Security

- **Access control feature enhancements on EX4500 switches**—EX4500 switches now support 802.1X authentication (port-based, multiple supplicant) and 802.1X authentication with VLAN assignment and VoIP VLAN support. [See [Access Control on EX Series Switches.](#)]
- **Port security feature enhancements on EX4500 switches**—EX4500 switches now support DHCP snooping, persistent storage for DHCP snooping, and IP source guard. [See [Port Security on EX Series Switches.](#)]

Class of Service (CoS)

- **Interface-specific IPv6 classifiers and rewrite rules**—On EX4500 switches and EX4500 Virtual Chassis you can now configure and apply IPv6 classifiers and rewrite rules for each interface. [See [Understanding CoS Classifiers](#) and [Understanding CoS Rewrite Rules](#).]

Converged Networks (LAN and SAN)

- **DCBX support for the application protocol TLV on EX4500 switches**—Support for DCBX on EX4500 switches has been expanded to include support for the application protocol TLV. This feature enables you to implement DCBX for other Layer 2 and Layer 4 applications in addition to implementing it for Fibre Channel over Ethernet (FCoE) applications. DCBX is required for FCoE applications. Although it is not required for other applications, it adds reliability for enterprise data storage. By default, the FCoE application is enabled on DCBX interfaces. To use this feature for other Layer 2 and Layer 4 applications, you must configure an application map and then associate it with the DCBX interface that is carrying the application's traffic. [See [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#).]

Ethernet Switching and Spanning Trees

- **Diagnostics and debugging enhancement**—A new command, **show pfe statistics bridge**, displays the number of packets received, the number of ingress packets discarded and the reasons for the discard, and the number of packets transmitted through the egress pipeline of the Packet Forwarding Engine. You can use this information for troubleshooting investigations. [See [show pfe statistics bridge](#).]
- **Edge virtual bridging**—Edge virtual bridging (EVB) enables multiple virtual machines to communicate with one another and with external hosts in an Ethernet network environment. Servers using a virtual Ethernet packet aggregator (VEPA) to provide bridging support between multiple virtual machines, and external networks do not send packets directly from one virtual machine (VM) to another. Instead, the packets are sent to virtual bridges on an adjacent switch for processing. EX Series switches use EVB as a virtual bridge to return the packets on the same interface that delivered the packets. [See [Example: Configuring Edge Virtual Bridging for Use with VEPA Technology](#).]
- **Ethernet ring protection switching for EX Series switches**—Ethernet ring protection switching (ERPS), defined by ITU-T G8032, is a mechanism for preventing unwanted loops in Ethernet networks. It is supported on EX2200, EX3200, and EX4200 switches. [See [Example: Configuring Ethernet Ring Protection Switching on EX Series Switches](#).]

FIPS

- **FIPS mode certified in Junos OS Release 12.1R6 on EX3300, EX4200, EX4500, EX6200, and EX8200 switches**—FIPS 140-2 defines security levels for hardware and software that perform cryptographic functions. By meeting the applicable overall requirements within the FIPS standard, Juniper Networks EX Series Ethernet Switches running the Juniper Networks Junos OS in FIPS mode comply with the FIPS 140-2 Level 1 standard. Operating EX Series switches in a FIPS 140-2 Level 1 environment requires

enabling and configuring FIPS mode on the switches from the Junos OS command-line interface (CLI).

FIPS mode is certified in Junos OS Release 12.1R6 and is supported on the following EX Series switches:

- EX3300 switch
- EX4200 switch
- EX4500 switch
- EX6210 switch
- EX8208 switch
- EX8216 switch

See [Junos OS for EX Series Ethernet Switches, Release 12.1R6: FIPS](#).

Firewall Filters

- **Support for IPv6 firewall filters on EX4500 switches**—On EX4500 Virtual Chassis and EX4500 standalone switches, you can apply match conditions to IPv6 traffic on Layer 3 interfaces and aggregated Ethernet interfaces. The following match conditions are now applicable to IPv6 traffic: destination-address, destination-port, icmp-code, icmp-type, next-header, source-address, source-port, tcp-established, tcp-flags, tcp-initial, and traffic-class. The following actions and action modifiers are applicable to IPv6 traffic: accept, analyzer, count, discard, forwarding-class, loss-priority, and policer. [See [Descriptions of Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches](#).]
- **Support for the vlan action on EX8200 switches and EX8200 Virtual Chassis**—In firewall filter configurations for EX8200 standalone switches, you can now apply the vlan action on ports and VLANs for IPv4 and IPv6 ingress traffic. However, the vlan action works properly only when the interface action modifier is also configured along with the vlan action. For EX8200 Virtual Chassis, you can apply the vlan action (provided that the interface action modifier is also configured) only on VLANs for IPv4 and IPv6 ingress traffic. You can specify the interface action modifier to forward matched packets to a specific interface, bypassing the switching lookup. You can specify the vlan action to forward matched packets to a specific VLAN in the Virtual Chassis. [See [Descriptions of Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches](#).]

High Availability

- **GRES for IGMP snooping on EX3300 Virtual Chassis, EX4500 Virtual Chassis, and EX6200 switches**—GRES is now supported for IGMP snooping on these indicated platforms. [See [High Availability Features for EX Series Switches Overview](#).]
- **Nonstop active routing for BGP, IGMP, IS-IS, OSPF, and RIP with BFD on EX3300 Virtual Chassis**—NSR for OSPF with BFD, RIP with BFD, IS-IS with BFD, BGP with BFD, and IGMP with BFD is now supported on EX3300 Virtual Chassis. You can now configure NSR to enable transparent switchover between the master and backup Routing Engines

without having to restart any of these protocols. [See [Understanding Nonstop Active Routing on EX Series Switches](#).]

- **Nonstop active routing for PIM on EX8200 switches and Virtual Chassis**—NSR for Protocol Independent Multicast (PIM) is now supported on EX8200 switches and Virtual Chassis. [See [Understanding Nonstop Active Routing on EX Series Switches](#).]
- **Nonstop bridging for spanning-tree protocols on EX4500 Virtual Chassis and EX8200 Virtual Chassis**—NSB for spanning-tree protocols is now supported on EX4500 Virtual Chassis and EX8200 Virtual Chassis. You can now configure NSB to enable transparent switchover between the master and backup Routing Engines without having to restart any spanning-tree protocol. [See [Understanding Nonstop Bridging on EX Series Switches](#).]
- **Nonstop bridging for spanning-tree protocols, LACP, LLDP, and LLDP-MED on EX6200 switches**—NSB for spanning-tree protocols, LACP, LLDP, and LLDP-MED is now supported on EX6200 switches. You can now configure NSB to enable transparent switchover between the master and backup Routing Engines without having to restart any of these protocols. [See [Understanding Nonstop Bridging on EX Series Switches](#).]
- **Nonstop software upgrade on EX4200 and EX4500 Virtual Chassis**—NSSU is now supported on EX4200 and EX4500 Virtual Chassis. [See [Understanding Nonstop Software Upgrade on EX Series Switches](#).]
- **Virtual Chassis fast failover for EX4500 Virtual Chassis and mixed EX4200 and EX4500 Virtual Chassis**—Virtual Chassis fast failover is now supported on Virtual Chassis ports (VCPs) in an EX4500 Virtual Chassis or in a mixed EX4200 and EX4500 Virtual Chassis. The Virtual Chassis fast failover feature is a hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link or switch failure. [See [Understanding Fast Failover in an EX3300, EX4200, or EX4500 Virtual Chassis](#).]

Infrastructure

- **Extended DHCP server and extended DHCP relay**—EX Series switches now support both extended DHCP server and extended DHCP relay, and the legacy version of DHCP. [See [Understanding DHCP Services for EX Series Switches](#).]
- **New software features for EX6200 switches**—The following software features are now supported for EX6200 switches:
 - BFD protocol for BGP, IS-IS, OSPF, PIM, and RIP
 - BGP for IPv6
 - Captive portal authentication for Layer 3 interfaces
 - CoS features for IPv6
 - CoS features, including DSCP, IEEE 802.1p, and IP precedence packet rewrites on ingress RVIs
 - Distributed BFD
 - Filter-based S-VLAN tagging
 - Firewall filters on management Ethernet interfaces

- IPv6 firewall filters
- IPv6 ping operation
- IPv6 static routing
- IPv6 traceroute
- IS-IS for IPv6
- Junos OS image rollback
- L2PT
- MVRP (IEEE 802.1ak)
- MBGP
- NDP
- OSPFv3
- Path MTU discovery
- PIM for IPv6 multicast
- Q-in-Q tunneling
- RPM enables hardware timestamps on RVIs
- RIPng
- RPM client and server on the same interface
- Self-signed digital certificates for enabling SSL services
- sFlow monitoring technology
- VRRP for IPv6

[See [Class of Service for EX Series Switches](#), [Ethernet Switching on EX Series Switches](#), [Layer 3 Protocols Supported on EX Series Switches](#), [Routing Policy and Packet Filtering for EX Series Switches](#), [Understanding Authentication on EX Series Switches](#), and [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch](#).]

- **wildcard range configuration mode command**—EX Series switches now support the **wildcard range** configuration mode command. The **wildcard range** command enables you to specify ranges in the **activate**, **deactivate**, **delete**, **protect**, **set**, **show**, and **unprotect** commands. You can use ranges to specify a range of interfaces, logical units, VLANs, and other numbered elements. The **wildcard range** command expands the command you entered into multiple commands, each of which corresponds to one item in the range. For example, the command **wildcard range interfaces deactivate ge-0/0/[1-3]** expands to the commands **deactivate interfaces ge-0/0/1**, **deactivate interfaces ge-0/0/2**, and **deactivate interfaces ge-0/0/3**. [See [Example: Using the Wildcard Command with the Range Option](#).]

Interfaces

- **Generic routing encapsulation**—EX3200 and EX4200 switches now support GRE, a tunneling protocol to transport packets over a network. You can use GRE tunneling services to encapsulate any network layer protocol over any other network layer protocol. Acting as a tunnel source router, the switch encapsulates a payload packet that is to be transported through a tunnel to a destination network. The switch first encapsulates the payload packet in a GRE packet and then encapsulates the resulting GRE packet in a delivery protocol. A switch performing the role of a tunnel remote router extracts the tunneled packet and forwards the packet to the destination network. GRE tunnels can be used to connect noncontiguous networks and to provide options for networks that contain protocols with limited hop counts. [See [Understanding Generic Routing Encapsulation](#).]
- **Uplink failure detection on EX8200 switches and XRE200 External Routing Engines**—Uplink failure detection enables an EX Series switch to detect link failure on uplink interfaces and to propagate the failure to the downlink interfaces so that servers connected to those downlinks can switch over to secondary interfaces. Switches can have up to 48 groups, each with up to 48 uplinks and 48 downlinks for uplink failure detection. [See [Understanding Uplink Failure Detection](#).]

J-Web Interface

- **J-Web interface configuration for EX2200-C, EX3300, and EX6210 switches**—You can now configure the EX2200-C, EX3300, and EX6210 switches in the J-Web interface. [See [J-Web User Interface for EX Series Switches Overview](#).]

Layer 2 and Layer 3 Protocols

- **New Layer 3 protocols for EX3300 switches**—Several new Layer 3 protocols are now supported on EX3300 switches. [See [EX Series Switch Software Features Overview](#).]

Management and Remote Monitoring

- **Support for sFlow monitoring technology on EX8200 Virtual Chassis**—EX8200 Virtual Chassis now support sFlow monitoring technology. sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. You can configure sFlow technology to continuously monitor traffic at wire speed on all interfaces simultaneously. [See [Understanding How to Use sFlow Technology for Network Monitoring on an EX Series Switch](#) and [Junos OS for EX Series Switches, Release 12.1](#).]

MPLS

- **MPLS enhancements on EX8200 switches**—EX8200 Virtual Chassis now support all the MPLS features that are supported on EX8200 switches. In addition, EX8200 switches and EX8200 Virtual Chassis now support the following features:

- IPv6 tunneling and IPv6 Layer 3 VPNs—You can now configure EX8200 switches to tunnel IPv6 over an MPLS-based IPv4 network. This configuration enables you to interconnect a number of smaller IPv6 networks over an IPv4-based network core, enabling you to provide IPv6 service without having to upgrade the switches in your core network.
- MPLS over RVIs or Layer 3 subinterfaces—You can now use an RVI or a Layer 3 subinterface as the MPLS core-facing interface. The RVI functions as a logical router, eliminating the need for having both a switch and a router. Layer 3 subinterfaces allow you to route traffic among multiple VLANs along a single trunk line that connects an EX Series switch to a Layer 2 switch.
- Routed VLAN interfaces—On EX8200 Virtual Chassis, you can now use a RVI or a Layer 3 subinterface as the MPLS core-facing interface. The RVI functions as a logical router, eliminating the need for having both a switch and a router. Layer 3 subinterfaces allow you to route traffic among multiple VLANs along a single trunk line that connects an EX Series switch to a Layer 2 switch.
- Static LSPs—For static LSPs, you must manually assign labels (ingress, transit, and egress) on all the switches that are part of the LSP. No signaling protocol is needed. Configuring static LSPs is similar to configuring static routes on individual switches. As with static routes, there is no error reporting, liveness detection, or statistics reporting.
- Ultimate-hop popping using explicit NULL labels—EX8200 switches now support ultimate-hop popping. With ultimate-hop popping enabled, EXP bits are carried through to the egress PE switch. The egress PE switch makes use of EXP bits to classify the packets and send them out from the MPLS network. By default, ultimate-hop popping is disabled.

[See [Day One: Exploring IPv6 and MPLS for EX Series Switches](#).]

- **MPLS CoS enhancements on EX8200 switches**—EX8200 switches, both standalone and Virtual Chassis, support MPLS enhancements that enable you to prioritize certain types of traffic during periods of congestion. The enhancements are provided through the following CoS configurations:
 - EXP classification—EX8200 switches now support EXP classification and rewriting. If you enable the MPLS protocol family on a logical interface, the default MPLS EXP classifier is automatically applied to that logical interface. The default MPLS classifier maps EXP bits to forwarding classes and loss priorities.
 - EXP rewriting—You can now configure rewrite rules on the egress PE switch to alter the CoS settings of the packets. Rewrite rules set the value of the CoS bits within the packet's header. Each rewrite rule reads the current forwarding class and loss priority information associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header.
 - LSP CoS for both Layer 3 VPNs and Layer 2 VPNs—You can now configure a fixed CoS value for each LSP or for all LSPs on the switch. A fixed CoS value ensures that all packets entering the LSP are assigned the same class of service.

[See [MPLS for EX Series Switches](#).]

Multicast Protocols

- **MLD snooping on EX Series switches**—MLD snooping enables the switch to monitor MLD messages between IPv6 multicast routers and hosts. MLD version 1 (MLDv1) and MLDv2 are supported. When MLD snooping is enabled, the switch can determine which interfaces in a VLAN have interested listeners and forward multicast traffic only to those interfaces instead of flooding all interfaces in the VLAN. [See [Understanding MLD Snooping on EX Series Switches](#).]

Power over Ethernet (PoE)

- **PoE firmware upgrade**—You can now upgrade the PoE controller firmware from the CLI using the new command **request system firmware upgrade poe**. [See [request system firmware upgrade poe](#).]

Software Installation and Upgrade

- **Advanced feature licenses on EX3300 switches**—EX3300 switches now require an advanced feature license (AFL) to run all the advanced software features on the switch. [See [Understanding Software Licenses for EX Series Switches](#).]

Virtual Chassis

- **Member switch support enhancement on EX8200 Virtual Chassis**—You can now configure up to eight EX8200 member switches in an EX8200 Virtual Chassis. [See [Understanding EX8200 Virtual Chassis Components](#).]
- **Ingress counters on RVIs for EX8200 Virtual Chassis**—EX8200 Virtual Chassis can now maintain an ingress counter on RVIs. [See [Understanding Routed VLAN Interfaces on EX Series Switches](#).]

Related Documentation

- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)
- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches

This section lists the changes in default behavior and syntax in Junos OS Release 12.1 for EX Series switches.

Infrastructure

- The following changes have been made to the system snapshot functionality, which you configure using the **request system snapshot** command:
 - By default, a snapshot backs up both partitions (**/** and **/altroot**) to the media (internal or external) that the device did not boot from.
 - The following partitions are backed up by a snapshot: **/**, **/altroot**, **/config**, **/var**, and **/var/tmp**.
 - You do not need to specify a media slice number for the location of a snapshot.
 - You can specify the alternate slice on the media the device booted from as the snapshot location by using the **slice alternate** option.
 - When you create a snapshot on the media that the switch did not boot from, you must use the **partition** option to partition the destination media.
 - The **show system snapshot** command displays information for the **/** and **/altroot** partitions.

[This issue was being tracked by PR/599708.]

- The following changes have been made to the **show system snapshot** command:
 - You do not need to specify a media slice number for the location of a snapshot because all information about both slices is displayed by default.
 - The command displays information about the backup of the root file system (**/**) and directories **/altroot**, **/var**, and **/var/tmp**.

[This issue was being tracked by PR/785373.]

- The switch monitors available disk space in the **/var** partition every 10 minutes. If disk space in the **/var** partition is more than 75 percent of the partition space, the switch displays a yellow alarm. If disk space in the **/var** partition is more than 90 percent of the partition space, the switch displays both a yellow alarm and a red alarm. To avoid getting these warnings, use the **request system storage cleanup** command to clear up the disk space.

Power over Ethernet (PoE)

- The **show poe telemetries interface** command now supports using the keyword **all** in place of an interface name. If you specify **all**, records are displayed for all interfaces on which telemetries are enabled. In addition, a new command, **clear poe telemetries interface**, enables you to clear telemetry records from all interfaces or the interface you specify. When you use this command, telemetry collection stops on the specified interfaces. To restart telemetry collection, you must reconfigure telemetries on those interfaces.

Related Documentation

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)

- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

Limitations in Junos OS Release 12.1 for EX Series Switches

This section lists the limitations in Junos OS Release 12.1 for EX Series switches. If the limitation is associated with an item in our bug database, the description is followed by the bug tracking number.

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

Access Control and Port Security

- On EX Series switches, you cannot configure 802.1X authentication on RTGs. [This is a known software limitation.]

Ethernet Switching and Spanning Trees

- On EX Series switches, only dynamically learned routes can be imported from one routing table group to another. [This is a known software limitation.]

Firewall Filters

- On EX3200 and EX4200 switches, when a very large number of firewall filters are included in the configuration, it might take a long time, possibly a few minutes, for the egress filter rules to be installed. [PR/468806: This is a known software limitation.]
- On EX3300 switches, if you add and delete filters with a large number of terms (on the order of 1000 or more) in the same commit operation, not all the filters are installed. As a workaround, add filters in one commit operation, and delete filters in a separate commit operation. [PR/581982: This is a known software limitation.]
- On EX8200 switches, if you configure an implicit or explicit discard action as the last term in an IPv6 firewall filter on a loopback (lo0) interface, all the control traffic from the loopback interface is dropped. To prevent this, you must configure an explicit accept action. [This is a known software limitation.]

Hardware

- On 40-port SFP+ line cards for EX8200 switches, the LEDs on the left of the network ports do not blink to indicate that there is link activity if you set the speed of the network ports to 10/100/1000 Mbps. However, if you set the speed to 10 Gbps, the LEDs blink. [PR/502178: This is a known limitation.]

- You cannot connect EX2200-C-12P-2G switches to other vendors' prestandard IP phones. As a workaround, use a crossover cable. [PR/726929: This is a known limitation.]

High Availability

- You cannot verify that NSB is synchronizing Layer 2 protocol information to the backup Routing Engine even when NSB is properly configured. [PR/701495: This is a known software limitation.]
- When an EX8200 Virtual Chassis is using NSSU to upgrade from Junos OS Release 12.1R1 to a later Release 12.1 release, all network traffic on the Virtual Chassis is dropped, and all Layer 3 protocol states go down. After the NSSU finishes, normal operations resume. [PR/753548: This is a known software limitation.]
- On EX Series Virtual Chassis using NSSU to upgrade from Junos OS Release 11.2 or earlier to Junos OS Release 11.3 or later, after the NSSU operation finishes, the same MAC address might be assigned to multiple Layer 2 or aggregated Ethernet interfaces on different member switches within the Virtual Chassis. To set all Layer 2 and aggregated Ethernet ports to have unique MAC addresses, reboot the Virtual Chassis after the upgrade operation. To avoid these MAC address assignment issues, upgrade to Junos OS Release 11.3 or later without performing an NSSU operation.

Unique MAC address assignment for Layer 2 and aggregated Ethernet interfaces in a Virtual Chassis was introduced in Junos OS Release 11.3. If you are upgrading to Junos OS Release 11.2 or earlier, you should expect to see the same MAC address assigned to multiple ports on different member switches within the Virtual Chassis. [PR/775203: This is a known software limitation.]

Infrastructure

- Using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 10.4 to Release 11.1 or later is not recommended if you have configured the IGMP, MLD, or PIM protocols on the switch. If you have configured these multicast protocols, use the **request system software add** command to upgrade the software on an EX8200 switch from Junos OS Release 10.4 to Release 11.1 or later, as described in [Installing Software on an EX8200 Switch with Redundant Routing Engines \(CLI Procedure\)](#). This issue does not apply to upgrades from Junos OS Release 11.1 or later. [This is a known software limitation.]
- On EX Series switches, the **show snmp mib walk etherMIB** command does not display any output, even though the **etherMIB** is supported. This occurs because the values are not populated at the module level—they are populated at the table level only. You can issue the **show snmp mib walk dot3StatsTable**, **show snmp mib walk dot3PauseTable**, and **show snmp mib walk dot3ControlTable** commands to display the output at the table level. [This is a known software limitation.]
- Momentary loss of an inter-Routing Engine IPC message might trigger an alarm that displays the message **Loss of communication with Backup RE**. However, no functionality is affected. [PR/477943: This is a known software limitation.]

- Routing between virtual routing instances for local direct routes is not supported. [PR/490932: This is a known software limitation.]
- On EX4500 switches, the maintenance menu is not disabled even if you include the **lcd maintenance-menu disable** statement in the configuration. [PR/551546: This is a known software limitation.]
- When you enable the filter-id attribute on the RADIUS server for a particular client, none of the required 802.1X authentication rules are installed in the IPv6 database. Therefore, IPv6 traffic on the authenticated interface is not filtered; only IPv4 traffic is filtered on that interface. [PR/560381: This is a known software limitation.]
- On EX8200 switches, if OAM link fault management (LFM) is configured on a member of a VLAN on which Q-in-Q tunneling is also enabled, OAM PDUs cannot be transmitted to the Routing Engine. [PR/583053: This is a known software limitation.]
- When you reconfigure the MTU value of a next hop more than eight times without restarting the switch, the interface uses the maximum value of the eight previously configured values as the next MTU value. [PR/590106: This is a known software limitation.]
- On EX8208 and EX8216 switches that have two Routing Engines, one Routing Engine cannot be running Junos OS Release 10.4 or later while the other one is running Release 10.3 or earlier. Ensure that both Routing Engines in a single switch run either Junos OS Release 10.4 or later or Release 10.3 or earlier. [PR/604378: This is a known software limitation.]
- On EX6210 and EX8200 switches that have two Routing Engines, and on EX8200 Virtual Chassis that have two XRE200 External Routing Engine modules, you cannot issue the **commit synchronize** command from the J-Web interface. As a workaround, issue this command from the CLI. [This is a known software limitation.]

Interfaces

- EX Series switches do not support IPv6 interface statistics. Therefore, all values in the output of the **show snmp mib walk ipv6IfStatsTable** command always display a count of 0. [PR/480651: This is a known software limitation.]
- On EX8216 switches, a link might go down momentarily when an interface is added to a LAG. [PR/510176: This is a known software limitation.]
- On EX Series switches, if you clear LAG interface statistics while the LAG is down, then bring up the LAG and pass traffic without checking for statistics, and finally bring the LAG interface down and check interface statistics again, the statistics might be inaccurate. As a workaround, use the **show interfaces interface-name** command to check LAG interface statistics before bringing down the interface. [PR/542018: This is a known software limitation.]
- PoE and PoE+ cannot be configured for EX8200 member switches in an EX8200 Virtual Chassis by using the XRE200 External Routing Engine.

Configure PoE or PoE+ on each EX8200 member switch before cabling the Virtual Chassis. See [Configuring PoE \(CLI Procedure\)](#).

To configure PoE and PoE+ on an EX8200 member switch in an operational EX8200 Virtual Chassis:

1. Power off the EX8200 member switch. See [Powering Off an EX8200 Switch](#).
2. Uncable the switch from the Virtual Chassis.
3. Power on the switch. See [Powering On an EX8200 Switch](#).
4. Log in to the switch. See [Connecting an EX Series Switch to a Management Console](#).
5. Configure PoE. See [Configuring PoE \(CLI Procedure\)](#).
6. Cable the EX8200 member switch back into the EX8200 Virtual Chassis. See [Connecting an EX8200 Switch to an XRE200 External Routing Engine](#).

J-Web Interface

- In the J-Web interface, you cannot commit some configuration changes in the Ports Configuration page or the VLAN Configuration page because of the following limitations for port-mirroring ports and port-mirroring VLANs:
 - A port configured as the output port for an analyzer cannot be a member of any VLAN other than the default VLAN.
 - A VLAN configured to receive analyzer output can be associated with only one interface.

[PR/400814: This is a known software limitation.]

- In the J-Web interface, the Ethernet Switching Monitor page (Monitor > Switching > Ethernet Switching) might not display monitoring details if the switch has more than 13,000 MAC entries. [PR/425693: This is a known software limitation.]
- In the J-Web interface, in the OSPF Global Settings table in the OSPF Configuration page, the Global Information table in the BGP Configuration page, or the Add Interface window in the LACP Configuration page, if you try to change the position of columns using the drag-and-drop method, only the column header moves to the new position instead of the entire column. [PR/465030: This is a known software limitation.]
- In the J-Web interface for EX4500 switches, the Port Configuration page (Configure > Interfaces > Ports), the Port Security Configuration page (Configure > Security > Port Security), and the Filters Configuration page (Configure > Security > Filters) display features that are not supported on EX4500 switches. [PR/525671: This is a known software limitation.]
- If you insert four or more EX8200-40XS line cards in an EX8208 or EX8216 switch, the Support Information page (Maintain > Customer Support > Support Information) in the J-Web interface might fail to load because the configuration might be larger than the maximum size of 5 MB. The error message **Configuration too large to handle** is displayed. [PR/552549: This is a known software limitation.]
- The J-Web interface does not support role-based access control—it supports only users in the super-user authorization class. So a user who is not in the super-user class, such as a user with view-only permission, is able to launch the J-Web interface and is

allowed to configure everything, but the configuration fails on the switch, and the switch displays access permission errors. [PR/604595: This is a known software limitation.]

- In a mixed EX4200 and EX4500 Virtual Chassis, the J-Web interface does not list the features supported by member switches in the backup or linecard roles if those features are not also supported by the master. [PR/707671: This is a known software limitation.]

Layer 2 and Layer 3 Protocols

- On EX 3200 and EX4200 switches, MPLS on Layer 3 tagged subinterfaces and RVIs is not supported, even though the CLI enables you to commit a configuration that enables these features. [PR/612434: This is a known software limitation.]

Management and RMON

- On EX Series switches, an SNMP query fails when the SNMP index size of a table is greater than 128 bytes, because the Net SNMP tool does not support SNMP index sizes greater than 128 bytes. [PR/441789: This is a known software limitation.]
- When MVRP is configured on a trunk interface, you cannot configure connectivity fault management (CFM) on that interface. [PR/540218: This is a known software limitation.]
- The connectivity fault management (CFM) process (cfmd) might create a core file. [PR/597302: This is a known software limitation.]

Virtual Chassis

- A standalone EX4500 switch with its PIC mode set to **virtual-chassis** has less bandwidth available for network ports than an EX4500 switch with its PIC mode set to **intraconnect**. The network ports on a standalone EX4500 switch with a **virtual-chassis** PIC mode setting often do not achieve line-rate performance.

The PIC mode on an EX4500 switch can be set to **virtual-chassis** in one of the following ways:

- The switch was ordered with a Virtual Chassis module installed and thus has its PIC mode set to **virtual-chassis** by default.
- You entered the **request chassis pic-mode virtual-chassis** operational mode command to configure the switch as a member of a Virtual Chassis.

You can check the PIC mode for your EX4500 switch that has a Virtual Chassis module installed by entering the **show chassis pic-mode** command.

You should always set the PIC mode on a standalone EX4500 switch to **intraconnect**. Set the PIC mode to **intraconnect** by entering the **request chassis pic-mode intraconnect** operational mode command.

[This is a known limitation.]

- The automatic software update feature is not supported on EX4500 switches that are members of a Virtual Chassis. [PR/541084: This is a known software limitation.]

- When an EX4500 switch becomes a member of a Virtual Chassis, it is assigned a member ID. If that member ID is a nonzero value, then if that member switch is downgraded to a software image that does not support Virtual Chassis, you cannot change the member ID to 0. A standalone EX4500 switch must have a member ID of 0. The workaround is to convert the EX4500 Virtual Chassis member switch to a standalone EX4500 switch before downgrading the software to an earlier release, as follows:

1. Disconnect all Virtual Chassis cables from the member to be downgraded.
2. Convert the member switch to a standalone EX4500 switch by issuing the **request virtual-chassis reactivate** command.
3. Renummer the member ID of the standalone switch to 0 by issuing the **request virtual-chassis renumber** command.
4. Downgrade the software to the earlier release.

[PR/547590: This is a known software limitation.]

- When you add a new member switch to an existing EX4200 Virtual Chassis, EX4500 Virtual Chassis, or mixed EX4200 and EX4500 Virtual Chassis in a ring topology, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. The member switch returns to the operational state with no user intervention. Network traffic to the member switch is dropped during the downtime. To avoid this issue, follow this procedure:

1. Cable one dedicated or user-configured Virtual Chassis port (VCP) on the new member switch to the existing Virtual Chassis.
2. Power on the new member switch.
3. Wait for the new switch to become operational in the Virtual Chassis. Monitor the **show virtual-chassis** command output to confirm that the new switch is recognized by the Virtual Chassis and is in the Prsnt state.
4. Cable the other dedicated or user-configured VCP on the new member switch to the Virtual Chassis.

[PR/591404: This is a known software limitation.]

Related Documentation

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

Outstanding Issues in Junos OS Release 12.1 for EX Series Switches

The following are outstanding issues in Junos OS Release 12.1R7 for EX Series switches. The identifier following the description is the tracking number in our bug database.

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

Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers](#).

Access Control and Port Security

- When an EX Series switch is configured as a Junos OS enforcer on an IC Series Unified Access Control Appliance, the Odyssey Access Client (OAC) status might change from open/authenticating to open and authenticated. [PR/742369]
- If a UAC infranet controller is unreachable, an 802.1X (dot1x) interface might not be able to access the server-fail VLAN. [PR/781586]

Converged Networks (LAN and SAN)

- When you configure class-of-service drop profiles, the commit operation might fail and might display the message **Missing mandatory statement: 'drop-probability'**. [PR/807885]

Ethernet Switching and Spanning Trees

- If the bridge priority of a VSTP root bridge is changed such that this bridge becomes a nonroot bridge, the transition might take more than 2 minutes, and you might see a loop during the transition. [PR/661691]
- When using VSTP, if you try to enable all VLANs on a physical interface that is a member of all the VLANs, a configuration error might be displayed. For more information, see [“Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches” on page 57](#). [PR/736488]
- You cannot configure a VLAN whose name contains a hyphen (-). As a workaround, use an underscore (_) in the name instead. [PR/753090]

Hardware

- On EX4200 switches, the serial number displayed in the log output does not include the leading 0 that is shown on the label. This problem is just a display issue—there is no functional impact. [PR/815950]
- On EX4550 switches, the backlight on the LCD panel does not turn on. [PR/820473]

High Availability (HA) and Resiliency

- On an EX4200 Virtual Chassis, performing an NSSU from Junos OS Release 12.1R5 to Junos OS Release 12.2R3 might cause a vmcore core file in linecard members. [PR/844519]
- On EX8200 Virtual Chassis, a nonstop software upgrade (NSSU) might fail. [PR/871288]

Infrastructure

- On EX8208 switches, when a line card that has no interface configurations and is not connected to any device is taken offline using the **request chassis fpc-slot slot-number offline** command, the Bidirectional Forwarding Detection process (bfd) starts and stops repeatedly. The same bfd process behavior occurs on a line card that is connected to a Layer 3 domain when another line card that is on the same switch and is connected to a Layer 2 domain is taken offline. [PR/548225]
- The output of the **show system users no-resolve** command displays the resolved hostname. [PR/672599]
- When you configure a static route that has two multihop paths, BFD might become unstable, and the routing protocol process (rpd) might crash. [PR/701966]
- When a core file is generated, the master Routing Engine might stop operating, the console might stop responding, and all line cards might go offline. [PR/707527]
- On XRE200 External Routing Engine Virtual Chassis, when you upgrade the software while traffic is transiting the device, a kernel panic might occur. [PR/742727]
- When a hostname is added as an NTP server, it is resolved to an IP address before it is added to the configuration file. If the resolved IP address becomes unreachable for any reason, the switch cannot reach the NTP server. Although public NTP servers can resolve a hostname into different IP addresses, Junos OS cannot leverage this facility because the configuration file does not accept the hostname as a string. [PR/755591]
- On EX3300 switches, when you execute the **request system scripts delete** command, a vmcore file might be created, and the switch might reboot. [PR/768570]
- On aggregated Ethernet (ae) interfaces, the LLDP might not work. [PR/781814]
- On EX4550 switches, if you configure the management (me0) interface and a static route, the switch is unable to connect to a gateway. [PR/786184]
- On EX8200 Virtual Chassis, an RVI might send gratuitous ARP requests. [PR/848852]
- On EX Series switches, if you have configured a LAG with link protection, ingress traffic does not pass through the backup port. [PR/886205]

Interfaces

- When you disable a static LAG on an aggregated Ethernet (ae) interface, Ethernet ring protection traffic traveling in one direction might be lost for 3 to 5 seconds, and traffic traveling in the other direction might contain extra packets. [PR/703091]
- On aggregated Ethernet (ae) interfaces, LLDP might not work. [PR/781814]
- On EX4550 switches, link autonegotiation does not work on 1-Gb SFP interfaces. [PR/795626]
- When interfaces over which traffic is actively flowing flap, an Ethernet switching process (eswd) core file might be created. [PR/802149]
- An interface on an EX4550-32F switch might go up and down randomly even when no cable is plugged in. [PR/803578]
- On EX3300 switches, when you configure VRRP with MD5 authentication with the **preempt** option on an IRB interface, a vmcore file might be created. As a workaround, delete the **preempt** option and disable MD5 authentication for VRRP. [PR/808839]
- On EX8200 Virtual Chassis, nonstop software upgrade (NSSU) with the no-reboot option is not supported. [PR/821811]
- If you delete an IPv6 configuration on an RVI, ARP requests might not be trapped to the CPU and are not resolved. As a workaround, delete the RVI and then reconfigure it, or reboot the switch after you delete the IPv6 configuration. [PR/826862]
- Multicast packets might be lost when the user switches from one IPTV channel to another. [PR/835538: This issue has been resolved.]
- On EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6210 switches, a firewall filter with **family** set to **ethernet-switching** and configured for IPv4 blocks specific transit IPv6 traffic if the **ether_type** match condition in the filter is not explicitly set to **ipv4**. As a workaround, set **ether_type** to **ipv4** in the filter. [PR/843336]

J-Web Interface

- In the J-Web interface, in the Port Security Configuration page, you are required to configure **action** when you configure **MAC limit** even though configuring an action value is not mandatory in the CLI. [PR/434836]
- If you configure an IPv6 address for a VLAN in the J-Web interface, you cannot then edit the VLAN configuration. [PR/466633]
- When a large number of static routes are configured and you have navigated to pages other than page 1 in the Route Information table in the Static Routing monitoring page in the J-Web interface (Monitor > Routing > Route Information), changing the Route Table to query other routes refreshes the page but does not return to page 1. For example, if you run a query from page 3 and the new query returns very few results, the Results table continues to display page 3 and shows no results. To view the results, navigate to page 1 manually. [PR/476338]

- When you use an HTTPS connection in the Microsoft Internet Explorer browser to save a report from the following pages in the J-Web interface, the error message **Internet Explorer was not able to open the Internet site** is displayed in the following pages:

- Files page (Maintain > Files)
- History page (Maintain > Config Management > History)
- Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
- Static Routing page (Monitor > Routing > Route Information)
- Support Information page (Maintain > Customer Support > Support Information)
- View Events page (Monitor > Events and Alarms > View Events)

[PR/542887]

- When you open a J-Web session using HTTPS, enter a username and password and then click the **Login** button, the J-Web interface takes 20 seconds longer to launch and load the Dashboard page than it does if you use HTTP. [PR/549934]
- In the J-Web interface, you cannot upload a software package using the HTTPS protocol. As a workaround, use either the HTTP protocol or the CLI. [PR/562560]
- In the J-Web interface, the link status might not be displayed correctly in the Port Configuration page or the LACP (Link Aggregation Control Protocol) Configuration page if the Commit Options preference is set to single commit (the Validate configuration changes option). [PR/566462]
- If you have accessed the J-Web interface using an HTTPS connection through the Microsoft Internet Explorer Web browser, you might not be able to download and save reports from some pages on the Monitor, Maintain, and Troubleshoot tabs. Some affected pages are at these locations:
 - Maintain > Files > Log Files > Download
 - Maintain > Config Management > History
 - Maintain > Customer Support > Support Information > Generate Report
 - Troubleshoot > Troubleshoot Port > Generate Report
 - Monitor > Events and Alarms > View Events > Generate Report
 - Monitor > Routing > Route Information > Generate Report

As a workaround, use the Mozilla Firefox Web browser to download and save reports using an HTTPS connection. [PR/566581]

- If you access the J-Web interface using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR/603669]
- In the J-Web interface, HTTPS access might work with an invalid certificate. As a workaround, after you change the certificate, issue the **restart web-management** command to restart the J-Web interface. [PR/700135]

- On EX4500 Virtual Chassis, if you use the CLI to switch from **virtual-chassis** mode to **intraconnect** mode, the J-Web dashboard might not list all the Virtual Chassis hardware components, and the image of the master and backup switch chassis might not be visible after an autorefresh occurs. The J-Web dashboard also might not list the vcp-0 and vcp-1 Virtual Chassis ports in the rear view of an EX4200 switch (in the linecard role) that is part of an EX4500 Virtual Chassis. [PR/702924]
- After you have disabled the LCD Maintenance Menu and rebooted the switch, the EZSetup option might be available. [PR/707279]
- In the PoE Monitoring page (Monitor > PoE), the Telemetry Graph shows no data for power and voltage. [PR/723564]
- On EX2200-C switches, if you have changed the media type and committed the change, the Ports Configuration page (Configure > Interfaces > Ports) might not list the uplink port. [PR/742847]
- In the J-Web interface, you cannot configure the TCP fragment flag for a firewall filter in the Filters Configuration page (Configure > Security > Filters). [PR/756241]
- In the J-Web interface, you cannot delete a term from a firewall filter and simultaneously add a new term to that filter in the Filters Configuration page (Configure > Security > Filters). [PR/769534]
- After you remove or reboot a Virtual Chassis member (either the backup or a line card), when you click other members in the J-Web interface, the chassis view for those members might not expand, and the dashboard might log the following error: **stackImg is null or not an object**. [PR/771415]
- If a Virtual Chassis contains more than six members, the Support Information page (Maintain > Customer Support > Support information) might not load. [PR/777372]
- Some component names shown by the tooltip in the Temperature in Health Status Panel of the dashboard might be truncated. As a result, you might see many components that have the same name displayed. For example, the components GEPHY Front Left, GEPHY Front Middle, and GEPHY Front Right might all be displayed as GEPHYFront. [PR/778313]
- In a mixed EX4200 and EX4500 Virtual Chassis, the master chassis view might display the temperature indicator of the backup. [PR/783052]
- On an EX Series Virtual Chassis that has more than five members, logging in to the J-Web interface dashboard might take more than 30 seconds. [PR/785300]
- In the J-Web interface on EX4500 and EX4550 switches, you can configure temporal and exact-temporal buffers, which are not supported by the Junos OS software. [PR/796719]
- In an EX4550 mixed-mode Virtual Chassis in which an EX4550 switch is the master and at least one Virtual Chassis member supports Power over Ethernet (PoE), if you click **Configure > POE** and then click another tab, a javascript error might be displayed. [PR/797256]
- In the J-Web interface on EX4200 switches, the model number specific to Juniper Networks is displayed instead of the IBM-specific model number in the header portion

of the IBM-OEM J-Web device. As a workaround, refer to the dashboard system information panel to view the model number specific to IBM. [PR/798447]

- In the J-Web interface on EX4550 switches, if you are using in-band management and select **EZSetup**, the error message **undefined configuration delivery failed** is displayed even though the configuration has been successfully committed. [PR/800523]
- In the J-Web interface on EX2200 switches, the flash memory utilization graph on the Dashboard displays an incorrect value. As a workaround, view the flash memory utilization from the System Information page (Monitor > System View > System Information) and then click the **Storage Media** tab. [PR/823795]

Layer 2 and Layer 3 Protocols

- No CLI command is available to verify that NSB is enabled. [PR/613452]
- If you try to configure a Layer 3 protocol such as IS-IS, OSPF, or RIP on a Layer 2 interface (that is, an interface configured with the family **ethernet-switching**), the commit operation fails. [PR/729993]
- In a scenario with 120 VLANs, 14,000 MAC addresses, 19,000 (s,g) entries, and 4200 (*,g) entries, if VRRP switches a few times from master to backup or backup to master on an STP root bridge, it can cause the rpd to crash. Also, MVPN configured on a scaled setup can generate an rpd core file. [PR/801104]
- After an NSSU operation, OSPF might remain in the INIT state because the flooding entry is not programmed correctly. [PR/811178]
- EX4500 and EX4550 switches do not support enhanced transmission selection (ETS) TLV negotiations from peer switches or routers that originate TLV elements. [PR/830205]
- When an invalid subnet configuration on a multicast group is committed using the **commit** or **commit check** command, the routing protocol daemon (rpd) crashes and creates a core file, and the commit operation fails. [PR/856925]

Management and RMON

- The CFM process (cfmd) might create a core file. [PR/597302]
- On EX8200 Virtual Chassis, when you perform an snmpwalk operation on the jnxPsmMIB, the output shows details only for the power supplies on a single member switch. [PR/689656]
- When an SNMP string is longer than 30 characters, it is not displayed in Junos OS command output. [PR/781521]
- EX4200 and EX4500 switches support 64 aggregated Ethernet interfaces even though the hardware can support 111 interfaces. [PR/746239]

Software Upgrade and Installation

- On EX4200 switches, when you upgrade Junos OS, the software build-time date might be reset. [PR/742861]

- EX4550 switches might not load the configuration file after you perform an automatic image upgrade. [PR/808964]

Virtual Chassis

- When you remove the hard drive from an XRE200 External Routing Engine, an SNMP trap and a system alarm might not be generated. [PR/710213]
- On XRE200 External Routing Engines, when you issue the **show chassis hardware (<get-chassis-inventory>)** command, duplicate occurrences of the <name> and <serial-number> tags under the <chassis> tag might result in malformed XML output. [PR/772507]
- The SNMP MIB Walk displays unwanted data for newly added objects. [PR/791848]
- On EX8200 Virtual Chassis, if you are using the wizard in the Firefox version 3.x browser, if you have selected more than six port pairs from the same member for conversion, the wizard might display the incorrect port conversion status. Also, if you double-click **Next** after deleting an active member on the members page, the J-Web interface might stop working. [PR/796584]
- On EX8200 Virtual Chassis, if you add and then delete a firewall filter for traffic that enters on one Virtual Chassis member and is transmitted out another member, IPv6 traffic might be dropped. If the ingress and egress interfaces are on the same member, the firewall filter works correctly. [PR/803845]
- On EX4550 Virtual Chassis, the **show chassis environment power-supply-unit** operational mode command does not show the power supply status of all member interfaces. Use the **show chassis hardware** command instead. [PR/817397]
- If you upgrade an EX8200 Virtual Chassis from Junos OS Release 11.1S6 or Release 11.2 to Release 12.1R4, there might be an outage of more than 300 seconds. [PR/818510]

Related Documentation

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

Resolved Issues in Junos OS Release 12.1 for EX Series Switches

The following are the issues that have been resolved in Junos OS Release 12.1 for EX Series switches. The identifier following the descriptions is the tracking number in our bug database.

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

Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers](#).

- [Issues Resolved in Release 12.1R1 on page 33](#)
- [Issues Resolved in Release 12.1R2 on page 42](#)
- [Issues Resolved in Release 12.1R3 on page 45](#)
- [Issues Resolved in Release 12.1R4 on page 47](#)
- [Issues Resolved in Release 12.1R5 on page 50](#)
- [Issues Resolved in Release 12.1R6 on page 51](#)
- [Issues Resolved in Release 12.1R7 on page 53](#)

Issues Resolved in Release 12.1R1

The following issues have been resolved since Junos OS Release 11.4. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- On EX2200, EX3300, and EX6200 switches, and on EX8200 Virtual Chassis, NetBIOS snooping does not work. [PR/706588: This issue has been resolved.]
- When you enable LLDP-MED autonegotiation on an EX Series switch, the autonegotiation bit in the LLDP-MED packet is set to not-supported, which might cause IP phones to discard LLDP-MED packets received from the switch. [PR/708752: This issue has been resolved.]
- If incoming LLDP packets contain multiple management address TLVs, EX Series switches discard them. [PR/718781: This issue has been resolved.]
- When DHCP snooping information is not learned, ARP request packets might add the following message to the system log (syslog) file: **ESWD_DAI_FAILED: 3 (null) received, interface**. [PR/719751: This issue has been resolved.]
- When an EX Series switch is reauthenticating users using 802.1X (dot1x), if the switch loses reachability to the RADIUS server, the dynamic filters that were installed when the same user was previously authenticated are not cleared, resulting in traffic issues. [PR/721124: This issue has been resolved.]
- On EX Series switches running Junos OS Release 11.x, LLDP packets might not be generated out of interfaces that are part of a LAG, causing LLDP neighbors not to form. As a workaround, follow these steps:
 1. Delete the LLDP-MED configuration.
 2. Commit the configuration.
 3. Delete the LLDP configuration.
 4. Commit the configuration.

5. Configure LLDP again.
6. Commit the configuration.
7. Optionally, configure LLDP-MED again.
8. Commit the configuration.

[PR/727627: This issue has been resolved.]

- EX3200 switches might repeatedly create 802.1X core files. As a workaround, if access accounting is enabled, disable it by issuing the **deactivate access profile *profile-name* accounting** command. [PR/739921: This issue has been resolved.]

Device Security

- If storm control is enabled, LACP might stop and then restart when Layer 2 packets are sent at a high speed. As a workaround, disable storm control for all multicast traffic on aggregated Ethernet interfaces by issuing the **set ethernet-switching-options storm-control interface *interface-name* no-multicast** command. [PR/575560: This issue has been resolved.]
- You cannot configure the level for storm control. [PR/734307: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- On EX Series switches, when you remove a VLAN that has a VLAN ID and then add the same VLAN ID but with a different VLAN name, the Ethernet switching process (eswd) might create a core file. [PR/668210: This issue has been resolved.]
- On an EX4200 switch, when you disable a Q-in-Q interface on which you have configured a large number (more than 500) of VLAN swap rules, control traffic might be affected for about 10 minutes. During this time, the forwarding process (pfem) can consume up to 98 percent of the CPU. The system resumes its normal state after the forwarding process completes its processing. [PR/678792: This issue has been resolved.]
- When you enable VLANs and Q-in-Q tunneling on a switch, the switch drops packets, and no MAC address learning occurs. [PR/685481: This issue has been resolved.]
- On a LAG interface on which Q-in-Q tunneling is enabled on a VLAN, packets entering the LAG might be dropped. As a workaround, explicitly configure the VLAN to allow the desired traffic. [PR/699940: This issue has been resolved.]
- When ingress and egress ports are on different member switches and a packet is routed from the default routing instance to another forwarding instance type, the VLAN ID might be modified in such a way that the traffic is redirected to the default routing instance for subsequent routing. [PR/721436: This issue has been resolved.]
- When you configure the same VLAN ID on both interface VLAN tagging and global tagging, ARP entries cannot be resolved on the VLAN interface. [PR/722815: This issue has been resolved.]
- RVIs might use the system MAC address instead of using the MAC address one greater than the system MAC address (that is, system MAC address + 1), and Layer 3 ports

might use their hardware MAC address instead of using the system MAC address. [PR/723643: This issue has been resolved.]

- When you change the spanning-tree protocol from RSTP or VSTP to MSTP, the Ethernet switching process (eswd) might create a core file. [PR/725436: This issue has been resolved.]

Firewall Filters

- On EX8200 switches, if you configure a **discard** term on an egress firewall filter, the filter might not block ARP broadcast packets. [PR/672621: This issue has been resolved.]
- For two-rate, three-color policers, the egress traffic might not flow at the configured peak information rate (PIR). [PR/687564: This issue has been resolved.]
- When you configure VLAN ID translation when using Q-in-Q tunneling, if you apply a tricolor marking (TCM) policer to the Q-in-Q interface, a Packet Forwarding Engine (pfem) core file might be created. [PR/688438: This issue has been resolved.]
- In an EX8200 Virtual Chassis that is configured with an implicit **deny** statement and that has VCCP traffic flowing through 10-Gigabit Ethernet ports configured as Virtual Chassis ports (VCPs), if you apply a loopback filter, then the line cards of member 0 and member 1 might lose contact with the master Routing Engine. [PR/688983: This issue has been resolved.]
- Firewall rules might not be installed in the TCAM, and you might see the following error message: **dfw_grph_merge_dfw_bind: rules for filter ACL will not be installed**. [PR/689288: This issue has been resolved.]
- When you configure a **syslog** action in a firewall filter on the me0 interface, an EX2200 switch might crash when you commit the configuration. [PR/694602: This issue has been resolved.]
- If you configure both a regular and a firewall filter-based analyzer, the traffic from the regular analyzer might exit from the output port you configured for the firewall filter-based analyzer. [PR/724795: This issue has been resolved.]
- If you configure a firewall filter on a loopback interface whose last term is **deny all**, static routes filtered with the **reject** action reach the CPU, and packets that failed the multicast trap and RPF are implicitly allowed to reach the CPU. [PR/740641: This issue has been resolved.]

Hardware

- On XRE200 External Routing Engines, the output of the **show chassis hardware** command might contain duplicate Routing Engine inventory information for members 8 and 9. [PR/663272: This issue has been resolved.]
- On EX6210 switches, traffic might not exit from the 10-Gigabit Ethernet interfaces on the Routing Engines. [PR/669330: This issue has been resolved.]
- For certain vendors' SFPs with Juniper part number 740-021308 and types SFP+ 10GE-SR, SFP+ 10GE-LR, or SFP+ 10GE-ER, when the low-power threshold is crossed, the power-low warning alarm is not set on extra-scale and PoE line cards. [PR/683732: This issue has been resolved.]

- On EX4500 switches, the LCD panel might not list the ADM (administrative status) or DPX (duplex) options in the Idle menu. Also, you might not be able to navigate through the status mode LEDs by pressing Enter. [PR/692341: This issue has been resolved.]
- On EX4200 switches, the EZSetup menu is not displayed on the LCD panel after you set the switch to the factory-default configuration. [PR/712322: This issue has been resolved.]
- On EX8208 switches, when the SRE module is in the spare state and you configure it to go offline and then come back online again, the module's ST LED does not turn back on. [PR/724455: This issue has been resolved.]

High Availability

- When you perform an NSSU operation on an EX8200 Virtual Chassis, if you do not include the **reboot** option when you request that the NSSU have the switch perform an automatic reboot, the upgrade might hang indefinitely after the Junos OS images have been pushed to the master Routing Engine. [PR/692422: This issue has been resolved.]
- After a GRES operation, clone routes might move into the reject state. [PR/724729: This issue has been resolved.]

Infrastructure

- The system log (syslog) files might contain the message **Juniper syscall not available**. This message is harmless, and you can ignore it. [PR/519153: This issue has been resolved.]
- The system log (syslog) file might contain the following message: **/var: filesystem full**. [PR/600145: This issue has been resolved.]
- On EX Series switches, the **request system snapshot** command mistakenly includes the **as-primary** option. [PR/603204: This issue has been resolved.]
- If you remove or change interfaces soon after completing an NSSU operation, the multicast snooping process (mcsnoopd) might create a core file. [PR/662065: This issue has been resolved.]
- Layer 3 next-hop entries might remain queued in the kernel of the backup Routing Engine and might never be installed in the forwarding table. [PR/670799: This issue has been resolved.]
- On EX8200 switches, when you run a failover operation on the Routing Engines, a vmcore file might be created. [PR/678465: This issue has been resolved.]
- The management process (mgd) might create a core file when reading long lines. For example, this can happen when you are displaying a Junos OS configuration file that contains long lines. When the mgd crashes, the command that you were executing does not finish, and the following errors appear in the **messages** file:
%KERN-3-BAD_PAGE_FAULT: pid 57182 (mgd), uid 0: pc 0x8870ab92 got a write fault at 0x8488000, x86 fault flags = 0x6 and %KERN-6: pid 57182 (mgd), uid 0: exited on signal 11 (core dumped). [PR/679992: This issue has been resolved.]

- On EX4500 switches, ICMPv6 packets might transit the Routing Engine even though IPv6 is not configured. [PR/682953: This issue has been resolved.]
- On an XRE200 External Routing Engine, the rescue configuration might not get synchronized with the backup external Routing Engine. [PR/687797: This issue has been resolved.]
- During a GRES operation between an EX4200 and an EX4500 switch, a Packet Forwarding Engine (pfem) core file might be created. [PR/688618: This issue has been resolved.]
- You might not be able to commit a configuration on an XRE200 External Routing Engine, and the switch might display the error **could not save to juniper.save+**. [PR/689764: This issue has been resolved.]
- An EX4200 switch might stop forwarding traffic, and a Packet Forwarding Engine (pfem) core file might be created. [PR/691504: This issue has been resolved.]
- When the same MAC address is learned on both the primary and community VLANs, an Ethernet switching process (eswd) core file might be created. [PR/693942: This issue has been resolved.]
- On EX4200 switches, if you connect and then disconnect the cable to the port on which BFD is running, a software forwarding process (sfid) core file might be created. [PR/694150: This issue has been resolved.]
- EX Series switches might not learn the MAC addresses of directly connected devices. [PR/695280: This issue has been resolved.]
- On EX3200, EX4200, EX4500, EX6200, EX8208, and EX8216 switches, the root user is allowed to telnet into the me0 interface, which does not comply with the default Junos OS behavior, as documented in [Telnet to JUNOS router fails with root login](#) and [Connecting and Configuring an EX Series Switch \(CLI Procedure\)](#). [PR/695346: This issue has been resolved.]
- On EX Series switches, when you are configuring DHCP option 82, the **use-interface-description** statement, which uses the interface description rather than the interface name (the default) in the circuit ID or remote ID value in the DHCP option 82 information, does not work. [PR/695712: This issue has been resolved.]
- When you commit a configuration, the following message might be placed in the system log (syslog) file: **kernel: PFE not configured, Juniper syscall not available**. This message is harmless, and you can ignore it. [PR/696471: This issue has been resolved.]
- When the switch is performing 802.1X (dot1x) authentication using MAC RADIUS, you might see the following message in the system log (syslog) file: **kmern type temp using 57344K, exceeding limit 57344K**. [PR/697815: This issue has been resolved.]
- When you add a new VLAN on a switch on which you have also configured loopback and other IPv4 firewall filters, a pfem core file might be created that contains the message **No space available in tcam**. [PR/701779: This issue has been resolved.]
- On EX8200 switches that have IPv6 and IPv4 configured on routing instances, when many interfaces go down and come back up repeatedly, the next-hop programming for some routes might fail, which can cause disruptions in traffic. [PR/701985: This issue has been resolved.]

- In rare cases, EX3300 switches running a Junos OS release earlier than Release 11.3R4 or Release 11.4R2 might experience some traffic loss or a link failure due to non-user-configurable settings that are not optimized. The issue can be resolved by upgrading to one of the following Junos OS releases:

- Junos OS Release 11.3R4 and later
- Junos OS Release 11.4R2 and later
- Junos OS Release 12.1R1 and later

[PR/703147: This issue has been resolved.]

- When you run the **commit check** command, the word *operation* in the command output might be misspelled. [PR/704910: This issue has been resolved.]
- If the egress interface is a trunk interface, frames that exceed 9216 bytes might not be fragmented. [PR/705905: This issue has been resolved.]
- If you power off and then restart an EX Series switch, the chassis process (chassisd) might not restart. This problem might also occur in switches that have redundant power supplies. [PR/708872: This issue has been resolved.]
- If you perform an NSSU operation that includes the **reboot** option, some traffic loss might occur. [PR/717662: This issue has been resolved.]
- After a MAC address moves, the ARP index might not point to the correct MAC address. As a workaround, run the **clear ethernet-switching table** command. [PR/718698: This issue has been resolved.]
- A destination with two equal-cost next hops might not be installed in the Packet Forwarding Engine when a virtual management Ethernet (VME) interface is configured and a default route with a reachable next hop is present. As a workaround, remove and then re-add one of the two equal-cost next hops. [PR/719745: This issue has been resolved.]
- On EX4200 switches, after you issue the **request system zeroize media** command, you might not be able to establish a connection with the switch using SSH, and you might not be able to issue the **commit** command on the switch. [PR/723918: This issue has been resolved.]
- New hosts and routes might not be installed in the Packet Forwarding Engine, and the system log (syslog) files might contain the following message: **Failed to jtm_alloc for mrvl_rt_nh_uc_install**. [PR/726043: This issue has been resolved.]
- On EX8200 switches, after you issue the **request system zeroize media** command, the line cards might not come online. [PR/728082: This issue has been resolved.]
- In the Login page and in the Help mapping file in the J-Web interface, the copyright date is set to 2011. [PR/731790: This issue has been resolved.]
- The Ethernet switching process (eswd) might create a core file. [PR/732263: This issue has been resolved.]

Interfaces

- The displayed bandwidth on unit 0 of an interface might be incorrect. As a workaround, configure an interface speed on unit 0. [PR/471628: This issue has been resolved.]
- When you configure the member interfaces of an aggregated Ethernet interface before you configure the aggregated Ethernet (ae) interface using **set** commands from the CLI, the aggregated Ethernet interface does not receive MAC updates. As a workaround, configure the aggregated Ethernet interface first, and then configure the member interfaces. [PR/680913: This issue has been resolved.]
- When you configure an aggregated Ethernet interface with LACP in fast mode, the logical interface of the aggregated interface might flap after a GRES operation. [PR/686585: This issue has been resolved.]
- When 10-Gigabit Ethernet interfaces flap frequently, a routing protocol process (rpd) core file might be created. [PR/692126: This issue has been resolved.]
- EX Series switches do not show the control packet counters in both directions on the logical units of aggregated Ethernet (ae) interfaces. [PR/693202: This issue has been resolved.]
- Interfaces might not come up, and the system log (syslog) file might contain the message **DCD_CONFIG_WRITE_FAILED:configuration write failed for an RT ADD: Cannot allocate memory**. [PR/697300: This issue has been resolved.]
- If you set an interface speed at the **[edit interfaces interface-range]** hierarchy level, this speed might not be applied to the individual interface. If you set different interface speeds directly on the interface and at the **[edit interfaces interface-range]** hierarchy level, both settings are applied on the interface. [PR/697478: This issue has been resolved.]
- On EX4500 switches with an interface in loopback mode, BPDUs might be processed instead of being dropped as expected, generating topology change notifications (TCNs). As a workaround, disable the interface that is in loopback mode. [PR/698077: This issue has been resolved.]
- When you configure the **no-preempt** and **interface-tracking** options on a switch that is a VRRP master router, if the VRRP mastership is taken over by a switch that is a VRRP backup router and the tracking interface on the original master router goes down, then if the tracking interface on the original master router comes back up and the master's original priority is restored, the new master's mastership might transition to the original master router. [PR/699243: This issue has been resolved.]

J-Web Interface

- In the J-Web interface, the dashboard does not display the uplink ports or uplink module ports unless transceivers are plugged into the ports. [PR/477549: This issue has been resolved.]
- If you have created dynamic VLANs by enabling MVRP from the CLI, then in the J-Web interface, the following features do not work with dynamic VLANs and static VLANs:

- In the Port Configuration page (Configure > Interface > Ports)—Port profile (select the interface, click **Edit**, and select **Port Role**) or the VLAN option (select the interface, click **Edit**, and select **VLAN Options**).
- VLAN option in the LACP (Link Aggregation Control Protocol) Configuration page (Configure > Interface > Link Aggregation)—Select the aggregated interface, click **Edit**, and click **VLAN**.
- In the 802.1X Configuration page (Configure > Security > 802.1x)—VLAN assignment in the exclusion list (click **Exclusion List** and select **VLAN Assignment**) or the move to guest VLAN option (select the port, click **Edit**, select **802.1X Configuration**, and click the **Authentication** tab).
- Port security configuration (Configure > Security > Port Security).
- In the Port Mirroring Configuration page (Configure > Security > Port Mirroring)—Analyzer VLAN or ingress or egress VLAN (click **Add** or **Edit** and then add or edit the VLAN).

[PR/669188: This issue has been resolved.]

- In the J-Web interface, if you navigate to any of the top panel tabs (such as Configure, Monitor, Maintain, and Troubleshoot) and then click **Help > Help Contents**, you might be directed to an undefined page. To display the correct help for a feature, first click the menu item corresponding to the feature to load the page, and then click **Help > Help Contents**. [PR/684958: This issue has been resolved.]
- In the J-Web interface, if you discard any available MIB profile, file, or predefined object from accounting-options in the Point and Click CLI Configuration page (Configure > CLI Tools > Point and Click CLI), the J-Web session times out. As a workaround, perform the same operation from the CLI. [PR/689261: This issue has been resolved.]
- On EX4500 switches, you cannot configure BGP in the BGP Configuration page (Configure > Routing > BGP). [PR/699308: This issue has been resolved.]
- In the J-Web interface, the dashboard might not be displayed. [PR/700274: This issue has been resolved.]
- In the J-Web interface on an EX4500 Virtual Chassis, if you configure four or more Virtual Chassis members in the Support Information page (Maintain > Customer Support > Support Information), you might see the error **Configuration of switch is too large**. [PR/704992: This issue has been resolved.]
- On SRX210 Services Gateways and EX Series switches, if you use the CLI to delete a DHCP pool, the J-Web interface Monitor page (Monitor > Services > DHCP > Pools) might not display the correct value in the Excluded address field. Instead, it might display the text **[objec object]** in the table. [PR/723555: This issue has been resolved.]
- On the J-Web dashboard the Total number of ports field in the Capacity Utilization page might show incorrect values for a mixed EX4200 and EX4500 Virtual Chassis. As a workaround, use the **show chassis hardware | match PIC | except Virtual** command to display the correct values. [PR/734766: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- On EX Series switches and M Series routers, IPv6 neighbor unreachability detection does not work. As a workaround, use the **clear ipv6 neighbor** command to initiate neighbor detection. [PR/613230: This issue has been resolved.]
- When a BGP interface is flapping quickly, BGP might unnecessarily withdraw prefixes even when a good route to that prefix still exists. [PR/677191: This issue has been resolved.]
- On EX3200 and EX4200 switches, no counters are available for MPLS statistics for circuit cross-connects (CCCs) because of a hardware limitation. As a result, no counters are incremented in MPLS statistics files for LSPs that are used for CCCs. [PR/724371: This issue has been resolved.]

Management and RMON

- When you use the snmpwalk application to get information about switch interfaces, it returns information about incorrect interfaces. [PR/664940: This issue has been resolved.]
- EX Series switches might not send jnxMIMst traps. [PR/707141: This issue has been resolved.]

Multicast Protocols

- You might not be able to delete stale multicast routes even though no corresponding (S, G) traffic exists. [PR/674419: This issue has been resolved.]
- Approximately every 300 seconds, a multicast route entry is deleted and added back again, resulting in a traffic loss for about 1 to 3 seconds. [PR/698129: This issue has been resolved.]

Software Upgrade and Installation

- When you upgrade Junos OS from a release earlier than Release 10.4R3, in which resilient dual-root partitioning was introduced, to a later release that supports resilient dual-root partitioning, the time required for the upgrade to finish is longer than upgrading either between two releases that are earlier than Release 10.4R3 or between two releases that are Release 10.4R3 or later. [PR/683337: This issue has been resolved.]
- On EX3300 switches, when you load the factory default settings, the last two ports of the uplink ports are configured as Virtual Chassis ports (VCPs). If you convert these ports to network ports, they might not pass traffic. As a workaround, reboot the switch after converting the ports. [PR/685300: This issue has been resolved.]
- After you upgrade Junos OS, a software forwarding process (sfid) core file might be created. [PR/691958: This issue has been resolved.]
- After you upgrade EX8208 switches from Junos OS Release 10.0S17 to Release 12.1B2, a Junos OS panic might occur. [PR/748744: This issue has been resolved.]

Virtual Chassis

- On EX8200 Virtual Chassis, the link status of an aggregated Ethernet (ae) interface managed by LACP goes down and comes back up when a GRES operation is performed between the XRE200 External Routing Engines. This switchover might have been initiated from the Junos OS CLI or because of a failure of the master Routing Engine. [PR/599772: This issue has been resolved.]
- On EX8200 Virtual Chassis, if the topology is formed such that ingress multicast traffic is routed first to the RP and then returns to the Virtual Chassis for egress through Layer 2 multicast, the multicast traffic is forwarded only to receivers connected to the Virtual Chassis member in which the returned multicast traffic is received. Multicast traffic is not forwarded to other receivers in other Virtual Chassis members. [PR/666355: This issue has been resolved.]
- When you interconnect EX4200 and EX4500 switches to form a mixed EX4200 and EX4500 Virtual Chassis, the switches might fail to form a Virtual Chassis. [PR/681072: This issue has been resolved.]
- On EX8200 Virtual Chassis, LACP and all Layer 3 protocols flap constantly when an EX8200 member switch's backup Routing Engine is rebooted. Because of this issue, the connection between the master XRE200 External Routing Engine and the EX8200 member switch might be lost. [PR/700295: This issue has been resolved.]
- In mixed EX4200 and EX4500 Virtual Chassis, if you configure class of service, some traffic might be dropped because it is mapped to the incorrect queue. [PR/711071: This issue has been resolved.]
- When the ingress and egress ports are on different member switches and a packet is routed from the default routing instance to another forwarding instance type, the VLAN ID might be modified in such a way that the traffic is redirected to the default routing instance for subsequent routing. [PR/721436: This issue has been resolved.]

Issues Resolved in Release 12.1R2

The following issues have been resolved since Junos OS Release 12.1R1. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- When you configure MVRP, the LLDP process might create a core file as the result of a memory leak. [PR/740793: This issue has been resolved.]

Hardware

- The XRE200 External Routing Engine temperature monitors, which you can view using the **show chassis environment** command, might report temperatures that are twice as high as the actual temperature. This temperature-reporting error has no impact on XRE200 External Routing Engine behavior. The fans and the system receive the correct temperature internally, so unwanted fan speed changes or an XRE200 External Routing Engine shutdown cannot occur as a result of this misreported temperature. However, the incorrect reported temperatures generate alarms and alarm messages. [PR/734233: This issue has been resolved.]

- On EX3300 switches, power supply failure errors might occur. To circumvent this problem, a software workaround has been provided. The software reads the power supply bit multiple times before it declares the power supply module to be down. [PR/743115: This issue has been resolved.]

High Availability

- After multiple GRES operations, the Virtual Management Ethernet (vme) interface goes down and comes back up again after you restart Ethernet switching. [PR/719424: This issue has been resolved.]
- When NSB is enabled on a switch, if you issue the **show spanning-tree interface msti msti-id** command on the backup Routing Engine, no output is displayed. [PR/732676: This issue has been resolved.]
- After a GRES operation with an NSB, the MSTP port boundary status might be displayed incorrectly. [PR/737179: This issue has been resolved.]

Infrastructure

- The **allow-configuration-regexps** statement at the **[edit system login class]** hierarchy level does not work exactly the same way as the deprecated **allow-configuration** statement at the same hierarchy level. [PR/720013: This issue has been resolved.]
- If you include the **autoinstallation** configuration statement at the **[edit system]** hierarchy level, the switch interfaces might not work correctly. [PR/728344: This issue has been resolved.]
- When you delete the VLAN mapping for an aggregated Ethernet (ae) interface, the Ethernet switching process (eswd) might crash and display the error message **No vlan matches vlan tag 116 for interface ae5.0**. [PR/731731: This issue has been resolved.]
- The **unlink** option in the **request system software add package** command does not work on EX Series switches. [PR/739795: This issue has been resolved.]
- When you quickly insert and then remove a line card, the chassis manager process (chassism) might become unstable. [PR/740730: This issue has been resolved.]
- On EX Series switches and SRX Series services gateways, if you enable "Change password every time the user logs out" in the active directory, you are unable to change your password. [PR/740869: This issue has been resolved.]
- On EX8200 switches, a chassism core file might be created. [PR/745964: This issue has been resolved.]
- When there is a large amount of NetBIOS traffic on the network, the switch might exhibit high latency while pinging between VLANs. [PR/748707: This issue has been resolved.]
- On all EX Series switches except EX8200 switches, if you have configured several policer settings in the same filter, they might all be overwritten when you change one of the settings. As a workaround, delete the setting and then add it back again with the desired changes. [PR/750497: This issue has been resolved.]

Interfaces

- When you perform a switchover between two XRE200 External Routing Engines, an LACP flap might occur on the aggregated Ethernet (ae) interfaces in a LAG. [PR/705772: This issue has been resolved.]

J-Web Interface

- In the J-Web interface, if you click the EX8200-48T, EX8200-48F, or EX8200-8XS line card in the chassis view in the dashboard, the expanded line card might not load its interfaces and might not display the interface status for both the EX8208 and EX8216 switches. As a workaround, first click the EX8200-40XS line card in the same chassis view and then close that line card. Then, click the EX8200-48T, EX8200-48F, or EX8200-8XS line card to display the status of all interfaces. [PR/742448: This issue has been resolved.]
- If you used the CLI to create an RTG group whose members are not trunk ports, you cannot edit this group from the J-Web interface. As a workaround, edit the group from the CLI. [PR/745458: This issue has been resolved.]
- For EX Series switches, when you use the J-Web interface software upload package, the **unlink** option does not work. [PR/746546: This issue has been resolved.]
- When a switch has no routed interfaces, you cannot use the J-Web interface to add OSPF areas. As a workaround, use the CLI to add these areas. [PR/746624: This issue has been resolved.]
- In the J-Web interface on an EX8200 switch that is set in **virtual-chassis** mode, when you expand the number of uplink modules, line cards that have no uplink module report an error or map ports to nonexistent modules. This problem happens the first time that you configure capacity utilization values. [PR/750854: This issue has been resolved.]

Software Upgrade and Installation

- When you use NSSU to upgrade from Junos OS Release 11.3R5 to Junos OS Release 12.1, all traffic across a LAG might be dropped. [PR/733050: This issue has been resolved.]

Virtual Chassis

- In a setup in which two XRE200 External Routing Engines (one acting as the master, the other as the backup) are connected to a member of an EX8200 Virtual Chassis that has two Routing Engines (one acting as the master, the other as the backup), if you remove the master Routing Engine or if you reboot this Routing Engine (for example, using the **request system reboot member 0 re0** command when re0 is the master Routing Engine), interfaces on which LACP is configured might flap. This interface flapping does not occur if you remove or reboot the backup Routing Engine. [PR/718857: This issue has been resolved.]
- On an EX2200 switch used as an intermediary switch in an EX8200 Virtual Chassis setup, a link down is not detected after you reboot the partner XRE200 External Routing Engine. [PR/726501: This issue has been resolved.]

- On EX4200 switches and EX4200 Virtual Chassis, the event process (eventd) process might create a core file. [PR/737893: This issue has been resolved.]
- In EX8200 Virtual Chassis, the switch might incorrectly send untagged packets. As a result, some hosts in the VLAN experience connectivity issues. [PR/752021: This issue has been resolved.]

Issues Resolved in Release 12.1R3

The following issues have been resolved since Junos OS Release 12.1R2. The identifier following the description is the tracking number in our bug database.

Access Control and Port Authority

- If a VLAN change occurs quickly, the client might not be able to get an IP address. [PR/746479: This issue has been resolved.]
- If you enable 802.1X with MAC RADIUS authentication, that is, by including the **mac-radius** statement in the configuration, the authentication manager process (authd) might reach a memory limit when there are approximately 250 users. As a workaround, reset the authd process when it reaches 85 percent of its RLIMIT_DATA value (that is, 85 percent of 130 MB). To check the amount of memory being used by the authd process, use the **show system processes extensive** operational mode command. [PR/783363: This issue has been resolved.]
- DHCP snooping might not enable DHCP Inform ACK packets to pass to the client. [PR/787161: This issue has been resolved.]
- If you configure a static MAC bypass for 802.1X (dot1x) and you add a new host to the exclusion list, the MAC addresses of existing hosts that have already been successfully authenticated using static MAC bypass might move to an incorrect MAC address. [PR/787679: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- When you enable Q-in-Q tunneling and MLD snooping, no snooping database is present on the switch. [PR/693224: This issue has been resolved.]

High Availability

- On an XRE200 External Routing Engine, when you perform an NSSU operation that includes the **reboot** option, the physical link might flap, which causes traffic loss and protocol flapping. [PR/718472: This issue has been resolved.]

Infrastructure

- If you enable gratuitous ARP by including the **gratuitous-arp-reply**, **no-gratuitous-arp-reply**, or **no-gratuitous-arp-request** statement in the configuration, the switch might process gratuitous ARP packets incorrectly. [PR/518948: This issue has been resolved.]
- In some cases, broadcast traffic that is received on the management port (me0) is broadcast to other subnets on the switch. [PR/705584: This issue has been resolved.]

- In previous releases, typing the alt-break sequence on the console put the console interface in debugger mode (the db> prompt). You can now configure the **system no-debugger-on-alt-break** statement to disable the alt-break sequence on the serial console. [PR/717491: This issue has been resolved.]
- NTP-related **show** commands such as **show ntp status** might display incorrect output. [PR/722528: This issue has been resolved.]
- When multicast traffic is transiting the switch, a kernel panic may occur on a new master Routing Engine, and the kernel may display the string **rn_clone_unwire parent unreferenced** during an NSSU operation or after multiple GRES operations. [PR/734295: This issue has been resolved.]
- On EX4200 switches, a Packet Forwarding Engine process (pfem) core file might be created while the switch is running the PFE internal support script and saving the output to a file. [PR/749974: This issue has been resolved.]
- You might see the following message in log files: **Kernel/ (COMPOSITE NEXT HOP) failed, err 6 (No Memory)**. [PR/751985: This issue has been resolved.]
- EX4500 Series switches and EX8200-40XS line cards do not forward IP UDP packets when their destination port is 0x013f (PTP) or when the fragmented packet has the value 0x013f at the same offset (0x2c). [PR/775329: This issue has been resolved.]
- When EX Series switches receive packets across a GRE tunnel, they might not generate and send ARP packets to the device at the other end of the tunnel. [PR/782323: This issue has been resolved.]
- After you remove an IPv6 interface configuration and then perform a rollback operation, the IPv4 label might change to explicit null. [PR/786537: This issue has been resolved.]
- If you configure IPv6 and VRRP, the IPv6 VRRP MAC address might be used incorrectly as the source MAC address when traffic is routed across VLANs. [PR/791586: This issue has been resolved.]

Interfaces

- When VRRP is running between two EX8200 Series switches on a VLAN, after a master switchover, both switches might act as master. [PR/752868: This issue has been resolved.]
- On EX8200 switches, the **master-only** configuration option does not work on the management interface (me0). [PR/753765]

J-Web Interface

- In some help files, the copyright date is set to 2011 instead of 2012. [PR/735607: This issue has been resolved.]
- The J-Web interface is vulnerable to HTML cross-site-scripting attacks, also called XST, or cross-site tracing. [PR/752398: This issue has been resolved.]

- When you configure the **no-tcp-reset** statement, the J-Web interface might be slow or unresponsive. [PR/754175: This issue has been resolved.]
- In the J-Web interface, the Help page for the Install package in the Software Maintenance page (Maintain > Software) might not appear. [PR/786654: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- When you are using IS-IS for forwarding only IPv6 traffic and IPv4 routing is not configured, if you perform an SNMP get/walk on an IS-IS routing database table, the RPD process might crash and restart, possibly causing a momentary traffic drop. [PR/753936: This issue has been resolved.]

Multicast Protocols

- If multicast clients are present across VRF instances, multicast traffic might not be received when you deactivate and then reactivate one of the VRF instances. [PR/769963: This issue has been resolved.]
- When an EX Series switch is routing multicast traffic, that traffic might not exit from the multicast router port in the source VLAN. [PR/773787: This issue has been resolved.]

Network Management and RMON

- After you upgrade to Junos OS Release 11.4R3, 11.4R4, or 12.1R2, EX Series switches might stop responding to SNMP ifIndex list queries. As a workaround, restart the switch. If restarting the switch is not an option, restart the shared-memory daemon (shm-rtssdbd). [PR/782231: This issue has been resolved.]

Virtual Chassis

- On EX3300 switches, when a Virtual Chassis is formed, the Virtual Chassis backup member's console CLI does not automatically redirect to the Virtual Chassis master's console CLI. As a workaround, manually log out from the Virtual Chassis backup member. [PR/744241: This issue has been resolved.]
- On EX8200 switch line cards, a Packet Forwarding Engine process (pfem) core file might be created as the result of a memory segmentation fault. [PR/757108: This issue has been resolved.]

Issues Resolved in Release 12.1R4

The following issues have been resolved since Junos OS Release 12.1R3. The identifier following the description is the tracking number in our bug database.

Access Control and Port Authority

- When access configuration is not required and the guest VLAN feature is configured, supplicants might not be authenticated using the guest VLAN, and they might remain in the *connecting* state. [PR/783606: This issue has been resolved.]
- When you configure DHCP snooping, DHCP Inform ACK packets might not pass to the client. [PR/787161: This issue has been resolved.]

Converged Networks (LAN and SAN)

- On EX4500 switches, the DCBX protocol does not work. [PR/795835: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- When you add a new VRF instance, existing firewall filters might not be applied to the new VRF instance. [PR/786662: This issue has been resolved.]
- Ethernet ring protection switching (ERPS; G.8032) does not block PVST BPDUs. [PR/793891: This issue has been resolved.]

High Availability (HA) and Resiliency

- After you perform an NSSU, you might notice a traffic outage of 150 seconds while the line cards are restarting. [PR/800460: This issue has been resolved.]

Infrastructure

- The **wildcard range unprotect** configuration statement might not be synchronized with the backup Routing Engine. [PR/735221: This issue has been resolved.]
- After you successfully install Junos OS, if you uninstall AI scripts, an mgd core file might be created. [PR/740554: This issue has been resolved.]
- On EX3300 switches, if you configure more than 20 BGPv6 neighbor sessions, the CLI might display the db> prompt. [PR/753261: This issue has been resolved.]
- The Junos OS kernel might crash because of a timing issue in the ttymodem() internal I/O processing routine. The crash can be triggered by simple remote access (such as Telnet or SSH) to the device. [PR/755448: This issue has been resolved.]
- When many packets are queued to have their next hop resolved, some packets might become corrupted. [PR/790201: This issue has been resolved.]
- After you upgrade Junos OS, a ppmd core file might be created, and protocols that use ppmd might not work correctly. [PR/802315: This issue has been resolved.]

Interfaces

- When you issue the **show vrrp brief** command, a VRRP process (vrrpd) core file might be created. [PR/782227: This issue has been resolved.]
- On EX Series switches, if you have configured a LAG bundle with link protection, an interface on the backup member might drop ingress traffic. [PR/796348: This issue has been resolved.]
- If you apply a policer to an interface, the policer might not work, and messages similar to the following are logged: **dfw_bind_policer_template_to_filter:205 Binding policer fails**. [PR/802489: This issue has been resolved.]

J-Web Interface

- If you issue the **set protocols rstp interface logical-interface-name edge** configuration command from the CLI, the J-Web interface might show that the configuration in the

Configuration detail for Desktop and Phone window is not applicable for the port profile. However, no functionality for the Desktop and Phone port profile is affected. [PR/791323: This issue has been resolved.]

- In the J-Web interface, if you enable a spanning-tree protocol (STP, RSTP, or MSTP) and then exclude some ports from the spanning tree, you might not be able to include these ports as part of an RTG. [PR/791759: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- If you have configured PIM NSR, a core file might be created on an upstream router because of high churn in unicast routes or a continuous clearing of PIM join-distribution in the downstream router. To prevent this possibility, disable NSR for PIM. [PR/707900: This issue has been resolved.]
- On EX3300 switches, when you are configuring BGP authentication, after you have configured the authentication key, BGP peering is never established. [PR/803929: This issue has been resolved.]

Management and RMON

- The incorrect ifType might be displayed for counters on physical interfaces. [PR/784620: This issue has been resolved.]
- After a Routing Engine switchover, LACP and MIB process (mib2d) core files might be created. [PR/790966: This issue has been resolved.]
- In logical systems, you cannot use snmpwalk for SNMP polling. As a workaround, configure the client at the **[edit snmp community logical-system]** hierarchy level. [PR/791859: This issue has been resolved.]

Multicast Protocols

- While multicast is resolving routes, the following SPF-related error might be displayed: **SPF:spf_change_sre(),383:jt_change() returned error-code (Not found:4)!** [PR/774675: This issue has been resolved.]
- On XRE200 External Routing Engines on which PIM is configured, an NSSU operation might fail when performed when an MSDP peer is not yet up. As a workaround, either disable NSR for PIM using the **set protocols pim nonstop-routing disable** configuration comment or ensure that MSDP has reached the Established state before starting an NSSU operation. [PR/799137: This issue has been resolved.]

Virtual Chassis

- On XRE200 External Routing Engines, a chassis core file might be created. [PR/791959: This issue has been resolved.]
- On XRE200 External Routing Engines on which DHCP snooping and dynamic ARP inspection are enabled, when packets are transmitting out a different line card type from the ingress interface, an SFID core file might be created. [PR/794293: This issue has been resolved.]

Issues Resolved in Release 12.1R5

The following issues have been resolved since Junos OS Release 12.1R4. The identifier following the description is the tracking number in our bug database.

Access Control and Port Authority

- On EX6200 switches, LLDP stops working if you execute the **set ethernet-switching-options voip interface access-ports vlan** command. [PR/829898: This issue has been resolved.]

Infrastructure

- After an EX8200 switch has been up for several days, the line cards might reach 100 percent CPU usage and then stay at 100 percent. [PR/752454: This issue has been resolved.]
- On EX8200 switches, when you issue the **request system reboot other-routing-engine** command, a timeout error might be displayed before the Routing Engine initiates its reboot operation. [PR/795884: This issue has been resolved.]
- On EX Series switches that have PoE capability, chassisd (the chassis daemon) might crash if you run SNMP requests (for example, SNMP get, get-next, or walk requests) on pethMainPse objects. This happens because the system tries to free memory that is already freed. As a workaround, do not run SNMP requests on pethMainPse objects. [PR/817311: This issue has been resolved.]
- On EX4200 switches, high CPU usage might be due to console cable noise. [PR/818157: This issue has been resolved.]
- On EX8200 Virtual Chassis, when both dscp and ieee-802.1 rewrite rules are applied on an RVI, deleting filters on the RVI and binding them again on the same RVI might create a pfem core file. [PR/828661: This issue has been resolved.]
- When an uplink module in a switch is operating in 1-gigabit mode, a chassism core file might be created if you remove an SFP transceiver from the module. As the chassism process restarts, all traffic passing through the interface is dropped. This problem happens with both copper and fiber SFPs. [PR/828935: This issue has been resolved.]
- Traffic leaks might occur for unknown unicast and broadcast traffic from multiple VLANs when a MAC-RADIUS-assigned VLAN is set on a switch interface through a server-initiated attribute change. If the 802.1X interface has VLAN 100 assigned and the RADIUS server sends a different VLAN attribute (for example, 200 rather than 100), after the interface is assigned in VLAN 200, the RADIUS server also sends egress unknown unicast and broadcast traffic that belongs to VLAN 100. [PR/829436: This issue has been resolved.]
- An SNMP poll might not return clear information for some FRUs, such as fans and power supplies. The FRU description might not indicate which switch contains the FRU. [PR/837322: This issue has been resolved.]
- If you reboot the switch with an RVI disabled, then even if you reenables the RVI, the RVI traffic is not routed in the Packet Forwarding Engine; the traffic is trapped in the

CPU and is policed by the rate limit in the Packet Forwarding Engine. [PR/838581: This issue has been resolved.]

Interfaces

- On EX8200 Virtual Chassis, when you swap the members of a LAG, a vmcore or ksyncd core file might be created on the backup Routing Engine. [PR/793778: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- On an EX4200 switch configured for VLAN translation, Microsoft Windows NetBIOS traffic might not be translated. [PR/791131: This issue has been resolved.]

Management and RMON

- Incorrect information might be displayed for output ports. [PR/784623: This issue has been resolved.]
- EX Series switches might drop sFlow monitoring technology packets if the packet size exceeds 1500 bytes. [PR/813879: This issue has been resolved.]
- In EX3300 Virtual Chassis, if you perform an SNMP poll of jnxOperatingState for fan operation, the information for the last two members in the Virtual Chassis is incorrect. [PR/813881: This issue has been resolved.]
- On EX8200 switches, sFlow monitoring technology packets were being generated with an incorrect source MAC address of 20:0b:ca:fe:5f:10. This issue has been resolved. EX8200 switches now use the outbound port's MAC address as the source MAC address for the sFlow monitoring technology traffic. [PR/815366: This issue has been resolved.]

Issues Resolved in Release 12.1R6

The following issues have been resolved since Junos OS Release 12.1R5. The identifier following the description is the tracking number in our bug database.

Access Control and Port Authority

- When you use the dynamic firewall filter allocation feature, if the filter ID string is very large and the reauthentication interval is very short, a memory leak might occur with the dot1x process. [PR/837183: This issue has been resolved.]
- On EX Series switches, the LLDP-MED media endpoint class is shown as invalid. This problem is just a display issue—there is no functional impact. [PR/840915: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- If an EX Series switch has an RTG link, a MAC Refresh message might be sent on a new active link of the RTG when RTG failover occurs. The switch sends the RTG MAC Refresh message with a VLAN tag even though RTGs are configured on access ports. [PR/853911: This issue has been resolved.]

Hardware

- On EX3200, EX4200, and EX8200 switches, the receiver signal average optical power is shown as 0.0000 in output for the **show interfaces diagnostics optics** command. The problem has been observed for SFP-SX and SFP-LX10 transceivers. [PR/854726: This issue has been resolved.]

Infrastructure

- If you use EZSetup to configure a root password that contains a comma (,), the characters after the comma are not checked during authentication, so it is possible to log in to the switch with several different passwords. As a workaround, configure the root password from the CLI. [PR/738310]
- For all EX Series switches except the EX8200 switch, hash index collisions were causing problems with the learning of MAC addresses in the forwarding database (FDB). You can now increase the maximum number of searchable hash indexes in increments of 4, from 4 to a maximum of 32 entries, using the **set ethernet-switching-options max-lookup-length** CLI command. [PR/842439: This issue has been resolved.]
- For EX4500 switches, queue counters are not updated for child members when the **monitor interface ae#** command is running. [PR/846059: This issue has been resolved.]
- In a mixed EX4200 and EX4500 Virtual Chassis, link aggregation might create a pfem core file in some member switches. [PR/846498: This issue has been resolved.]
- When you boot up an EX2200 or EX3300 switch with Junos OS Release 12.2R1 or later, the message **dog: ERROR - reset of uninitialized watchdog** appears. The message appears even if you reboot the switch by using the proper reboot procedure. The error does not cause a system reset; thus, you can ignore this message. [PR/847469: This issue has been resolved.]
- On EX Series switches, EXP CoS classification does not occur if EXP CoS classifiers are deleted and then added. [PR/848273: This issue has been resolved.]
- On EX4200 standalone switches and EX4200 Virtual Chassis, an alarm **/var partition is full** might be logged in the system log file even though the **/var** partition is not full and **CHASSISD_RE_CONSOLE_ME_STORM** log messages might also be logged in the system log file. [PR/866863: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- On EX Series switches, under rare conditions, an interface might not learn MAC addresses when it is authenticated by the 802.1X protocol. [PR/837970: This issue has been resolved.]
- L2TP for CDP and VTP does not work properly. [PR/842852]

- On EX Series switches, the Q-BRIDGE-MIB OID 1.3.6.1.2.1.17.7 reports the VLAN internal index instead of the VLAN ID. [PR/850299: This issue has been resolved.]

Management and RMON

- On EX Series switches, a configured OAM threshold value might get reset when the chassis is rebooted. [PR/829649]
- On EX4200 and EX4500 switches, adaptive sampling is triggered on interfaces configured for sFlow monitoring technology even though the sampling rate is less than 300 pps. [PR/840858: This issue has been resolved.]
- The SNMP query or walk on ipNetToMediaPhysAddress does not match **show arp** command output. [PR/850051: This issue has been resolved.]

Virtual Chassis

- On EX Series switches, if you configure a physical interface's MTU with a large value and you do not reconfigure the family inet MTU, OSPF packets might be dropped when they reach the internal logical interface if the packet size exceeds 1900 bytes. All communications traffic between Routing Engines and between FPCs passes through the internal logical interface. The OSPF neighbor does not receive the OSPF transmissions and ends the OSPF session. The switch displays the error message **bmeb_rx failed**. [PR/843583: This issue has been resolved.]

Issues Resolved in Release 12.1R7

The following issues have been resolved since Junos OS Release 12.1R6. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- On EX Series switches, DHCP snooping binding does not renew the lease time when IPv6 is configured on the client VLAN. When DHCP snooping is configured with ARP inspection and when a client renews the lease, the switch does not update the DHCP snooping table with the new lease time. The lease eventually times out of the DHCP snooping table, and the client still has a valid lease. The client's ARP request eventually times out of the switch, and the client loses connectivity because ARP inspection blocks the transmission because the client has no entry in the DHCP snooping table. [PR/864078: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- The **show spanning-tree interface vlan-id detail** command displays all the interfaces instead of displaying only those interfaces associated with the VLAN ID. [PR/853632: This issue has been resolved.]
- When a topology change is detected on an RSTP-enabled port, a BPDU packet is sent immediately with the topology change flag in all other ports. But for an MSTP-enabled port, a delay of a few seconds occurs before the BPDU packet is sent with the topology change flag. [PR/860748: This issue has been resolved.]

High Availability

- On EX8200 Virtual Chassis, an NSSU might fail. [PR/871288: This issue has been resolved.]
- NSSU on an EX8200 Virtual Chassis is not recommended for Junos OS Release 12.2R5. [PR/893440: This issue has been resolved.]

Infrastructure

- Rate limiting for management traffic (namely, SSH and Telnet) arriving on network ports causes file transfer speeds to be slow. [PR/831545: This issue has been resolved.]
- On EX4500 switches, multicast packet fragments might be dropped. [PR/835855: This issue has been resolved.]
- The sFlow monitoring technology feature is not supported on EX2200, EX2200-C, and EX3300 switches. [PR/872292: This issue has been resolved.]
- On EX3200 and EX4200 switches, high traffic on management Ethernet (me0) interfaces might affect switch control and management plane functions. [PR/876110: This issue has been resolved.]
- On EX8200 Virtual Chassis, NetBIOS traffic might be dropped when it crosses the Virtual Chassis port extension (VCPe) connections. The NetBIOS traffic is dropped because of a conflict on the Packet Forwarding Engine of the Virtual Chassis member with the VCPe ports. [PR/877503: This issue has been resolved.]
- On EX2200 switches, storm control does not limit traffic to the set value when that traffic enters through uplink ports; the traffic is limited at 10 times the set value. [PR/879798: This issue has been resolved.]
- On EX4200 switches, an aggregated Ethernet interface is not supported as a match condition in a firewall filter. [PR/886476: This issue has been resolved.]

Interfaces

- On EX Series switches, configuration of a static LACP system-ID is not supported. [PR/889318: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- If an invalid PIM-SSM multicast group is configured on the routing device, then when you issue the **commit** or **commit check** command, a routing protocol daemon (rpd) core file is created. There is no traffic impact because the main rpd process spawns another rpd process to parse the corresponding configuration changes, and the new rpd process crashes and creates a core file. When this problem occurs, you might see the following messages:

```
user@router# commit check
error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped
core(0x86)

error: configuration check-out failed

user@router# commit
```

error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86)

error: configuration check-out failed

[PR/856925: This issue has been resolved.]

Management and RMON

- On EX Series switches, when the ARP table is cleared from the CLI, the SNMP MIB ipNetToMediaPhysAddress might have more entries than the ARP table does. [PR/853536: This issue has been resolved.]

Software Upgrade and Installation

- On an EX2200-24T-DC-4G switch model, autoinstallation is not activated during initial installation because this model is missing a configuration file. [PR/873689: This issue has been resolved.]

Virtual Chassis

- On an EX8200 Virtual Chassis, when a Virtual Chassis Port (VCP) is physically down and GRES is performed on the LCC Routing Engine, a Virtual Chassis cannot be formed because the VCP is down, and traffic between the chassis may be affected. However, the VCP interface appears on the CLI. This problem may be detected during NSSU. [PR/730336: This issue has been resolved.]
- On an EX8200 Virtual Chassis, the **request system scripts add** command does not install the AI-Scripts bundle package on all nodes of the chassis. [PR/832975: This issue has been resolved.]
- On EX4200 Virtual Chassis, if the **mac-persistence-timer** value is set to 0, the system MAC base address is changed during a master switchover by the **request chassis routing-engine master switch** command. We recommend that you configure a value of 1 or more than 1 for **mac-persistence-timer** to avoid this issue. [PR/858330]

Related Documentation

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)
- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches

- [Changes to Junos OS for EX Series Switches Documentation on page 56](#)
- [Errata on page 56](#)

Changes to Junos OS for EX Series Switches Documentation

The following changes have been made to the documentation for Junos OS Release 12.1 for EX Series switches since it was published:

- The [EX Series Switch Software Features Overview](#) now correctly indicates that the EX4200 switches and EX8200 switches support virtual router-aware DHCP (VR-aware DHCP).

Errata

This section lists outstanding issues with the published documentation for Junos OS Release 12.1 for EX Series switches.

- The uplink failure detection feature is supported on EX3300 switches and on EX3300 Virtual Chassis starting in Junos OS Release 12.1R2.

The [EX Series Switch Software Features Overview](#) and [EX Series Virtual Chassis Software Features Overview](#) documents previously stated that initial support for the uplink failure detection feature started in Junos OS Release 12.2R1. These documents have been updated.

- The EX Series Redundant Power System (RPS) documentation incorrectly states that only one RPS can be connected to an EX3300 Virtual Chassis. However, multiple RPSs can be used to back up an EX3300 Virtual Chassis.
- The Junos OS documentation for the **route configuration** statement at the **[edit interfaces name unit number family inet address address vrrp-group group-id track]** hierarchy level shows that the statement is supported on EX Series switches as of Junos OS Release 12.1. The statement is supported in releases earlier than Release 12.1.
- The following changes have been made to the system snapshot functionality, which you configure using the **request system snapshot** command:
 - By default, a snapshot backs up both partitions (**/** and **/altroot**) to the media (internal or external) that the device did not boot from.
 - The following partitions are backed up by a snapshot: **/**, **/altroot**, **/config**, **/var**, and **/var/tmp**.
 - You do not need to specify a media slice number for the location of a snapshot.
 - You can specify the alternate slice on the media the device booted from as the snapshot location by using the **slice alternate** option.
 - When you create a snapshot on the media that the switch did not boot from, you must use the **partition** option to partition the destination media.
 - The **show system snapshot** command displays information for the **/** and **/altroot** partitions.

[This issue was being tracked by PR/599708.]

- In the J-Web interface, you cannot configure interface ranges and interface groups.
[This issue was being tracked by PR/600559.]

- The documentation for the **request system zeroize** command indicates that it was supported on EX Series switches as of Junos OS Release 11.2. Support for the **request system zeroize** command on the switches began in Release 11.1.
- In the topic *Example: Setting Up Bridging with Multiple VLANs for EX Series Switches*, one of the interfaces in the *support* VLAN is incorrectly shown as ge-0/0/0.24. The correct interface name is ge-0/0/24.0.
- You can configure Ethernet OAM link fault management (LFM) on aggregated interfaces.
- The documentation for the **vlan**s configuration statement incorrectly states the required privilege levels as routing and routing-control. The correct privilege levels for this statement are system and system-control.
- The documentation for the **request system software add** command incorrectly states that the **validate** option is supported on EX Series switches. This option is not supported on any EX Series switches. [This issue is being tracked by PR/821244.]
- The documentation for the **request system software validate** command incorrectly states that this command is supported on EX Series switches. This command is not supported on any EX Series switches. [This issue is being tracked by PR/803185.]

**Related
Documentation**

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)
- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches on page 57](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.1 for EX Series Switches

This section discusses the following topics:

- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 57](#)
- [Upgrading to Release 12.1R2 or Later Releases, with Existing VSTP Configurations on page 58](#)
- [Upgrading from Junos OS Release 10.4R3 or Later on page 58](#)
- [Upgrading from Junos OS Release 10.4R2 or Earlier on page 60](#)
- [Upgrading EX Series Switches Using NSSU on page 60](#)

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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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 <http://www.juniper.net/support/eol/junos.html>.

Upgrading to Release 12.1R2 or Later Releases, with Existing VSTP Configurations

If you are upgrading to Junos OS Release 12.1R2 or later releases from Release 12.1R1 or earlier releases, ensure that any VSTP configurations on the switch meet the following guidelines. If the VSTP configurations do not meet these guidelines and you run the upgrade, the upgrade fails and you have to connect the console, change the invalid VSTP configurations, and commit the changed configurations through the console. Guidelines for VSTP configurations are:

- If you have specified physical interfaces for VSTP-configured VLANs, ensure that those interfaces are members of the VLANs specified in the VSTP configuration. If the VSTP configuration specifies `vlan all`, then the interfaces configured under `vstp vlan all` must be members of all VLANs.
- If the interfaces are not members of the VLANs in the VSTP configurations but are already added to the VSTP configurations, remove them from those configurations, add them to the VLANs, and then add them back to the VSTP configurations.

This issue is being tracked by PR/736488 in our bug database.

Upgrading from Junos OS Release 10.4R3 or Later

This section contains the procedure for upgrading from Junos OS Release 10.4R3 or later to Junos OS Release 12.1. You can use this procedure to upgrade Junos OS on a standalone EX Series switch with a single Routing Engine and to upgrade all members of a Virtual Chassis or a single member of a Virtual Chassis.

To upgrade Junos OS on a EX6200 or EX8200 switch with dual Routing Engines, see [Installing Software on an EX Series Switch with Redundant Routing Engines \(CLI Procedure\)](#).

On switches with dual Routing Engines or on Virtual Chassis, you might also be able to use NSSU to upgrade Junos OS. See [“Upgrading EX Series Switches Using NSSU” on page 60](#) for more information.

To upgrade Junos OS on a switch with a single Routing Engine or on a Virtual Chassis:

1. Download the software package as described in [Downloading Software Packages from Juniper Networks](#).
2. (Optional) Back up the current software configuration to a second storage option. See the [Junos OS Installation and Upgrade Guide](#) for instructions.
3. (Optional) Copy the software package to the switch. We recommend that you use FTP to copy the file to the `/var/tmp` directory.

This step is optional because you can also upgrade Junos OS using a software image that is stored at a remote location.

4. Install the new software package on the switch:

```
user@switch> request system software add package
```

Replace *package* with one of the following paths:

- `/var/tmp/package.tgz`—For a software package in a local directory on the switch
- `ftp://hostname/pathname/package.tgz` or `http://hostname/pathname/package.tgz`—For a software package on a remote server

package.tgz is the name of the package; for example, `jinstall-ex-4200-11.4R1.8-domestic-signed.tgz`.

To install software packages on all switches in a mixed EX4200 and EX4500 Virtual Chassis, use the `set` option to specify both the EX4200 package and the EX4500 package:

```
user@switch> request system software add set [package package]
```

To install the software package on only one member of a Virtual Chassis, include the `member` option:

```
user@switch> request system software add package member member-id
```

Other members of the Virtual Chassis are not affected. To install the software on all members of the Virtual Chassis, do not include the `member` option.



NOTE: To abort the installation, do not reboot your device. Instead, finish the installation and then issue the `request system software delete package.tgz` command, where *package.tgz* is the name of the package; for example, `jinstall-ex-8200-11.4R1.8-domestic-signed.tgz`. This is the last chance to stop the installation.

5. Reboot the switch to start the new software:

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

To reboot only a single member in a Virtual Chassis, include the `member` option:

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

6. After the reboot has completed, log in and verify that the new version of the software is properly installed:

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

- Once you have verified that the new Junos OS version is working properly, copy the version to the alternate slice to ensure that if the system automatically boots from the backup partition, it uses the same Junos OS version:

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

To update the alternate root partitions on all members of a Virtual Chassis, include the **all-members** option:

```
user@switch> request system snapshot slice alternate all-members
```

Upgrading from Junos OS Release 10.4R2 or Earlier

To upgrade to Junos OS Release 12.1 from Release 10.4R2 or earlier, first upgrade to Junos OS Release 11.4 by following the instructions in the Junos OS Release 11.4 release notes. See *Upgrading from Junos OS Release 10.4R2 or Earlier* or *Upgrading from Junos OS Release 10.4R3 or Later* in the [Junos OS 11.4 Release Notes](#).

Upgrading EX Series Switches Using NSSU

You can use NSSU to upgrade Junos OS releases on EX8200 standalone switches and EX4200, EX4500, and EX8200 Virtual Chassis. For instructions on how to perform an upgrade using NSSU, see:

- [Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, EX4550 Virtual Chassis, or Mixed Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\)](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\)](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\)](#)

Table 1 on page 60 details the Junos OS releases on which NSSU is supported.

Table 1: Platform and Junos OS Upgrade Support for NSSU

Switch Platform	Upgrade from Junos OS Release x.x	Upgrade to Junos OS Release 12.1R1	Upgrade to Junos OS Release 12.1R2 or Later
EX4200 Virtual Chassis, EX4500 Virtual Chassis, and mixed EX4200 and EX4500 Virtual Chassis	Releases earlier than 12.1R1	Not supported	Not supported
	12.1R1	—	Supported
EX8200 standalone switch	Please contact JTAC before you use NSSU to upgrade EX8200 standalone switches to Junos OS Release 12.1R1 or later.		
EX8200 Virtual Chassis	Please contact JTAC before you use NSSU to upgrade EX8200 Virtual Chassis to Junos OS Release 12.1R1 or later.		

On an EX8200 Virtual Chassis, an NSSU operation can be performed only if you have configured the XRE200 External Routing Engine member ID to be 8 or 9.



NOTE: If you are using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 10.4 or Release 11.1 and sFlow technology is enabled, disable sFlow technology before you perform the upgrade using NSSU. After the upgrade is complete, you can reenables sFlow technology. If you do not disable sFlow technology before you perform the upgrade with NSSU, sFlow technology does not work properly. This issue does not affect upgrades from Junos OS Release 11.2 or later.



NOTE: If you are using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 11.1 and NetBIOS snooping is enabled, disable NetBIOS snooping before you perform the upgrade using NSSU. After the upgrade is complete, you can reenables NetBIOS snooping. If you do not disable NetBIOS snooping before you perform the upgrade with NSSU, NetBIOS snooping does not work properly. This issue does not affect upgrades from Junos OS Release 11.2 or later.

**Related
Documentation**

- [New Features in Junos OS Release 12.1 for EX Series Switches on page 8](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for EX Series Switches on page 18](#)
- [Limitations in Junos OS Release 12.1 for EX Series Switches on page 20](#)
- [Outstanding Issues in Junos OS Release 12.1 for EX Series Switches on page 26](#)
- [Resolved Issues in Junos OS Release 12.1 for EX Series Switches on page 32](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.1 for EX Series Switches on page 55](#)

Junos OS Release Notes for Branch SRX Series Services Gateways and J Series Services Routers

Powered by Junos OS, Juniper Networks SRX Series Services Gateways provide robust networking and security services. SRX Series Services Gateways range from lower-end branch devices designed to secure small distributed enterprise locations to high-end devices designed to secure enterprise infrastructure, data centers, and server farms. The branch SRX Series Services Gateways include the SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, and SRX650 devices. The high-end SRX Series Services Gateways include the SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

Juniper Networks J Series Services Routers running Junos OS provide stable, reliable, and efficient IP routing, WAN and LAN connectivity, and management services for small to medium-sized enterprise networks. These routers also provide network security features, including a stateful firewall with access control policies and screens to protect against attacks and intrusions, and IPsec VPNs. The J Series Services Routers include the J2320, J2350, J4350, and J6350 devices.

Junos OS 12.1 software is not currently available for J Series Services Routers. If you would like to receive notification when Junos OS 12.1 is available for J Series Services Routers, please subscribe to the Juniper Networks technical bulletin service. If required, please contact your Juniper Networks customer service representative for more information.



NOTE:

- [Release Notes Updates on page 62](#)
- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)
- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)
- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Release Notes Updates

The technical documentation for Junos OS Release 12.1R1 for SRX Series and J Series devices has been rereleased in the Juniper Networks standard format. This new tabbed

webpage format simplifies browsing by enabling you to navigate to content based on the job you need to do: Overview, Configuration, Administration, or Troubleshooting.

- Access to related information and to PDF versions of each document are available from each tabbed page.
- For Junos OS Release 12.1R1, "legacy" PDFs are also provided for those customers who have already bookmarked the original Junos OS Release 12.1R1 documentation.
- The Junos OS Release Notes for SRX Series and J Series devices will continue to be available in both HTML and PDF format.
- See [Junos OS for SRX Series: Release 12.1](#)

New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

The following features have been added to Junos OS Release 12.1. Following the description is the title of the manual or manuals to consult for further information.



NOTE: For the latest updates about support and issues on Junos Pulse, see the [Junos Pulse Release Notes](#).

- [Release 12.1R4 New Features on page 63](#)
- [Release 12.1R3 New Features on page 65](#)
- [Release 12.1R2 Software Features on page 67](#)
- [Release 12.1 R2 Hardware Features on page 72](#)
- [Release 12.1 R1 Software Features on page 74](#)
- [Hardware Features—SRX550 Services Gateways on page 81](#)
- [Software Features—SRX550 Services Gateways on page 83](#)

Release 12.1R4 New Features

Hardware Features---SRX240 Services Gateways

[Table 2 on page 64](#) lists the hardware features supported on various models of the SRX240 Services Gateway

Table 2: SRX240 Services Gateway Hardware Features

Features	SRX240 Services Gateway with AC Power Supply and No PoE Support	SRX240 Services Gateway with AC Power Supply and PoE Support	SRX240 Services Gateway with DC Power Supply
DDR memory	- For SRX240B: 512 MB - For SRX240B2: 1 GB - For SRX240H: 1 GB - For SRX240H2: 2 GB - For SRX240H-TAA: 1 GB - For SRX240H2-TAA: 2 GB	- For SRX240H-POE: 1 GB - For SRX240H2-POE: 2 GB - For SRX240H-POE-TAA: 1 GB - For SRX240H2-POE-TAA: 2 GB	- For SRX240H-DC: 1 GB - For SRX240H2-DC: 2 GB
PoE support	No	Yes	No
Power supply rating	150 watts	360 watts	190 watts
Input voltage	100 to 240 VAC	100 to 240 VAC	–48 VDC Operating range: –40.5 V to –72 V
Gigabit Ethernet ports	16	16	16
Console ports	1	1	1
Universal Serial Bus (USB) ports	2	2	2
Mini-PIM slots	4	4	4
LEDs	Status, Alarm, HA, Power, Mini-PIMs, and Port (TX/RX/Link and PoE)	Status, Alarm, HA, Power, Mini-PIMs, and Port (TX/RX/Link and PoE)	Status, Alarm, HA, Power, Mini-PIMs, Port (TX/RX/Link and PoE), DC power feed LEDs.

NOTE: The PoE LED is enabled only on the PoE variant of the SRX240 Services Gateway.

Table 2: SRX240 Services Gateway Hardware Features (*continued*)

Features	SRX240 Services Gateway with AC Power Supply and No PoE Support	SRX240 Services Gateway with AC Power Supply and PoE Support	SRX240 Services Gateway with DC Power Supply
Internal flash memory	- For SRX240B: 1 GB - For SRX240B2: 2 GB - For SRX240H: 1 GB - For SRX240H2: 2 GB - For SRX240H-TAA: 1 GB - For SRX240H2-TAA: 2 GB	- For SRX240H-POE: 1 GB - For SRX240H2-POE: 2 GB - For SRX240H-POE-TAA: 1 GB - For SRX240H2-POE-TAA: 2 GB	- For SRX240H-DC: 1 GB - For SRX240H2-DC: 2 GB
Fans	6	6	6
Air filters	One (Separately orderable)	None	One
NEBS-compliance	No	No	Yes



NOTE: An air filter is not shipped with the SRX240 Services Gateway with AC power supply models. To meet NEBS requirements, you must order the air filter separately and install it. Contact your Juniper Networks customer service representative for more information.

Release 12.1R3 New Features

Network Address Translation (NAT)

- Network Address Translation support for port mapping**—This feature is supported on all branch SRX Series and J Series devices.

Network Address Translation (NAT) now provides **destination-port low to high** and **mapped-port low to high** statements to allow static NAT to map ports as follows:

- To map multiple IP addresses to the same IP addresses on a specified range of ports
- To map a specific IP address and port to a different IP address and port

Examples of configuring static NAT with port mapping:

```
set security nat static rule-set rs from zone trust rule r1 match destination-address 1.1.1.1/32
set security nat static rule-set rs from zone trust rule r1 destination-port 100 to 200
set security nat static rule-set rs from zone trust rule r1 then static-nat prefix 10.1.1.1/32
set security nat static rule-set rs from zone trust rule r1 then static-nat prefix mapped-port 300 to 400
```

```
set security nat static rule-set rs from zone trust rule r2 match destination-address 1.1.1/32
set security nat static rule-set rs from zone trust rule r2 destination-port 300 to 400
set security nat static rule-set rs from zone trust rule r2 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs from zone trust rule r2 then static-nat prefix mapped-port 300
to 400
set security nat static rule-set rs from zone trust rule r3 match destination-address 1.1.1.3/32
set security nat static rule-set rs from zone trust rule r3 destination-port 300
set security nat static rule-set rs from zone trust rule r3 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs from zone trust rule r3 then static-nat prefix mapped-port 200
```

Dynamic Host Configuration Protocol (DHCP)

- **Dynamic Host Configuration Protocol (DHCP)** — This feature is supported on branch SRX Series devices.

The DHCP feature has been enhanced to include routing instance support. To configure a DHCP server, DHCP relay agent, or DHCP client in a routing instance, use the **[edit routing-instances]** hierarchy.

The following considerations apply when you enable a DHCP server, DHCP relay agent, or DHCP client in a routing instance:

- The DHCP server, DHCP relay agent, and DHCP client can be configured in one routing instance, but the functionality is mutually exclusive on one interface. If the DHCP client is enabled on one interface, the DHCP server or the DHCP relay agent cannot be enabled on that interface.
- The DHCP client, DHCP relay agent and DHCP server services act independently in their respective routing instance. The following features can function simultaneously on a device:
 - DHCP client and DHCP server
 - DHCP client and DHCP relay agent
- Multiple routing instances. Each instance can have a DHCP server, DHCP relay agent, or DHCP client, or each routing instance can have a DHCP client and DHCP server or a DHCP client and DHCP relay agent.



NOTE: There are two types of DHCP services available:

- DHCP services in a default routing instance
- DHCP services in a routing instance.

From Junos OS Release 12.1 R3 onwards, we recommend using DHCP services in a routing instance.

Junos OS CLI Reference, Junos OS Initial Configuration Guide for Security Devices

Release 12.1R2 Software Features

Interfaces

- **8-Port Serial GPIM**—The 8-Port synchronous serial GPIM is supported on SRX550 and SRX650 devices. This GPIM provides 8 ports that operate in sync mode and supports a line rate of 64 Mbps or 8 Mbps per port.

The 8-Port synchronous serial GPIM supports the following features:

- Operation modes (auto-select based on cable, no configuration required)
 - DTE (data terminal equipment)
 - DCE (data communication equipment)
- Clocking
 - Tx clock modes
 - DCE clock (only valid in DTE mode)
 - Baud clock (internally generated)
 - Loop clock (external)
 - Rx clock modes
 - Baud clock (internally generated)
 - Loop clock (external)
- Clock rates (baud rates) from 1.2 KHz to 8.0 MHz



NOTE: RS-232 serial interfaces might cause an error with a clock rate greater than 200 KHz.

- MTU—9192 bytes, default value is 1504 bytes
- HDLC features
 - Idle flag/fill (0x7e or all ones), default idle flag is (0x7e)
 - Counters—giants, runts, FCS error, abort error, align error
- Line encoding—NRZ and NRZI
- Invert data
- Line protocol—EIA530/EIA530A, X.21, RS-449, RS-232, V.35
- Data cables—Separate cable for each line protocol (both DTE/DCE mode)
- Error counters (conformance to ANSI specification)
- Alarms and defects

- Rx clock absent
- Tx clock absent
- DCD absent
- RTS/CTS absent
- DSR/DTR absent
- Data signal—Rx clock
- Control signals
 - To DTE: CTS, DCD, DSR
 - From DTE: DTR, RTS
- Serial autoresync
 - Configurable resync duration
 - Configurable resync interval
- Diagnostic features
 - Loopback modes—local, remote, and dce-local loopback
 - Ability to ignore control signals
- Layer 2 features
 - Encapsulation
 - PPP (PAP/CHAP)
 - Cisco HDLC
 - Frame Relay
 - MLPPP
 - MLFR
- SNMP features
 - SNMP information should be receivable at each port
 - IF-MIB - rfc2863a.mib
 - jnx-chassis.mib
- Anticounterfeit check

Interfaces Configuration Guide for Security Devices

- **802.3ah OAM LFM**—This feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650 devices.

The Ethernet interfaces on these SRX Series devices support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on

point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

The following OAM LFM features are supported:

- Discovery and link monitoring—The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The device performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote fault detection—Remote fault detection uses flags and events. Flags convey Link Fault (a loss of signal), Dying Gasp (an unrecoverable condition such as a power failure), and Critical Event (an unspecified vendor-specific critical event). You can specify the periodic OAM PDU sending interval for fault detection. SRX Series devices use the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.
- Remote loopback—Remote loopback mode ensures link quality between the device and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

Junos OS can place a remote data terminal equipment (DTE) into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

The following interfaces are supported:

1. Physical interface (fe/ge)
 - IFD encapsulations
 - ethernet-ccc
 - extended-vlan-ccc

- ethernet-tcc
 - extended-vlan-tcc
 - vlan-tagging mode
 - family
 - inet
 - mpls
 - ccc
 - tcc
 - iso
 - ethernet-switching
 - inet6
2. Aggregated Ethernet interface
- Static lag
 - LACP enabled lag
 - vlan-tagging mode
 - family
 - ethernet-switching
 - inet
 - mpls
 - iso
 - inet6
 - IFD encapsulations
 - ethernet-ccc
 - extended-vlan-ccc
 - vlan-ccc

Junos OS CLI Reference

- **Multi-VLAN over VDSL (PT Mode only)**—This feature is supported on SRX110, SRX210, SRX220, SRX240, and SRX550 devices.

In a Multi-VLAN over VDSL environment, a pt-interface is enabled to configure VLAN tagging. A single pt-interface contains different logical units that are tagged to different VLANs. There is a VLAN tag attached in any packet that is sent to a particular logical unit and any logical unit, that receives a packet will have a similar VLAN tag attached to it.

*Interfaces Configuration Guide for Security Devices***IP Monitoring**

- **IP Monitoring**—This feature is supported on SRX100, SRX210, SRX220, SRX240, and SRX650 devices.

Enhanced the IP Monitoring feature to support the following actions:

- Specifying the no-preempt option. If the no-preempt option is specified, the policy does not perform preemptive failback when it is in a failover state or when the RPM probe test recovers from a failure.
- Setting preferred metric values. If the preferred metric value is set, during failover, the route is injected with the set preferred metric value.
- Enabling and disabling interfaces.
 - **Interface-Enable** – On a physical or logical interface, when the interface-enable action is configured, the initial state of the interface is disable after startup, and it continues to remain in the disable state as long as the associated RPM probe is in the pass state. When the associated RPM probe fails, the configured physical and logical interfaces are enabled.
 - **Interface-Disable** – On a physical or logical interface, when the interface-disable action is configured, the interface state remains unchanged. When the associated RPM probe fails, the physical and logical interfaces are disabled.

Security Policies

- New **show services unified-access-control counters** command—This feature is supported on all branch SRX Series devices.

The **show services unified-access-control counters** command displays the number of sessions allowed, denied, and terminated by the Unified Access Control (UAC) service when invoked by a firewall policy that specifies the uac-policy action. Counts are reported for each action taken by UAC.

On branch SRX Series devices, UAC counts are accumulated by device only. There is no PIC specification on these devices. Sessions that were allowed, denied, or terminated by some other firewall policy action are not included in these statistics.

```
user@host> show services unified-access-control counters
Sessions allowed
Policy action: 0
```

Timeout action: 0
 Sessions denied
 Unauthenticated: 0
 Policy action: 0
 Policy not matched: 0
 Timeout action: 0
 Sessions terminated
 Reevaluation: 0
 Signout: 0

Release 12.1 R2 Hardware Features

Hardware Features—8-Port Serial GPIM

The 8-Port Serial Gigabit-Backplane Physical Interface Module (GPIM) is a network interface card (NIC) that installs in the front slots of the SRX650 and SRX550 Services Gateways to provide physical connections to a WAN. It can be plugged into all the GPIM slots of SRX650 and SRX550 Services Gateways. After you install the GPIM in the services gateway, you configure a network interface on the GPIM. It enhances port density and provides flexible support for universal serial ports including a wide array of interface options such as synchronous support.



NOTE: Only SRX650 Services Gateway supports hot-swappable functionality for GPIMs.

All the ports of the 8-Port Serial GPIM are connected to two dual stacked 136-pin VHDCI connectors. Custom cables are used to connect the VHDCI connectors to the data terminal equipment (DTE) or the data circuit-terminating equipment (DCE). Different cables are required to support different protocols. The protocols that a cable supports are identified by the cable ID. The cables are labeled with a letter (refers to the cable type) and four digits (unique serial numbers for the cables).

Table 3 on page 72 lists the hardware specifications of the 8-Port Serial Gigabit-Backplane Physical Interface Module (GPIM).

Table 3: 8-Port Serial GPIM Hardware Specifications

Description	Value
Dimensions (H x W x L)	0.756 in. x 6.723 in. x 8.234 in. (1.921 cm x 17.079 cm x 20.916 cm)
Weight	16.6 oz (0.470 kg)
Connector type	Two dual stacked 136-pin VHDCI connectors. Custom cables are used to connect the VHDCI connectors to the data terminal equipment (DTE) or the data circuit-terminating equipment (DCE)
Form factor	Single-high, single-wide GPIM

Table 3: 8-Port Serial GPIM Hardware Specifications (*continued*)

Description	Value
Environmental operating temperature	32°F through 104°F (0°C through 40°C)
Storage temperature	-40°F through 158°F (-40°C through 70°C)
Relative humidity	5% to 90% noncondensing
Altitude	Up to 10,000 ft (3,000 m)

Table 4 on page 73 lists the network interface specifications of the 8-Port Serial Gigabit-Backplane Physical Interface Module (GPIM).

Table 4: 8-Port Serial GPIM Network Interface Specifications

Description	Specification
Network Interface Specifications	
Line encoding	NRZ and NRZI
Line Protocols	- EIA-530 - EIA-530A - EIA-449 - RS-232 - V.35 - X.21
Modes	- Data terminal equipment (DTE) - Clear to send (CTS) - Data carrier detect (DCD) - Data set ready (DSR) - Data circuit-terminating equipment (DCE) - Request to send (RTS) - Data terminal ready (DTR)
Synchronous Tx/Rx rate/port	Up to 8 Mbps
Error counters	rx_error_counts, tx_abort_counts, rx_giant_counts, rx_runt_counts, and crc_err_counts
HDLC Features	
Mode	56K or 64K
PPP, Frame Relay, and Cisco-HDLC encapsulation	Supported
Shared flag	Supported

Table 4: 8-Port Serial GPIM Network Interface Specifications (*continued*)

Description	Specification
Idle flag/fill (0x7e or all ones)	Supported
Counters	Runts, giants, fcs error, abort error, and align error

Release 12.1 R1 Software Features

AppSecure

- **Synchronizing application identification package in a chassis cluster**—This feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550, SRX650, and J Series devices.

The existing application identification package download feature is modified to download the security package on the primary node and synchronize it on the secondary node.

When you download the application signature package on a device operating in chassis cluster mode, this feature enables you to download and install the package to the primary node, after which the primary and secondary nodes are automatically synchronized. This synchronization helps maintain the same version of the security package on both the primary node and the secondary node.

You can use the command **request services application-identification download** to download the application identification package to the primary node and use the command **request services application-identification download status** to verify the download status of the application package and the status of the synchronization to the secondary node.

[Junos OS Security Configuration Guide]

Chassis Cluster

- Chassis cluster is now supported on SRX550 devices in addition to existing support on SRX100, SRX210, SRX220, SRX240, and SRX650 devices.

To form a chassis cluster, two of the same model of supported SRX550 devices combine to act as a single system that enforces the same overall security. The chassis cluster design requires two ports to manage the cluster. The cluster uses one port for chassis control and one port for data plane traffic. A maximum of 9 Flexible PIC Concentrator (FPC) slots is supported on SRX550 devices. SRX550 devices use separate high-availability (HA) control and fabric ports.



NOTE: SRX550 devices also support fiber-optic ports.

On SRX550 devices, ge-0/0/1 port on node 0 and ge-9/0/1 port on node 1 are used as HA control ports.

SRX550 devices support the following features in chassis cluster mode:

- All interfaces types supported in noncluster mode
- Data plane
- Firewall filters
- IPsec
- QoS
- Redundancy groups
- Redundant Ethernet Interface
- SNMP
- WAN interfaces

Interfaces that require pseudointerface support are not supported in chassis cluster mode on SRX550 devices.

[*Junos OS Security Configuration Guide*]

Denial of Service (DoS)

- **Whitelists**—This feature is supported on all branch SRX Series devices.

You can configure a whitelist of IP addresses that are to be exempt from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process. To do so, use the new **white-list** statement at the [**edit security screen ids-option screen-name tcp syn-flood**] hierarchy level.

Both IP version 4 (IPv4) and IP version 6 (IPv6) whitelists are supported. Addresses in a whitelist must be all IPv4 or all IPv6. Each whitelist can have up to 32 IP address prefixes.

[*Junos OS CLI Reference, Junos OS Security Configuration Guide*]

Intrusion Detection and Prevention (IDP) and AppSecure

- **Synchronizing IDP security package in a chassis cluster**— This feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550, SRX650, and J Series devices.

The existing IDP security package download feature has been modified to download the security package on the primary node and synchronize it on secondary node.

When you download the IDP security package on a device operating in chassis cluster mode, this feature enables you to download and install the IDP signature database to the primary node, after which the primary and secondary nodes are automatically synchronized. This synchronization helps maintain the same version of the security package on both the primary node and the secondary node.

You can use the **request security idp security-package download** command to download the IDP signature database to the primary node.

When you check the security package download status using the **request security idp security-package download status** command, a message is displayed confirming that the downloaded security package is being synchronized on the primary and secondary nodes.

[*Junos OS Security Configuration Guide*]

J-Web

- **J-Web support for IPv6 stage 3 security features**—This feature is supported on all branch SRX Series and J Series devices.

The following J-Web pages have been added to support the IPv6 feature:

- Proxy ND Configuration page
- DS-Lite Configuration page

The following J-Web pages have been updated to include IPv6 support:

- Source NAT Pool Configuration page
- Static NAT Rule Configuration page

- [Source NAT Rule Configuration page](#)
- [Destination NAT Pool Configuration page](#)
- [Destination NAT Rule Configuration page](#)

Routing

- **Equal-cost multipath (ECMP) flow-based forwarding**—This feature is supported on all branch SRX Series and J Series devices.

An equal-cost multipath (ECMP) set is formed when the routing table contains multiple next-hop addresses for the same destination with equal cost. (Routes of equal cost have the same preference and metric values.) If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose *one* of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. On Juniper Networks devices, per-flow load balancing can be performed to spread traffic across multiple paths between routing devices. On Juniper Networks security devices, source and destination IP addresses and protocols are examined to determine individual traffic flows. Packets for the same flow are forwarded on the same interface; the interface does not change when there are additions or changes to the ECMP set. This is important for features such as source NAT, where the translation is performed only during the first path of session establishment; IDP; ALG; and route-based VPN tunnels. If a packet arrives on a given interface in an ECMP set, the security device ensures that reverse traffic is forwarded through the same interface.



NOTE: ECMP flow-based forwarding on security devices applies only to IPv4 unicast traffic flows. Multicast and IPv6 flows are not supported.

On Juniper Networks security devices, the maximum number of next-hop addresses in an ECMP set that can be installed in the forwarding table is 16. If there are more than 16 next-hop addresses in an ECMP set, only the first 16 addresses are used.

In a chassis cluster deployment, a *local* interface is an interface that is on the same node as the interface on which a packet arrives, and a *remote* interface is an interface that is on the other chassis cluster node. If an ECMP route has both local and remote interfaces in a chassis cluster, then the local interface is favored for the next hop.

To configure ECMP flow-based forwarding on Juniper Networks security devices, first define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the [edit **policy-options**] hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** configuration statements at the [edit **routing-options**] hierarchy level.

[Junos OS Routing Protocols and Policies Configuration Guide for Security Devices]

Security Policy

- **Binary format available for security log messages**—This feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550 and SRX650 devices. The feature requires a minimum of 1 gigabyte of memory.

Binary format for security log messages is available on all SRX Series devices. Because security logging can produce large amounts of data, text-based log files can quickly consume storage and CPU resources. Depending on your implementation of security logging, a log file in a binary-based format can provide more efficient use of on- or off-device storage and improved CPU utilization.

When configured in event mode, security log messages stored in binary format are maintained on the device in a separate log file from system log messages. The **show security log file** operational command supports decoding, converting, and viewing binary log files stored locally on the device. (Events stored in binary format are not accessible with advanced log scripting commands intended for text-based formats.)

When configured in stream mode, security messages generated in the data plane in binary format are streamed directly to external log collection clients in a Juniper Networks-specific binary format. The third-party client handles decoding, converting, and viewing binary log files stored on the remote device.

A programmatic interface provides the tools for developing third-party collection, decoding, conversion, and viewing applications. For information about the programmatic interface, contact your Juniper Networks representative.

[Junos OS Monitoring and Troubleshooting Guide for Security Devices]

- **Hit-count tracking**—This feature is supported on all branch SRX Series devices.

The new **show security policies hit-count** command displays the utility rate of security policies according to the number of hits they receive. You can use this feature to determine which policies are being used on the device, and how frequently they are used. Depending on the command options that you choose, the number of hits can be listed without order or sorted in either ascending or descending order, and they can be

restricted to the number of hits that fall above or below a specific count or within a range. Data is shown for all zones associated with the policies or named zones.

[*Junos OS CLI Reference, Junos OS Security Configuration Guide*]

UAC Authentication

- **Security enhancement between a Junos Pulse Access Control Service and an Infranet Enforcer**—This feature is supported on all branch SRX Series devices.

To improve security for a UAC configuration, a message appears if the Access Control Service's certificate cannot be verified on the SRX Series device. To ensure that the Access Control Service and the Enforcer are connected securely, configure the ca-profile on the SRX Series device to verify the Access Control Service's certificate.

[*Junos OS Security Configuration Guide*]

User Role Firewalls

- **Automating the Single Sign-On for User Role FW**—This feature is supported on SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650 devices.

To allow users to silently authenticate and access network resources over any protocol (not just HTTP), a hidden browser can be launched to start SPNEGO.

When a user logs on to the AD domain, the authentication server (MS Windows Domain Controller) pushes a log on script to the endpoint. When invoked at the endpoint, the script opens a browser and navigates to the URL of the Junos Pulse Access Control Service.

The following sample Jscript file defines a variable for the URL of the Access Control Service, defines the navigation function, and creates a procedure to open a browser and navigate to the specified URL.



NOTE: This is only a sample script. Your implementation will vary depending on your configuration, browser choice, administrative options, and scripting language.

```
//Variable definitions
var done = false;
var url = WScript.Arguments.Count() ? WScript.Arguments.Item(0) :
"www.your-acs.com" ;

//Event Handlers
function objIE_NavigateComplete2(pdisp, URL) {
    done = true;
}

//Main section
var objIE = WScript.CreateObject("InternetExplorer.Application", "objIE_");
objIE.Visible = false;
objIE.Navigate(url);

while (!done) {
```

```
WScript.Sleep(1000);  
}
```

Save the file using the file extension for the scripting language, and copy it to the Sysvol\Scripts location on the Active Directory server.

Add the file in the Logon Properties window in the Group Policy Operations path. See your Microsoft Active Directory administrators guide for further information.

[*Junos Pulse Administrator Guide*]

- **User role firewall providing flexibility and higher security**—This feature is supported on all branch SRX Series devices.

Network security enforcement, monitoring, and reporting based solely on IP information soon will not be sufficient for today's dynamic and mobile workforce. By integrating user firewall policies, administrators can permit or restrict network access of employees, contractors, partners, and other users based on the roles they are assigned.

A new match criteria, source-identity, defines applicable roles for each policy. In this way, traffic can be permitted or denied access based on the role of the user, as well as the zone pair, source and destination IP addresses, and application.

The SRX Series device can be configured to:

- Work with existing Junos Pulse Access Control Service features.
- Add direct support for agentless Active Directory domain controller authentication, and use MS single sign-on protocols for dynamic role processing.
- Implement local database support.

[*Junos OS Security Configuration Guide, Unified Access Control Solution Guide for SRX Series Services Gateways*]

- **User role integration into AppTrack logs**—This feature is supported on all branch SRX Series devices.

User identity details such as username and user role have been added to the AppTrack session create, session close, and volume update logs. These fields will contain the username and role associated with the policy match. The logging of username and roles are enabled only for security policies that provide user firewall enforcement. For security policies without user firewall enforcement, the username and roles are displayed as N/A. The username is displayed as unauthenticated-user and the user role is displayed as N/A, if the device cannot retrieve information for that session because there is no authentication table entry for that session or because logging of this information is disabled. The user-role field in the log will contain the list of all the roles performed by the user if match criteria is specific, authenticated-user, or any and the username field in the log contains the correct username. The user-role field in the log will contain N/A if the match criteria and the username field in the log contains un-authenticated user or unknown user.

[*Junos OS Security Configuration Guide*]

Hardware Features—SRX550 Services Gateways

Introduction

Junos OS Release 12.1 supports the Juniper Networks SRX550 Services Gateway. The SRX550 Services Gateway is a mid-range dynamic services gateway that consolidates network infrastructure and security applications for regional offices, large branch offices, and small to medium enterprises. The services gateway provides cost-effective, scalable integration of routing, security, and other mid-range applications for these sites.

Hardware Features

The SRX550 Services Gateway provides the following features:

- Symmetric Multiprocessing-based data forwarding.
- Hardware-based control and data plane separation.
- Six on-board 10/100/1000Base-T Gigabit Ethernet ports.
- Four on-board SFP Gigabit Ethernet ports.
- Support for dual AC or DC power supplies with a redundant configuration in the chassis. The 645 W AC and DC power supplies with or without Power over Ethernet (PoE) support. The AC and DC power supplies are hot-swappable.
- Junos OS support for advanced security and routing services on the Services and Routing Engine (SRE).
- The services gateway supports Gigabit-Backplane Physical Interface Modules (GPIMs) and also Mini Physical Interface Modules (Mini-PIMs). For details about the supported GPIMs and Mini-PIMs, see the *SRX Series Services Gateway for the Branch Physical Interface Modules Hardware Guide*.

Physical Specifications

[Table 5 on page 81](#) provides information on the physical specifications of the SRX550 Services Gateway.

Table 5: SRX550 Services Gateway Specifications

Specification	Value
Chassis height	2 rack units (U)
Chassis width	17.5 in. (44.4 cm)
Chassis depth	18.2 in. (46.2 cm)
Chassis weight (includes one power supply without any GPIMs or Mini-PIMs)	21.96 lb (9.96 kg)

Table 5: SRX550 Services Gateway Specifications (*continued*)

Specification	Value
Power supply output/consumption	645 W AC and 645 W DC power supply can provide: • 390 W at 12 V • 255 W PoE on a single power supply, or with redundancy using the two power supplies • 510 W PoE using the two power supply option operating as nonredundant NOTE: Using the two power supply option operating as nonredundant for up to 510 W PoE power, the administrator has the ability to prioritize the PoE ports that will receive power if an outage should occur to either the power source or to one of the power supplies.
Altitude	No performance degradation to 13,000 ft (3962.4m)
Relative humidity	Normal operation ensured in relative humidity range of 5% to 90%, noncondensing
Temperature	Normal operation ensured in temperature range of 32°F (0°C) to 104°F (+40°C) Nonoperating storage temperature in shipping container: –40° F (–40°C) to 158° F (70°C)
Seismic	Designed to meet Telcordia Technologies Zone 4 earthquake requirements
Maximum thermal output	AC power: 4400 BTU/hour DC power with one 645 W power supply unit: 2200 BTU/hour DC power with two 645 W power supply units, nonredundant: 4400 BTU/hour NOTE: These specifications are estimates and subject to change.
AC input voltage	100 to 240 VAC
AC input line frequency	50 to 60 Hz
AC system current rating	7.6 to 3.8 A



NOTE: Install the services gateway only in restricted areas, such as dedicated equipment rooms and equipment closets, in accordance with Articles 110–16, 110–17, and 110–18 of the National Electrical Code, ANSI/NFPA 70.

Software Features—SRX550 Services Gateways

Software Features

- On SRX550 devices, the following software features are supported:
 - AppSecure
 - Chassis cluster
 - Cascading style sheet (CSS) integration test
 - Cellular broadband data bridge (CX111)
 - Data and voice Application Layer Gateways (ALGs)
 - Device Management Interface (DMI), Network and Security Manager (NSM), and Junos Space
 - Dynamic VPN configuration
 - Flow
 - Intrusion Detection and Prevention (IDP) Security
 - Interfaces
 - Internet Protocol Security Virtual Private Network (IPsec VPN)
 - Internet Protocol version 4 (IPv4), encapsulation, and quality of service (QoS)
 - Internet Protocol version 6 (IPv6)
 - IPv6 CSS
 - J-Web
 - Layer 2 Transparent Mode (L2TM), standalone and high availability
 - Layer 2 Switching Support
 - Network address translation
 - Policies
 - Screen
 - Simple Network Management Protocol (SNMP)
 - Syslog
 - Uboot
 - Unified Threat Management (UTM)
 - VPN
 - Zones

Related Documentation

- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)

- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)
- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

The following current system behavior, configuration statement usage, and operational mode command usage might not yet be documented in the Junos OS documentation:

AppSecure

- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, application tracking is enabled by default. You can disable application tracking with the `set security application-tracking disable` command. This command allows you to disable and re-enable application tracking without modifying your existing zone selections.

Chassis Cluster

- On all SRX Series devices, in Junos OS Release 12.1R5 and earlier releases, when you configured some interfaces using the `interface-range` command, if you later requested the committed configuration using the `<get-configuration>` RPC with the `inherit="inherit"` and `database="committed"` attributes, the device returned `junos:changed-localtime` and `junos:changed-seconds` in the RPC reply instead of `junos:commit-localtime` and `junos:commit-seconds`.

Command-Line Interface (CLI)

- The range in the `maximum-call-duration` command has been changed from 7200 minutes to 720 minutes.
- On all SRX Series devices, the `show security flow session extensive` command has been updated to show the predefined application name.
- On all SRX Series devices, the `reauthentication` command has been introduced to specify the reauthentication interval in seconds.
- On all branch SRX Series devices, when the device is in low-memory condition and IDP is processing operations such as security package install or policy compilation, IDP aborts the operation.

To disable this functionality, a new command, `disable-low-memory-handling`, has been introduced to continue IDP operations in low-memory condition.



NOTE: The device might reboot suddenly if you enable the device to continue IDP operations in low-memory condition.

- On SRX550 devices, the **request system firmware upgrade re bios** command is supported to initiate the BIOS image upgrade.
- On all SRX Series devices, the **no-recommended** command has been introduced to specify non recommended attack objects in the dynamic attack group.
- On all branch SRX Series and J Series devices, the following commands are now supported:

CLI Command	Description
show pppoe interfaces	List all Point-to-Point Protocol over Ethernet (PPPoE) sessions.
request pppoe connect	Connect to all sessions that are down.
request pppoe connect <i>pppoe interface name</i>	Connect only to the specified session.
request pppoe disconnect	Disconnect all sessions that are up.
request pppoe disconnect <i>session id or pppoe interface name</i>	Disconnect only the specified session, identified by either a session ID or a PPPoE interface name.

- On all branch SRX Series devices, in Junos OS Release 11.4 and earlier releases, if an invalid value of `reboot_reason` was detected, the **show chassis routing-engine** command would indicate a normal shutdown. Because of this, a flash memory corruption would go undetected. This issue is now fixed in Junos OS Release 12.1. The behavior is changed, and the reboot reason displays "Unknown reboot_reason", with the invalid value of `reboot_reason`, in hexadecimal format.

Deprecated Items for Security Hierarchy

[Table 6 on page 86](#) lists deprecated items (such as CLI statements, commands, options, and interfaces).

CLI statements and commands are deprecated—rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration. We strongly recommend that you phase out deprecated items and replace them with supported alternatives.

Table 6: Items Deprecated in Release 12.1

Deprecated Item	Replacement	Hierarchy Level or Command Syntax	Additional Information
download-timeout	-	download-timeout timeout	On all branch SRX Series devices, the download-timeout command is deprecated. If the configuration is present, it will be ignored. The idpd daemon internally triggers the security package to install when an automatic download is completed. There is no need to configure any download timeout.
node	-	request security idp security-package download	On all branch SRX Series devices operating in a chassis cluster, the request security idp security-package download command with the node option is not supported: request security idp security-package download node primary request security idp security-package download node local request security idp security-package download node all

Flow and Processing

- On SRX Series devices, TCP session configuration has been enhanced. In previous releases, the TCP session was only invalidated after 2 seconds of receiving a FIN or ACK packet. As a result, if there were new incoming SYN packets (with the same five tuples) within this 2 seconds, the packets would use the same TCP session, which would be terminated within 2 seconds. Therefore a new session could not be established.

A new configuration attribute, **fin-invalidate-session**, has been introduced at the **[edit security flow tcp-session]** hierarchy level. This command invalidates a TCP session immediately on reception of a FIN or ACK packet. New incoming SYN packets will need to establish a new session.

To end a session immediately on reception of a FIN or ACK packet, use the **set security flow tcp-session fin-invalidate-session** command:



NOTE: The in-service software upgrade (ISSU) procedure is not supported while upgrading from a Junos OS release supporting this command to a Junos OS release that does not support it.

ISSU is supported in the following scenarios:

- While upgrading from a Junos OS release that does not support the command to a Junos OS release that supports the command.
 - While upgrading from a Junos OS release that supports the command to a Junos OS release that also supports the command.
 - If the `fin-invalidate-session` option is not configured.
-
- The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.

Hardware

- On SRX550 devices, the mini-USB console cable provides a “break” message to the Windows application whenever the console cable is unplugged and re-plugged. If you have configured “debugger-on-break”, the system goes to the `db>` prompt because the system receives a break character. This behavior is specific to the mini-USB console.

Interfaces and Routing

- On SRX240 and SRX650 devices, for the Layer 2 link aggregation group (LAG) interface, the hash algorithm for load balancing is now based on source IP address and destination IP address instead of source MAC address and destination MAC address.
- On SRX 550 device with DS3/E3 interface, the external clocking option is disabled to overcome from the limitation present in the hardware to support this clocking option.

Internet Protocol Security (IPsec)

- On all branch SRX Series devices, for Path Maximum Transmission Unit (PMTU) calculations, the IPsec authentication data length is fixed at 16 bytes. However, the authentication data length for packets going through the IPsec tunnel is in accordance with the authentication algorithm negotiated for that tunnel.

The authentication data lengths for the different algorithms are:

- hmac-md5-96 (12 bytes)
 - hmac-sha-256-128 (16 bytes)
 - hmac-sha1-96 (12 bytes)
- On all branch SRX Series devices, any device running Junos OS Release 11.4 and earlier and configured with SHA-256 authentication must be upgraded immediately. Earlier releases truncated the SHA-256 message digest to 96 bits from the correct size of 128 bits. The `hmac-sha-256-96` parameter fixes this compatibility issue.

Intrusion Detection and Prevention

- On all branch SRX Series devices, the idpd child process stops and exits the process when the device memory is low. However, the exit process might create a core file.

J-Web

- On all branch SRX Series and J Series devices, the username field does not accept HTML tags or the '<' and '>' characters. The following error message appears: A username cannot include certain characters, including < and >.

Remote ID Check for Site-to-Site VPNs

- As of Junos OS Release 11.4, checks are performed to validate the IKE ID received from the VPN peer device. By default, SRX Series and J Series devices validate the IKE ID received from the peer with the IP address configured for the IKE gateway. In certain network setups, the IKE ID received from the peer (which can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address) does not match the IKE gateway configured on the SRX Series or J Series device. This can lead to a Phase 1 validation failure.

To modify the configuration of the SRX Series or J Series device or the peer device for the IKE ID that is used:

- On the SRX Series or J Series device, configure the **remote-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level to match the IKE ID that is received from the peer. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.



NOTE: If you do not configure **remote-identity**, the device uses the IPv4 or IPv6 address that corresponds to the remote peer by default.

- On the peer device, ensure that the IKE ID is the same as the **remote-identity** configured on the SRX Series or J Series device. If the peer device is an SRX Series or J Series device, configure the **local-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.

Security Policy Rules



NOTE: If a *policy* is configured with multiple applications, and more than one of the applications match the traffic, then the *application* that best meets the match criteria is selected.

System Logging

The following system log messages have been introduced in Junos OS Release 12.1R6.

- When **sensor-configuration disable-low-memory-handling** is not set and IPD receives a low-memory signal, the following message appears: **Processing low memory signal**
- When **sensor-configuration disable-low-memory-handling** is set and IPD receives a low-memory signal, the following message appears: **Ignoring low memory signal**

Related Documentation

- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)
- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

ADSL Mini-PIM

- **SRX Series - ADSL Mini-PIM** - It takes more than 5 minutes for ATM interface to show up when CPE is configured in ANSI-DMT mode and CO is configured in auto mode. This occurs only with ALU 7300 DSLAM, due to limitation in current firmware version running on ADSL Mini – PIM.

AppSecure

- **Junos OS application identification**—When you create custom application or nested application signatures for Junos OS application identification, the order value must be unique among all predefined and custom application signatures. The order value determines the application matching priority of the application signature.

The order value is set with the **set services application-identification application application-name signature order** command. You can also view all signature order values by entering the **show services application-identification | display set | match order** command. You will need to change the order number of the custom signature if it conflicts with another application signature.

- J-Web pages for AppSecure are preliminary.
- Custom application signatures and custom nested application signatures are not currently supported by J-Web.
- When ALG is enabled, AppID includes the ALG result to identify the application of the control sessions. AppFW permits ALG data sessions whenever control sessions are

permitted. If the control session is denied, there will be no data sessions. When ALG is disabled, AppID relies on its signatures to identify the application of the control and data sessions. If a signature match is not found, the application is considered unknown. AppFW handles them based on the AppID result.

AX411 Access Points

- On SRX210, SRX240, and SRX650 devices, up to four access points (maximum) can be configured and managed.
- On all branch SRX Series devices, managing AX411 WLAN Access Points through a Layer 3 Aggregated Ethernet (ae) interface is not supported.

Chassis Cluster

- SRX100, SRX210, SRX220, SRX240, SRX550, and SRX650 devices have the following chassis cluster limitations:
 - Virtual Router Redundancy Protocol (VRRP) is not supported.
 - In-service software upgrade (ISSU) is not supported.
 - The 3G dialer interface is not supported.
 - On SRX Series device failover, access points on the Layer 2 switch reboot and all wireless clients lose connectivity for 4 to 6 minutes.
 - On very-high-bit-rate digital subscriber line (VDSL) mini-PIM, chassis cluster is not supported for VDSL mode.
 - Queuing on the aggregated Ethernet (**ae**) interface is not supported.
 - Group VPN is not supported.
 - Sampling features like J-Flow, packet capture, and port mirror on the **reth** interface are not supported.
 - Switching is not supported in chassis cluster mode for SRX100 devices.
 - The Chassis Cluster MIB is not supported.
 - Any packet-based services like MPLS and CLNS are not supported.
 - Isq-0/0/0—Link services Multilink Point-to-Point Protocol (MLPPP), Multilink Frame Relay (MLFR), and Compressed Real-Time Transport Protocol (CRTP) are not supported.
 - ip-0/0/0—IP-over-IP (IP-IP) encapsulation is not supported.
 - It-0/0/0—CoS for real-time performance monitoring (RPM) is not supported.
- Packet-based forwarding for MPLS and International Organization for Standardization (ISO) protocol families is not supported.
- On all J Series devices, a Fast Ethernet port from a 4-port Ethernet PIM cannot be used as a fabric link port in a chassis cluster.

- On all branch SRX Series devices, only redundant Ethernet interfaces (*reth*) are supported for IKE external interface configuration in IPsec VPN. Other interface types can be configured, but IPsec VPN might not work.
- On all J Series devices, the ISDN feature on chassis cluster is not supported.

Command-Line Interface (CLI)

- On all branch SRX and all J Series devices, the **clear services flow** command is not supported.
- On all J Series devices, RADIUS accounting is not supported.
- On SRX210 and SRX240 devices, J-Web crashes if more than nine users log in to the device by using the CLI. The number of users allowed to access the device is limited as follows:
 - For SRX210 devices: four CLI users and three J-Web users
 - For SRX240 devices: six CLI users and five J-Web users
- On J6350 devices, there is a difference in the power ratings provided by user documentation (*J Series Services Routers Hardware Guide* and PIM, uPIM, and ePIM Power and Thermal Calculator) and the power ratings displayed by CLI (by a unit of 1). The cause of this issue is a round off error, where the CLI display rounds off the value to a lower integer and the ratings provided in user documentation rounds off the value to the higher integer. As a workaround, follow the user documentation for accurate ratings.

DOCSIS Mini-PIM

- On SRX210 devices, the DOCSIS Mini-PIM delivers speeds up to a maximum of 100 Mbps throughput in each direction.

Dynamic Host Configuration Protocol (DHCP)

- On all branch SRX Series and J Series devices DHCP is not supported in a chassis cluster.
- On all branch SRX Series and J Series devices do not support DHCPv6 client authentication is not supported.

Dynamic VPN

SRX100, SRX210, and SRX240 devices have the following limitations:

- The IKE configuration for the Junos Pulse client does not support the hexadecimal preshared key.
- The Junos Pulse client IPsec does not support the Authentication Header (AH) protocol and the Encapsulating Security Payload (ESP) protocol with NULL authentication.
- When you log in through the Web browser (instead of logging in through the Junos Pulse client) and a new client is available, you are prompted for a client upgrade even

if the **force-upgrade** option is configured. Conversely, if you log in using the Junos Pulse client with the **force-upgrade** option configured, the client upgrade occurs automatically (without a prompt).

- On all branch SRX Series devices, DH-group 14 is not supported for dynamic VPN.
- On all branch SRX Series devices, when you download the Pulse client using the Mozilla browser, the “Launching the VPN Client” page is displayed when Junos Pulse is still downloading. However, when you download the Pulse client using Internet Explorer, “Launching the VPN Client” page is displayed after Junos Pulse has been downloaded and installed.

Firewall Filters

- On all SRX Series devices, firewall filters are not supported in Layer 2 transparent mode.

Flow and Processing

- On all branch SRX Series devices, sessions going through GRE over IPsec are not marked as backup on backup node, instead, they are marked as active on backup node. This behavior causes the real active session to be deleted. Nesting tunnel, like GRE over IPsec, is not supported on chassis cluster.
- On all branch SRX Series and J Series devices, a mismatch between the Firewall Counter Packet and Byte Statistics values, and between the Interface Packet and Byte Statistics values, might occur when the rate of traffic increases above certain rates of traffic.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, due to a limit on the number of large packet buffers, Routing Engine based sampling might run out of buffers for packet sizes greater than or equal to 1500 bytes and hence those packets will not be sampled. You could run out of buffers when the rate of the traffic stream is high.
- On SRX100 and SRX240 devices, the data file transfer rate for more than 20 megabits per second is reduced by 60 percent with the introduction of Junos Pulse1.0 client as compared to the Acadia client that was used before Junos OS Release 11.1.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the default authentication table capacity is 10,000; the administrator can increase the capacity to a maximum of 15,000.
- On all branch SRX Series and J Series devices, when devices are operating in flow mode, the Routing Engine side cannot detect the path maximum transmission unit (PMTU) of an IPv6 multicast address (with a large size packet).
- On all J Series devices, even when forwarding options are set to drop packets for the ISO protocol family, the device forms End System-to-Intermediate System (ES-IS) adjacencies and transmits packets because ES-IS packets are Layer 2 terminating packets.
- On all branch SRX Series and J Series devices, high CPU utilization triggered for reasons such as CPU intensive commands and SNMP walks causes the Bidirectional Forwarding Detection protocol (BFD) to flap while processing large BGP updates.

- On SRX210, SRX240, and J Series devices, broadcast TFTP is not supported when flow is enabled on the device.
- Maximum concurrent SSH, Telnet, and Web sessions — On SRX210, SRX240, and SRX650 devices, the maximum number of concurrent sessions is as follows:

Sessions	SRX210	SRX240	SRX650
ssh	3	5	5
telnet	3	5	5
Web	3	5	5



NOTE: These defaults are provided for performance reasons.

- On SRX210 and SRX240 devices, for optimized efficiency, we recommend that you limit use of CLI and J-Web to the numbers of sessions listed in the following table:

Device	CLI	J-Web	Console
SRX210	3	3	1
SRX240	5	5	1

- On SRX100 devices, Layer 3 control protocols (OSPF, using multicast destination MAC address) on the VLAN Layer 3 interface work only with access switch ports.

Group VPN Interoperability with Cisco's GET VPN for Juniper Networks Security Devices that Support Group VPN

Cisco's implementation of the Group Domain of Interpretation (GDOI) is called *Group Encryption Transport (GET) VPN*. While group VPN in Junos OS and Cisco's GET VPN are both based on RFC 3547, *The Group Domain of Interpretation*, there are some implementation differences that you need to be aware of when deploying GDOI in a networking environment that includes both Juniper Networks security devices and Cisco routers. This topic discusses important items to note when using Cisco routers with GET VPN and Juniper Networks security devices with group VPN.

Cisco GET VPN members and Juniper Group VPN members can interoperate as long as the server role is played by a Cisco GET VPN server, Juniper Networks security devices are group members, and with the following caveats:

The group VPN in Release 12.1 of Junos OS has been tested with Cisco GET VPN servers running Version 12.4(22)T and Version 12.4(24)T.

To avoid traffic disruption, do not enable rekey on a Cisco server when the VPN group includes a Juniper Networks security device. The Cisco GET VPN server implements a proprietary ACK for unicast rekey messages. If a group member does not respond to the

unicast rekey messages, the group member is removed from the group and is not able to receive rekeys. An out-of-date key causes the remote peer to treat IPsec packets as bad security parameter indexes (SPIs). The Juniper Networks security device can recover from this situation by reregistering with the server to download the new key.

Antireplay must be disabled on the Cisco server when a VPN group of more than two members includes a Juniper Networks security device. The Cisco server supports time-based antireplay by default. A Juniper Networks security device will not interoperate with a Cisco group member if time-based antireplay is used because the timestamp in the IPsec packet is proprietary. Juniper Networks security devices are not able to synchronize time with the Cisco GET VPN server and Cisco GET VPN members because the sync payload is also proprietary. Counter-based antireplay can be enabled if there are only two group members.

According to Cisco documentation, the Cisco GET VPN server triggers rekeys 90 seconds before a key expires, and the Cisco GET VPN member triggers rekeys 60 seconds before a key expires. When interacting with a Cisco GET VPN server, a Juniper Networks security device member needs to match Cisco behavior.

A Cisco GET VPN member accepts all keys downloaded from the GET VPN server. Policies associated with the keys are dynamically installed. A policy does not have to be configured on a Cisco GET VPN member locally, but a deny policy can optionally be configured to prevent certain traffic from passing through the security policies set by the server. For example, the server can set a policy to have traffic between subnet A and subnet B be encrypted by key 1. The member can set a deny policy to allow OSPF traffic between subnet A and subnet B not to be encrypted by key 1. However, the member cannot set a permit policy to allow more traffic to be protected by the key. The centralized security policy configuration does not apply to the Juniper Networks security device.

On a Juniper Networks security device, the **ipsec-group-vpn** configuration statement in the permit tunnel rule in a scope policy references the group VPN. This allows multiple policies referencing a VPN to share an SA. This configuration is required to interoperate with Cisco GET VPN servers.

Logical key hierarchy (LKH), a method for adding and removing group members, is not supported with group VPN on Juniper Networks security devices.

GET VPN members can be configured for cooperative key servers (COOP Ks), an ordered list of servers with which the member can register or reregister. Multiple group servers cannot be configured on group VPN members.

Interfaces and Routing

- On all branch SRX Series devices, IPv6 traffic transiting over IPv4 based IP over IP tunnel (for example, IPv6-over-IPv4 using ip-x/x/x interface) is not supported.
- On all Branch SRX Series devices, subnet directed broadcast feature is not supported.
- On SRX100, SRX110, SRX210, SRX220, SRX240, and SRX550 devices, Link Aggregation Control Protocol (LACP) is not supported on the 1-Port Gigabit Ethernet Small Form-Factor Pluggable (SFP) Mini-PIM.

- On SRX650 devices, Link Layer Detection Protocol (LLDP) is not supported on base ports and on 2-Port 10 Gigabit Ethernet XPIM.
- On SRX650 devices, you can only create a maximum of 63 physical interface devices with 1GB RAM capacity. Therefore, we recommend that you use only seven octal serial cards to create physical interface devices. To optimally use the 8xoctal serial cards, and to create 64 physical interface devices, you require an SRX-650 device with 2 GB RAM capacity.
- On SRX100 and J Series devices, dynamic VLAN assignments and guest VLANs are not supported.
- On SRX650 devices, Ethernet switching is not supported on Gigabit Ethernet interfaces (**ge-0/0/0** through **ge-0/0/3** ports).
- On SRX210, SRX220, SRX240, and SRX650 devices, logs cannot be sent to NSM when logging is configured in the stream mode. Logs cannot be sent because the security log does not support configuring of the source IP address for the **fxp0** interface and the security log destination in stream mode cannot be routed through the **fxp0** interface. This implies that you can not configure the security log server in the same subnet as the **fxp0** interface and the route the log server through the **fxp0** interface.
- On all branch SRX Series devices, the number of child interfaces per node is restricted to 4 on the reth interface and the number of child interfaces per reth interface is restricted to 8.
- On SRX240 High Memory devices, traffic might stop between the SRX240 device and the Cisco switch due to link mode mismatch. We recommend setting autonegotiation parameters on both ends to the same value.
- On SRX100 devices, the link goes down when you upgrade FPGA on 1xGE SFP. As a workaround, run the **restart fpc** command and restart the FPC.
- On SRX210 devices with VDLS2, ATM COS VBR-related functionality cannot be tested.
- On SRX210 devices, Internet Group Management Protocol version 2 (IGMPv2) JOINS messages are dropped on an integrated routing and bridging (IRB) interface. As a workaround, enable IGMP snooping to use IGMP over IRB interfaces.
- On all J Series devices, the DS3 interface does not have an option to configure multilink-frame-relay-uni-nni (MFR).
- On SRX210, SRX220, and SRX240 devices, every time the VDSL2 PIM is restarted in the asymmetric digital subscriber line (ADSL) mode, the first packet passing through the PIM is dropped.
- On SRX240 Low Memory devices and SRX240 High Memory devices, the RPM server operation does not work when the probe is configured with the option **destination-interface**.
- On all J Series devices, Link Layer Discovery Protocol (LLDP) is not supported on routed ports.
- In J Series xDSL PIMs, mapping between IP CoS and ATM CoS is not supported. If the user configures IP CoS in conjunction with ATM CoS, the logical interface level shaper

matching the ATM CoS rate must be configured to avoid congestion drops in segmentation and reassembly (SAR).

Example:

set interfaces at-5/0/0 unit 0 vci 1.110

set interfaces at-5/0/0 unit 0 shaping cbr 62400 ATM COS

set class-of-service interfaces at-5/0/0 unit 0 scheduler-map sche_map IP COS

set class-of-service interfaces at-5/0/0 unit 0 shaping-rate 62400 ADD IFL SHAPER

- On SRX210, SRX220, and SRX240 devices, 1-port Gigabit Ethernet SFP mini-PIM does not support switching in Junos OS Release 12.1.
- On SRX650 devices, MAC pause frame and frame check sequence (FCS) error frame counters are not supported for the interfaces **ge-0/0/0** through **ge-0/0/3**.
- On SRX240 and SRX650 devices, the VLAN range from 3967 to 4094 falls under the reserved VLAN address range, and the user is not allowed any configured VLANs from this range.
- On SRX650 devices, the last four ports of a 24-Gigabit Ethernet switch GPIM can be used either as RJ-45 or SFP ports. If both are present and providing power, the SFP media is preferred. If the SFP media is removed or the link is brought down, then the interface will switch to the RJ-45 medium. This can take up to 15 seconds, during which the LED for the RJ-45 port might go on and off intermittently. Similarly, when the RJ-45 medium is active and an a small form-factor pluggable transceiver (SFP) link is brought up, the interface will transition to the SFP medium, and this transition could also take a few seconds.
- On SRX210 devices, the USB modem interface can handle bidirectional traffic of up to 19 Kbps. On oversubscription of this amount (that is, bidirectional traffic of 20 Kbps or above), **keepalives** do not get exchanged, and the interface goes down.
- On SRX100, SRX210, SRX240, and SRX650 devices, on the Layer 3 **ae** interface, the following features are not supported:
 - Encapsulations (such as CCC, VLAN CCC, VPLS, and PPPOE) on Layer 3 **ae** interfaces
 - J-Web
 - Layer 3 **ae** for 10-Gigabit Ethernet
- On SRX100 devices, the multicast data traffic is not supported on IRB interfaces.
- On SRX240 High Memory devices, when the **system login deny-sources** statement is used to restrict the access, it blocks a remote copy (rcp) between nodes, which is used to copy the configuration during the commit routine. Use a firewall filter on the lo0.0 interface to restrict the Routing Engine access. However, if you choose to use the **system login deny-sources** statement, check the private addresses that were automatically on lo0.x and sp-0/0/0.x and exclude them from the denied list.
- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, on VLAN-tagged routed interfaces, LLDP is not all supported.

Internet Key Exchange Version 2 (IKEv2)

On all branch SRX Series devices, IKEv2 does not include support for:

- Policy-based tunnels
- Dial-up tunnels
- Network Address Translation-Traversal (NAT-T)
- VPN monitoring
- Next-Hop Tunnel Binding (NHTB) for stO—Reusing the same tunnel interface for multiple tunnels
- Extensible Authentication Protocol (EAP)
- IPv6
- Multiple child SAs for the same traffic selectors for each QoS value
- Proposal enhancement features
- Reuse of Diffie-Hellman (DH) exponentials
- Configuration payloads
- IP Payload Compression Protocol (IPComp)
- Dynamic Endpoint (DEP)

Internet Protocol Security (IPsec)

- On all branch SRX Series devices, when you enable VPN, overlapping of the IP addresses across virtual routers is supported with following limitations:
 - An IKE external interface address cannot overlap with any other virtual router.
 - An internal/trust interface address can overlap across virtual routers.
 - An **stO** interface address cannot overlap in route-based VPN in point-to-multipoint tunnels such as NHTB.
 - An **stO** interface address can overlap in route-based VPN in point-to-point tunnels.

Intrusion Detection and Prevention (IDP)

- On all branch SRX Series devices, from Junos OS Release 11.2 and later, the IDP security package is based on the Berkeley database. Hence, when the Junos OS image is upgraded from Junos OS Release 11.1 or earlier to Junos OS 11.2 or later, a migration of IDP security package files needs to be performed. This is done automatically on upgrade when the IDP daemon comes up. Similarly, when the image is downgraded, a migration (secDb install) is automatically performed when the IDP daemon comes up, and previously installed database files get deleted.

However, migration is dependent on the XML files for the installed database to be present on the device. For first-time installation, full update files are required. If the last update on the device was an incremental update, migration might fail. In such a

case, you have to manually download and install the IDP security package using the **download** or **install** CLI command before using the IDP configuration with predefined attacks or groups.

As a workaround, use the following CLI commands to manually download the individual components of the security package from the Juniper Security Engineering portal and install the full update:

- **request security idp security-package download full-update**
- **request security idp security-package install**
- On all branch SRX Series devices, IDP does not allow header checks for nonpacket contexts.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the maximum supported number of entries in the ASC table for is 100,000 entries. However, because the user land buffer has a fixed size of 1 MB as a limitation, it displays a maximum of 38,837 cache entries.
- The maximum number of IDP sessions supported is 16,384 on SRX210 devices, 32,768 on SRX240 devices, and 13,1072 on SRX650 devices.
- On all branch SRX Series devices, IDP deployed in both active/active and active/passive chassis clusters has the following limitations:
 - No inspection of sessions that fail over or fail back.
 - The IP action table is not synchronized across nodes.
 - The Routing Engine on the secondary node might not be able to reach networks that are reachable only through a Packet Forwarding Engine.
 - The SSL session ID cache is not synchronized across nodes. If an SSL session reuses a session ID and it happens to be processed on a node other than the one on which the session ID is cached, the SSL session cannot be decrypted and will be bypassed for IDP inspection.
- On all branch SRX Series devices, IDP deployed in active/active chassis clusters has a limitation that for time-binding scope source traffic, if attacks from a source (with more than one destination) have active sessions distributed across nodes, then the attack might not be detected because time-binding counting has a local-node-only view. Detecting this sort of attack requires an RTO synchronization of the time-binding state that is not currently supported.



NOTE: On SRX100 devices, IDP high availability (HA) is supported in active/backup mode.

IPv6 IPsec

The IPv6 IPsec implementation has the following limitations:

- IPv6 routers do not perform fragmentation. IPv6 hosts should either perform path maximum transmission unit (PMTU) discovery or send packets smaller than the IPv6 minimum MTU size of 1280 bytes.
- Because IPv6 addresses are 128 bits long compared to IPv4 addresses, which are 32-bits long, IPv6 IPsec packet processing requires more resources. Therefore, a small performance degradation is observed.
- IPv6 uses more memory to set up the IPsec tunnel. Therefore, the IPsec IPv4 tunnel scalability numbers might drop.
- The addition of IPv6 capability might cause a drop in the IPsec IPv4-in-IPv4 tunnel throughput performance.
- The IPv6 IPsec VPN does not support the following functions:
 - 4in6 and 6in4 policy-based site-to-site VPN, IKE
 - 4in6 and 6in4 route-based site-to-site VPN, IKE
 - 4in6 and 6in4 policy-based site-to-site VPN, Manual Key
 - 4in6 and 6in4 route-based site-to-site VPN, Manual Key
 - 4in4, 6in6, 4in6, and 6in4 policy-based dial-up VPN, IKE
 - 4in4, 6in6, 4in6, and 6in4 policy-based dial-up VPN, Manual Key
 - Remote Access—XAuth, config mode, and shared IKE identity with mandatory XAuth
 - IKE authentication—public key infrastructure/digital signature algorithm (PKI/DSA)
 - IKE peer type—Dynamic IP
 - Chassis cluster for basic VPN features
 - IKE authentication—PKI/RSA
 - Network Address Translation-Traversal (NAT-T)
 - VPN monitoring
 - Hub-and-spoke VPNs
 - Next Hop Tunnel Binding Table (NHTB)
 - Dead Peer Detection (DPD)
 - Simple Network Management Protocol (SNMP) for IPsec VPN MIBs
 - Chassis cluster for advanced VPN features
 - IPv6 link-local address

Layer 2 Transparent Mode

- DHCP server propagation is not supported in Layer 2 transparent mode.

IPv6 Support

- **NSM**—Consult the Network and Security Manager (NSM) release notes for version compatibility, required schema updates, platform limitations, and other specific details regarding NSM support for IPv6 addressing on SRX Series and J Series devices.

J-Web

- **SRX Series and J Series browser compatibility**
 - To access the J-Web interface, your management device requires the following software:
 - Language support—English-version browsers
 - Supported OS—Microsoft Windows XP Service Pack 3
 - Supported browsers

Device	Application	Supported Browsers	Recommended Browser
SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, SRX650	J-Web	• Mozilla Firefox version 3.x • Microsoft Internet Explorer version 7.0 NOTE: The PPPoE Wizard works best with Mozilla Firefox version 15.x or later.	Mozilla Firefox version 3.x

- If the device is running the worldwide version of the Junos OS and you are using the Microsoft Internet Explorer Web browser, you must disable the Use SSL 3.0 option in the Web browser to access the device.
- To use the Chassis View, a recent version of Adobe Flash that supports ActionScript and AJAX (Version 9) must be installed. Also note that the Chassis View is displayed by default on the Dashboard page. You can enable or disable it using options in the Dashboard Preference dialog box, but clearing cookies in Internet Explorer also causes the Chassis View to be displayed.
- On all branch SRX Series devices, in the J-Web interface, there is no support for changing the T1 interface to an E1 interface or vice versa. As a workaround, use the CLI to convert from T1 to E1 and vice versa.
- On all branch SRX Series and J Series devices, users cannot differentiate between Active and Inactive configurations on the System Identity, Management Access, User Management, and Date & Time pages.

- On SRX210 devices, there is no maximum length when the user commits the hostname in CLI mode; however, only 58 characters, maximum, are displayed in the J-Web System Identification panel.
- On all J Series devices, some J-Web pages for new features (for example, the Quick Configuration page for the switching features on J Series devices) display content in one or more modal pop-up windows. In the modal pop-up windows, you can interact only with the content in the window and not with the rest of the J-Web page. As a result, online Help is not available when modal pop-up windows are displayed. You can access the online Help for a feature only by clicking the **Help** button on a J-Web page.
- On all branch SRX Series devices, you cannot use J-Web to configure a VLAN interface for an IKE gateway. VLAN interfaces are not currently supported for use as IKE external interfaces.

Network Address Translation (NAT)

- Maximum capacities for source pools and IP addresses have been extended on SRX650 devices, as follows:

Devices	Source NAT Pools	PAT Maximum Address Capacity	Pat Port Number	Source NAT rules number
SRX650	1024	1024	64M	1024

Increasing the capacity of source NAT pools consumes memory needed for port allocation. When source NAT pool and IP address limits are reached, port ranges should be reassigned. That is, the number of ports for each IP address should be decreased when the number of IP addresses and source NAT pools is increased. This ensures NAT does not consume too much memory. Use the **port-range** statement in configuration mode in the CLI to assign a new port range or the **pool-default-port-range** statement to override the specified default.

Configuring port overloading should also be done carefully when source NAT pools are increased.

For source pool with port address translation (PAT) in range (64,510 through 65,533), two ports are allocated at one time for RTP/RTCP applications, such as SIP, H.323, and RTSP. In these scenarios, each IP address supports PAT, occupying 2048 ports (64,512 through 65,535) for Application Layer Gateway (ALG) module use.

- **NAT rule capacity change**—To support the use of large-scale NAT (LSN) at the edge of the carrier network, the device-wide NAT rule capacity has been changed.

The number of destination and static NAT rules has been incremented as shown in [Table 7 on page 102](#). The limitation on the number of destination-rule-set and static-rule-set has been increased.

[Table 7 on page 102](#) provides the requirements per device to increase the configuration limitation as well as to scale the capacity for each device.

Table 7: Number of Rules on SRX Series and J Series Devices

NAT Rule Type	SRX100	SRX210	SRX240	SRX650	J Series
Source NAT rule	512	512	1024	1024	512
Destination NAT rule	512	512	1024	1024	512
Static NAT rule	512	512	1024	6144	512

The restriction on the number of rules per rule set has been increased so that there is only a device-wide limitation on how many rules a device can support. This restriction is provided to help you better plan and configure the NAT rules for the device.

Power over Ethernet (PoE)

- On SRX210-PoE devices, SDK packages might not work.

Security

- J Series devices do not support the authentication order **password radius** or **password ldap** in the **edit access profile *profile-name* authentication-order** command. Instead, use **order radius password** or **ldap password**.
- On all branch SRX Series and J Series devices, the limitation on the number of addresses in an address-set has been increased. The number of addresses in an address-set now depends on the device and is equal to the number of addresses supported by the policy.

Table 8: Number of Addresses in an address-set on SRX Series and J Series Devices

Device	address-set
Default	1024
SRX100 High Memory	1024
SRX100 Low Memory	512
SRX210 High Memory	1024
SRX210 Low Memory	512
SRX240 High Memory	1024
SRX240 Low Memory	512
SRX650	1024
J Series	1024

Simple Network Management Protocol (SNMP)

- On all J Series devices, the SNMP NAT-related MIB is not supported in Junos OS Release 12.1.

Switching

- **Layer 2 transparent mode support**—On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the following features are not supported for Layer 2 transparent mode:

- Gratuitous Address Resolution Protocol (GARP) on the Layer 2 interface
- Spanning Tree Protocol (STP)
- IP address monitoring on any interface
- Transit traffic through integrated routing and bridging (IRB)
- IRB interface in a routing instance
- IRB interface handling of Layer 3 traffic



NOTE: The IRB interface is a pseudointerface and does not belong to the reth interface and redundancy group.

- On SRX100, SRX210, SRX240, and SRX650 devices, Change of Authorization is not supported with 802.1x.
- On SRX100, SRX210, SRX240, and SRX650 devices, on the routed VLAN interface, the following features are not supported:
 - IPv6 (family inet6)
 - ISIS (family ISO)
 - Class of service
 - Encapsulations (Ether circuit cross-connect [CCC], VLAN CCC, VPLS, PPPoE, and so on) on VLAN interfaces
 - CLNS
 - Protocol Independent Multicast (PIM)
 - Distance Vector Multicast Routing Protocol (DVMRP)
 - VLAN interface MAC change
 - Gratuitous Address Resolution Protocol (ARP)
 - Change VLAN-Id for VLAN interface

Unified Threat Management (UTM)

- On all J Series devices, UTM requires 1 GB of memory. If your J2320, J2350, or J4350 device has only 512 MB of memory, you must upgrade the memory to 1 GB to run UTM.

USB

- On SRX100, SRX110, SRX210, SRX220, SRX240, SRX550, and SRX650 devices, frequent plug and play of USB keys is not supported. You must wait for the device node creation before removing the USB Key.

Upgrade and Downgrade

- On all J Series devices, the Junos OS upgrade might fail due to insufficient disk space if the CompactFlash is smaller than 1-GB in size. We recommend using a 1-GB CompactFlash for Junos OS Release 10.0 and later.
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, when you connect a client running Junos Pulse 1.0 to an SRX Series device that is running a later version of Junos Pulse, the client will not be upgraded automatically to the later version. You must uninstall Junos Pulse 1.0 from the client and then download the later version of Junos Pulse from the SRX Series device.

Virtual Private Networks (VPNs)

- On SRX100, SRX210, SRX240, and SRX650 devices, while configuring dynamic VPN using the Junos Pulse client, when you select the authentication-algorithm as sha-256 in the IKE proposal, the IPsec session might not get established.

Unsupported CLI for Branch SRX Series Services Gateways and J Series Services Routers

Accounting-Options Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the **accounting**, **source-class**, and **destination-class** statements in the **[accounting-options]** hierarchy level are not supported.

Chassis Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following chassis hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set chassis craft-lockout
```

```
set chassis routing-engine on-disk-failure
```

Class-of-Service Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following class-of-service hierarchy CLI commands are not supported. However, if you enter

these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set class-of-service classifiers ieee-802.1ad
set class-of-service interfaces interface-name unit 0 adaptive-shaper
```

Ethernet-Switching Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following Ethernet-switching hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set ethernet-switching-options bpdu-block disable-timeout
set ethernet-switching-options bpdu-block interface
set ethernet-switching-options mac-notification
set ethernet-switching-options voip interface access-ports
set ethernet-switching-options voip interface ge-0/0/0.0 forwarding-class
```

Firewall Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following Firewall hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set firewall family vpls filter
set firewall family mpls dialer-filter dl term
```

Interfaces CLI Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following interface hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

- [Aggregated Interface CLI on page 106](#)
- [ATM Interface CLI on page 106](#)
- [Ethernet Interfaces on page 106](#)
- [GRE Interface CLI on page 107](#)
- [IP Interface CLI on page 107](#)
- [LSQ Interface CLI on page 107](#)
- [PT Interface CLI on page 107](#)

- [TI Interface CLI on page 108](#)
- [VLAN Interface CLI on page 108](#)

Aggregated Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
request lacp link-swtchover ae0
set interfaces ae0 aggregated-ether-options lacp link-protection
set interfaces ae0 aggregated-ether-options link-protection
```

ATM Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces at-1/0/0 container-options
set interfaces at-1/0/0 atm-options ilmi
set interfaces at-1/0/0 atm-options linear-red-profiles
set interfaces at-1/0/0 atm-options no-payload-scrambler
set interfaces at-1/0/0 atm-options payload-scrambler
set interfaces at-1/0/0 atm-options plp-to-clp
set interfaces at-1/0/0 atm-options scheduler-maps
set interfaces at-1/0/0 unit 0 atm-l2circuit-mode
set interfaces at-1/0/0 unit 0 atm-scheduler-map
set interfaces at-1/0/0 unit 0 cell-bundle-size
set interfaces at-1/0/0 unit 0 compression-device
set interfaces at-1/0/0 unit 0 epd-threshold
set interfaces at-1/0/0 unit 0 inverse-arp
set interfaces at-1/0/0 unit 0 layer2-policer
set interfaces at-1/0/0 unit 0 multicast-vci
set interfaces at-1/0/0 unit 0 multipoint
set interfaces at-1/0/0 unit 0 plp-to-clp
set interfaces at-1/0/0 unit 0 point-to-point
set interfaces at-1/0/0 unit 0 radio-router
set interfaces at-1/0/0 unit 0 transmit-weight
set interfaces at-1/0/0 unit 0 trunk-bandwidth
```

Ethernet Interfaces

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces ge-0/0/1 gigether-options ignore-l3-incompletes
```

```
set interfaces ge-0/0/1 gigether-options mpls
set interfaces ge-0/0/0 stacked-vlan-tagging
set interfaces ge-0/0/0 radio-router
set interfaces ge-0/0/0 unit 0 interface-shared-with
set interfaces ge-0/0/0 unit 0 layer2-policer
set interfaces ge-0/0/0 unit 0 accept-source-mac
set interfaces fe-0/0/2 fastether-options source-address-filter
set interfaces fe-0/0/2 fastether-options source-filtering
set interfaces ge-0/0/1 passive-monitor-mode
```

GRE Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces gr-0/0/0 unit 0 ppp-options
set interfaces gr-0/0/0 unit 0 layer2-policer
```

IP Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces ip-0/0/0 unit 0 layer2-policer
set interfaces ip-0/0/0 unit 0 ppp-options
set interfaces ip-0/0/0 unit 0 radio-router
```

LSQ Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces lsq-0/0/0 unit 0 layer2-policer
set interfaces lsq-0/0/0 unit 0 family ccc
set interfaces lsq-0/0/0 unit 0 family tcc
set interfaces lsq-0/0/0 unit 0 family vpls
set interfaces lsq-0/0/0 unit 0 multipoint
set interfaces lsq-0/0/0 unit 0 point-to-point
set interfaces lsq-0/0/0 unit 0 radio-router
```

PT Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces pt-1/0/0 gratuitous-arp-reply
set interfaces pt-1/0/0 link-mode
set interfaces pt-1/0/0 no-gratuitous-arp-reply
set interfaces pt-1/0/0 no-gratuitous-arp-request
```

```
set interfaces pt-1/0/0 vlan-tagging
set interfaces pt-1/0/0 unit 0 radio-router
set interfaces pt-1/0/0 unit 0 vlan-id
```

T1 Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces t1-1/0/0 receive-bucket
set interfaces t1-1/0/0 transmit-bucket
set interfaces t1-1/0/0 encapsulation ether-vpls-ppp
set interfaces t1-1/0/0 encapsulation extended-frame-relay
set interfaces t1-1/0/0 encapsulation extended-frame-relay-tcc
set interfaces t1-1/0/0 encapsulation frame-relay-port-ccc
set interfaces t1-1/0/0 encapsulation satop
set interfaces t1-1/0/0 unit 0 encapsulation ether-vpls-fr
set interfaces t1-1/0/0 unit 0 encapsulation frame-relay-ppp
set interfaces t1-1/0/0 unit 0 layer2-policer
set interfaces t1-1/0/0 unit 0 radio-router
set interfaces t1-1/0/0 unit 0 family inet dhcp
set interfaces t1-1/0/0 unit 0 inverse-arp
set interfaces t1-1/0/0 unit 0 multicast-dlci
```

VLAN Interface CLI

- The following CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set interfaces vlan unit 0 family tcc
set interfaces vlan unit 0 family vpls
set interfaces vlan unit 0 accounting-profile
set interfaces vlan unit 0 layer2-policer
set interfaces vlan unit 0 ppp-options
set interfaces vlan unit 0 radio-router
```

Protocols Hierarchy

- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the following CLI commands are not supported. However, if you enter these commands in the CLI editor, they will appear to succeed and will not display an error message.

```
set protocols bfd no-issu-timer-negotiation
set protocols bgp idle-after-switch-over
set protocols l2iw
set protocols bgp family inet flow
```

```
set protocols bgp family inet-vpn flow
set protocols igmp-snooping vlan all proxy
```

Routing Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following routing hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set routing-instances < instance_name > services
set routing-instances < instance_name > multicast-snooping-options
set routing-instances < instance_name > protocols amt
set routing-options bmp
set routing-options flow
```

Services Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following services hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set services service-interface-pools
```

SNMP Hierarchy

- On SRX100, SRX210, SRX220, SRX240, SRX650, and all J Series devices, the following SNMP hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set snmp community < community_name > logical-system
set snmp logical-system-trap-filter
set snmp trap-options logical-system
set snmp trap-group d1 logical-system
```

System Hierarchy

- On all SRX100, SRX210, SRX220, SRX240, and SRX650 devices, the following system hierarchy CLI commands are not supported. However, if you enter these commands in the CLI editor, they appear to succeed and do not display an error message.

```
set system diag-port-authentication
```

Related Documentation

- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)

- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

The following problems currently exist in Juniper Networks branch SRX Series Services Gateways and J Series Services Routers. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Outstanding Issues in Junos OS Release 12.1R7 for Branch SRX Series Services Gateways

Chassis Cluster

- The counter for incoming traffic on a fabric interface (used for chassis cluster) always shows zero (0). [PR/520962]

Flow and Processing

- After enabling IPv6 in flow mode, IPv6 routes are not active. [PR/824563]
- The SRX Series cluster is used as a VPN concentrator that is connected to remote VPN clients. The Internet key exchange process (daemon) tries to reuse the IP address that was previously assigned to an XAuth client. But the original Xauth attributes are overwritten when the Auth reply is received from Authd. This causes IKEd to assign a new IP address every time a Phase 1 Security Association (SA) is negotiated. As a result, multiple remote clients cannot connect through VPN. [PR/854922]

Interfaces and Routing

- When you use the aggregated redundant Ethernet (reth) interface, traffic loss might be observed when the link member is down. [PR/858519]
- When there is high CPU utilization, Link Aggregation Control Protocol (LACP) might flap and consequently the aggregated Ethernet (ae) interface link might sometimes go down. [PR/860129]
- The point-to-multipoint (P2MP) interface does not accept any multicast packets. It can lead to interoperability issues with the SSG. [PR/895090]

J-Web

- The **Global options > Proxy** screen is blank for the first time when you access it using Internet Explorer version 7.0. [PR/737675]
- The Layer 2 transparent mode feature does not work with the groups configuration. [PR/815225]

Network Address Translation (NAT)

- NAT-T does not work when the VPN is with Cisco and the VPN is initiated from the Cisco peer. The VPN negotiates using port UDP 500 instead of UDP 4500 when NAT is involved. [PR/869458]

Unified Threat Management (UTM)

- Invalid notification options are displayed in the antivirus fallback-block notification. [PR/787063]
- On branch SRX Series devices, RT core files are generated and syslogs are missing for block action by enhanced web filtering. As a workaround, run **show security utm web-filtering statistics** in the CLI mode to view the message displayed. [PR/904720]

Virtual Private Network (VPN)

- Network Address Translation-Traversal (NAT-T) might not work when VPN is with Cisco and if VPN is initiated from a Cisco peer. The VPN negotiates using port UDP 500 instead of UDP 4500 when NAT is involved. [PR/869458]

Related Documentation

- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)
- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

The following are the issues that have been resolved since Junos OS Release 12.1 for Juniper Networks branch SRX Series Services Gateways and J Series Services Routers. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Resolved Issues in Junos OS Release 12.1R7 for Branch SRX Series Services Gateways

Command-Line Interface (CLI)

- On SRX220 power over ethernet (PoE) devices, the **smtp-profile junos-as-defaults** failed to load. [PR/791575: This issue has been resolved.]

Flow and Processing

- On SRX 240 High Memory devices, segmentation faults in the syslog process caused flow core files to be generated. [PR/859346: This issue has been resolved.]
- On devices with increased ALG or proxy traffic, memory leaks in global data plane memory were observed, and traffic (FTP, MSRPC, ApplD, and so on) dropped. [PR/859956: This issue has been resolved.]
- If Virtual Router Redundancy Protocol (VRRP) was configured with the preempt option on an aggregated Ethernet link aggregation group (LAG) interface, the device did not send Gateway-Address Resolution Protocol (G-ARP). [PR/863549: This issue has been resolved.]
- Incorrect port numbers were displayed for the traffic logs. [PR/885829: This issue has been resolved.]

Interfaces and Routing

- On the asymmetric digital subscriber line (ADSL) Mini-PIM, the Asynchronous Transfer Mode (ATM) Operation, Administration, and Management (OAM) feature was not supported. [PR/835677: This issue has been resolved.]
- In an invalid subnet configuration on a multicast group, when you perform a commit or commit check, the routing protocol process (rpd) crashes and generates core files. [PR/856925: This issue has been resolved.]

J-Web

- On devices in a chassis cluster, the security zone was not populated properly on the J-Web interface port configuration page. [PR/859200: This issue has been resolved.]
- In J-Web, information on routes was not listed under the **Configure > Routing > Static Routing** section. [PR/864324: This issue has been resolved.]
- In J-Web, when 200 or more users were listed under Access Profile, all the users were not displayed. [PR/872103: This issue has been resolved.]

Unified Threat management (UTM)

- On a device with UTM antivirus enabled, the forwarding daemon generated a core file when the file transfer using FTP took a long time or when the file size exceeded 1 GB. [PR/846655: This issue has been resolved.]
- On devices in a chassis cluster, the antivirus database was not synchronized on both the cluster nodes. [PR/863181: This issue has been resolved.]
- On all branch SRX Series devices with Unified Threat Management (UTM) full antivirus (Kaspersky lab engine) enabled, traffic was dropped intermittently when there was heavy traffic load to antivirus. This was because the cache space of antivirus (MFS disk) was marked as full once it was filled and the full flag was never cleared later even though the cache space was 100 percent free. As a result, traffic to the antivirus engine was flagged as out-of-resource and the connection was reset. [PR/864775: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R6 for Branch SRX Series Services Gateways

Application Layer Gateway (ALG)

- When a specific packet was sent through an SRX Series firewall, the flowd process crashed. This issue occurred only when the SIP ALG feature was enabled, and could cause the device to drop all existing network traffic. The device recovered on its own after the flowd process restarted. [PR/840238: This issue has been resolved.]

Authentication

- The access time remaining for firewall authentication users was associated with the IP address instead of the IP address and username. [PR/843591: This issue has been resolved.]

Chassis Cluster

- When you executed the **request system zeroize** command, the configuration was not deleted. As a result, the rescue configuration was loaded instead of the factory default configuration. [PR/835687: This issue has been resolved.]

Command-Line Interface (CLI)

- You could not configure **0.0.0.0/0** in **dialer-options watch-list**. The **set interfaces dlo unit 0 dialer-options watch-list 0.0.0.0/0** command failed to commit. [PR/841371: This issue has been resolved.]

Flow and Processing

- During configuration and maintenance of an SRX Series device, occasionally the security match policies were out of sync between the Packet Forwarding Engine and the Routing Engine. In most cases, an error message was displayed during an attempt to commit the configuration. [PR/836489: This issue has been resolved.]

- On SRX240 devices, when a non standard HTTPS port was set, the uniform resource identifier (URI) was changed to the IP address and port. [PR/851741: This issue has been resolved.]

Interfaces and Chassis

- On all branch SRX Series devices, the password recovery process did not work. [PR/585092: This issue has been resolved.]

Interfaces and Routing

- On J Series devices, E1 LCP links could not go up after BERT test. [PR/600846: This issue has been resolved.]

Intrusion Detection Prevention (IDP)

- If you configured only custom attacks without installing the IDP security package, the default detector was used. If the default detector version contained the date 110307, the detector was not compatible with the engine and resulted in the generation of flowd core files. [PR/795400: This issue has been resolved.]
- The forwarding module crashed as a result of IDP processing. [PR/832608: This issue has been resolved.]
- Signatures with negate attacks did not work as expected with hardware DFA-based pattern matching and led to many false positives. [PR/848659: This issue has been resolved.]

J-Web

- In J-Web, the session expired when the idle-timeout value was set too low. [PR/830644: This issue has been resolved.]
- In J-Web, when more than one security policy was configured on a device, the first policy was not listed in the Apply-Policy section. [PR/837799: This issue has been resolved.]
- In J-Web, custom-defined applications were presented as predefined. [PR/837820: This issue has been resolved.]
- When security policies were configured using the CLI or J-Web, the value of POLO was not visible. [PR/839749: This issue has been resolved.]
- In J-Web, when Junos-self zone was added or deleted and the configuration was verified using the CLI, you were not be able to delete the policy. [PR/839950: This issue has been resolved.]
- When the policy name was "0", the preprocessor hypertext function treated it as empty, so the log output could not be viewed. [PR/853093: This issue has been resolved.]

Security Policy

- A security policy with antivirus showed incorrect count of bytes and packets in the policy statistics. [PR/841923: This issue has been resolved.]

- You could not specify the Global Address Book object in a security policy in an untrust zone. [PR/853325: This issue has been resolved.]

SNMP Traps

- On all SRX and J Series devices, the SNMP jnxJsScreenCfgChange traps were rebooted even if there were no changes to the screen configuration. [PR/835290: This issue has been resolved.]

Unified Access Control (UAC)

- On a device deployed as a Unified Access Control (UAC) enforcer in a UAC network, if session logging was enabled for UAC-enforced security policies and the UAC authentication table contained users that had many roles associated, traffic matched these policies and flowd processes crashed and core files were generated. [PR/849805: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R5 for Branch SRX Series Services Gateways

Application Layer Gateway

- The SIP ALG dropped SIP refer messages because the SIP refer message did not have a SIP URI in its Refer-To header. [PR/807090: This issue has been resolved.]

Chassis Cluster

- After reboot of both cluster nodes and RG1 failover, a reth member was held in STP learn state and even if the cable was unplugged multiple times, the link remained in the hard down state and did not recover. [PR/811002: This issue has been resolved.]

Command-Line Interface

- On SRX550 devices, the **request system firmware upgrade re bios** command to upgrade the BIOS was missing. [PR/809921: This issue has been resolved.]

Flow and Processing

- When selective packet mode was used for multicast traffic, the traffic was dropped. [PR/805432: This issue has been resolved.]
- Destination port information was missing for IPv6 packets when the firewall was in packet mode. [PR/805986: This issue has been resolved.]
- When the web-filtering type was changed from **surfcontrol-integrated** to **juniper-enhanced** and if the surfcontrol license was not installed then the output of the **show system license** command displayed that the **wf_key_surfcontrol_cpa** license was available for 29 days. [PR/812037: This issue has been resolved.]
- SRX550 devices were not sending traffic logs when the **set system syslog file filename user info** CLI statement was configured. [PR/817323: This issue has been resolved.]
- On SRX210 and SRX240 High Memory devices, kernel issues caused the devices to reboot. [PR/818905: This issue has been resolved.]

- SRX100 devices failed to send ARP requests on link aggregation interface under certain conditions. [PR/819816: This issue has been resolved.]
- In J-Web, the Dashboard refresh rate changed from 15, 30, and 60 seconds to a minimum refresh rate of 2 minutes. [PR/826053: This issue has been resolved.]
- Erroneous messages were printed from library licenses during commit. [PR/826158: This issue has been resolved.]

J-Web

- In J-Web, the Security -> Logging tab did not work. [PR/806442: This issue has been resolved.]
- the firewall policy wizard did not work. [PR/816393: This issue has been resolved.]
- If the zone name was more than 11 characters, you could not view the complete zone name while creating another policy. [PR/824684: This issue has been resolved.]
- When J-Web was opened using Internet Explorer, the non-cluster mode was not displayed. [PR/825952: This issue has been resolved.]

System Logs

- The system log message reference for Junos OS Release 12.1 and the **help syslog** command on the terminal were different from the actual message implementation. [PR/794743: This issue has been resolved.]
- When multicast sessions were closed, the **RT_FLOW_SESSION_CLOSE** entry was not entered into the system log. [PR/804569: This issue has been resolved.]

Unified Access Control (UAC)

- On all branch SRX Series devices, the device acted as a Unified Access Control (UAC) enforcer in a UAC network to ensure only qualified end users could access protected resources scenario. However, when there were many users requiring authentication, users were redirected to the login portal and the IC server reported redirect loops. [PR/817764: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R4 for Branch SRX Series Services Gateways

Application Layer Gateway

- On SRX650 devices, the SQL ALG did not work correctly when TNS RESEND (type 11) was only 8 bytes long. [PR/806893: This issue has been resolved.]

Class of Service (CoS)

- On SRX210 devices, traffic shaping did not work as expected when the shaping rate was configured on Virtual-Channels. [PR/769244: This issue has been resolved.]

Command-Line Interface

- On SRX240 High Memory devices, the **services ip-monitoring** command failed when a large number of policies were configured. [PR/771344: This issue has been resolved.]

- On all branch SRX Series and J Series devices, the **set chassis usb storage disable** command did not work when the device was configured in a chassis cluster. [PR/793844: This issue has been resolved].

DHCP

- On SRX240 High Memory devices, when the DHCP client was configured with VLAN, the DHCP leases were not acquired by the client and unicast messages were dropped. [PR/776525: This issue has been resolved.]
- On SRX240 devices, when a DHCP client was configured in a routing instance in a JSRP setup, after failover, the device remained in secondary-hold indefinitely. [PR/790872: This issue has been resolved.]
- On SRX210 High Memory devices, the dhcp-service subsystem did not work as expected when the DHCP client and server were configured in a routing instance and a default routing instance at the same time. [PR/810912: This issue has been resolved.]

Flow and Processing

- On all branch SRX Series devices, when the syn-cookie feature was enabled along with syn-flood screen with a low timeout value, high latency TCP sessions failed to establish successfully. The clients sessions received unresponsive connections because the SRX device had timed out the flow for the session. The device also dropped subsequent packets from the client due to no state found. [PR/692484: This issue has been resolved.]
- On SRX240 High-Memory devices, when the DNS ALG was enabled, the re-write rules applied on the egress interface did not work for DNS messages. [PR/785099: This issue has been resolved.]
- On SRX210 High-Memory devices, flowd core files were generated during the IDP security-package update. [PR/793417: This issue has been resolved.]
- On SRX650 devices, the routing engine failover occurred due to possible out-of-sync information about already allocated SNMP interface index values, duplicate SNMP interface index values allocated. As a result, the mib2d process crashed or the SNMP interface index value of zero was allocated for newly created interfaces. [PR806098: This issue has been resolved.]

Infrastructure

- On J Series devices, the top utility with "ores" options was not sorting the output based on resident memory size. [PR/507675: This issue has been resolved.]

Interfaces and Routing

- On all branch SRX Series and J Series devices, the RSVP interface state might show as DOWN for point-to-point interfaces, if there is a change in the Maximum Transmission Unit (MTU) after enabling RSVP on the interface. [PR/772807: This issue has been resolved.]
- On SRX240 High Memory devices, the ICMP redirect did not work for FTP traffic. [PR/776388: This issue has been resolved.]

- On SRX210 devices, after reboot, sometimes the VLAN interface was down while its physical interface member was up. [PR/791610: This issue has been resolved.]
- On SRX240 High Memory devices, after reboot, occasionally interface VLAN was down while its member physical interface was up. [PR/795363: This issue has been resolved.]
- On SRX240 High Memory devices in a chassis cluster, long pauses and timeouts were seen for SNMP walk/query. This was caused by a delay in querying the gr-0/0/0 (GRE) interface by the kernel. [PR/800735: This issue has been observed.]
- On SRX210-LM devices, lower bandwidth was configured through a scheduler queue with small buffer size on an interface faster than 2Mbps. [PR/806745: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On SRX210 High Memory devices, when the device was in low memory condition on the control plane memory, the device rebooted suddenly during the IDP security-package update. [PR/776947: This issue has been fixed.]
- On all branch SRX Series devices, IDP had the following listed issues:
 - The detector was not updated in the control plane when the update-attack-database-only flag was used during security-package installation.
 - IDP memory leaked during IDP policy load, which caused a kernel heap fragmentation issue.

[PR/778816: This issue has been resolved.]

- On SRX210 High Memory devices, during IDP policy compile the failure message "idp policy parser compile failed" was displayed due to a memory leak in the application identification configuration load. [PR/787970: This issue has been resolved.]
- On SRX240 High Memory devices, IDP policy load failed though there was sufficient memory (heap) available. This issue occurred when there was not enough contiguous memory block available in kernel heap memory. [PR/789146: This issue has been resolved.]
- On SRX240 High Memory devices, when you changed the configuration, the show security idp policy-commit-status command showed the message "Failed to add connection for dataplane". [PR/789542: This issue has been resolved.]

J-Web

- On all branch SRX Series devices, using J-Web, when you clicked Enable Log on the **Monitor > Security > IDP > Attacks** page, the page was disabled and not accessible. [PR/768559: This issue has been resolved.]
- On all branch SRX Series devices, the Help page was not available for the **Configure > Interface > Ports** page. [PR/792544: This issue has been resolved.]
- On all branch SRX devices, the "Default" radio buttons did not work after you configured the **Configure > Security > UTM > Web Filtering > Add profile > Fallback options**. [PR/794441: This issue has been resolved.]

- On all branch SRX Series and J Series devices, the Help page was not available for the **Troubleshoot > CLI terminal** page. [PR/806027: This issue has been resolved.]
- On SRX240-HM devices, the http task was high. [PR/809061: This issue has been resolved.]

Logical Systems

- On all branch SRX Series devices running logical systems, you cannot use snmpwalk for SNMP polling. [PR/791859: This issue has been resolved.]

NAT

- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices and on all J Series devices, the commit of static NAT rules failed to commit when logical system interfaces, security zone, and NAT were committed at the same time. Similarly, there were problems with committing static rules when you committed security zone and NAT at the same time. [PR/756240: This issue has been resolved.]

Unified Threat management (UTM)

- On SRX240 High Memory devices, UTM mbuf leaks were observed after several hours of traffic load. [PR/795681: This issue has been resolved.]

VPN

- On all branch SRX Series devices, IPSec phase 2 negotiation failed between Junos OS Release 11.4R2 and 11.4R3, when the device was configured with authentication-algorithm hmac-sha-256-128. [PR/793760: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1 R3 for Branch SRX Series Services Gateways

Application Layer Gateway

- On SRX240 devices, EPRT command with port as 0 causes FTP session to hang. [PR/769444: This issue has been resolved.]
- On SRX240 devices, ALG processing of traffic can sometimes result in a core dump on the device. [PR/780007: This issue has been resolved.]

Authentication

- On SRX240 devices, the web authentication page does not display properly when you try to re-authenticate after an idle time. [PR/741973: This issue has been resolved.]
- On SRX220 devices, if the local or radius user's password contains a percentage character (%), firewall authentication through the web portal will fail due to an issue in processing the percentage sign. [PR/778891: This issue has been resolved.]

AX411

- On SRX550 devices, AX411 was not supported. [PR/750288: This issue has been resolved.]

Command-Line Interface

- On SRX210 devices, the **show interface at < > extensive** command may not display the correct value when the “at” interface is **UP** on the SHDSL MPIM. [PR/738322: This issue has been resolved.]

Dynamic Host Configuration Protocol (DHCP)

- On SRX210 and SRX240 devices, the devices acting as DHCP servers may send responses to DHCP requests to an incorrect UDP destination port, if DHCP requests are forwarded to the SRX device by a DHCP relay. [PR/774541: This issue has been resolved.]

Flow and Processing

- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, activating and deactivating logical interfaces a number of times might result in generation of a flowd core file. [PR/691907: This issue has been resolved.]
- On all branch SRX Series devices, the captive portal redirect does not work with strict SYN checking option enabled in the firewall. [PR/743466: This issue has been resolved.]
- On all branch SRX Series devices, commands after STARTTLS are encrypted, and can not be understood by the SMTP parser, are causing the session to hang till TCP session is closed, with no packets forwarded. [PR/750047: This issue has been resolved.]
- On SRX220 devices, when the device sent a broadcast ARP to an L3 VLAN interface, which was restarting, it caused the forwarding to restart, resulting in traffic loss and flowd_octeon_hm core file. [PR/755204: This issue has been resolved.]
- On all branch SRX Series devices, when SYN flood packets per second (pps) over SCREEN attack-threshold, SYN cookie will be triggered by default. As design, the SRX Series device sends SYN ACK to client with ISN, and once the correct ACK is received, the device sends SYN to server. But, the ACK packet (from client) might create a session and forwarded to server. Due to no SYN packet received instead of ACK from client, server will send RST and the RST will be forwarded to client, and connection would be reset. [PR/755727: This issue has been resolved.]
- On SRX240, and SRX650 devices, in certain cases, when the device is processing a large amount of traffic, performing AppID security package update may cause the flowd process to core dump. [PR/769832: This issue has been resolved.]
- On SRX650 devices, in HA, the forwarding module may become unresponsive if reth interface is deleted while traffic is flowing through the box. There is a possibility that flowd may crash. [PR/771273: This issue has been resolved.]
- On SRX240 devices, when passing GVPN multicast traffic, flowd core dump might occur when the GVPN packet is encapsulated in PIM register message. [PR/774133: This issue has been resolved.]

- On SRX550 and SRX650 devices, flowd core files get generated in chassis cluster mode with L2 Transparent configuration when system is being shutdown. [PR/782579: This issue has been resolved.]
- On SRX650 devices, the dot1x:mode:Multiple:Suplicants are authenticated even after sending the disconnect message from radius server. [PR/786731: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On SRX240, SRX550, and SRX650 devices, in the output of the **show services application-identification application-system-cache** command, the "application-system-cache" table for P2P encrypted traffic is incorrectly marked as "Encrypted: No" instead of "Encrypted:Yes". [PR/704023: This issue has been resolved.]
- On SRX550 and SRX650 devices, if IDP or AppID is enabled, packets can be dropped in heavy traffic due to a problem with packet serialization.[PR/737353: This issue has been resolved.]
- On all branch SRX Series devices, the detector does not get updated in control plane when update-attack-database-only flag is used during security-package install. [PR/778816: This issue has been resolved.]

Interfaces and Routing

- On all branch SRX Series devices, egress queues are not supported on VLAN or IRB interfaces. [PR/510568: This issue has been resolved.]
- On SRX240 and SRX650 devices, the Track IP (ipmon) feature does not work for VLAN tagged RETH interface. (PR/575754: This issue has been resolved.)
- On J2320 devices, in some occasions when using ISDN connections, a BAD_PAGE_FAULT error might occur and ISDN connection stops working. [PR/669297: This issue has been resolved.]
- On SRX550 devices, hotswap of GPIM's and XPIM's is not supported. [PR/719882: This issue has been resolved.]
- On J2320,J2350,J4350, and J6350 devices, RPD may crash while processing configuration changes that delete virtual-router instances that include multicast interfaces. [PR/727357: This issue has been resolved.]
- On J2350 devices, ISO/CLNS prefixes with more than or /152 VPN prefix length when advertised by BGP across VPN core causes BGP adjacency flap since the remote BGP rejects the same prefix as an invalid address. This is because the ISOVPN draft allows only up to /152 prefixes. [PR/742491: This issue has been resolved.]
- On all branch SRX Series devices, if the current service status of 3G modem is stuck at "Emergency calls only", as a workaround, run the **set interfaces cl-0/0/8 modem-options init-command-string "AT+CPIN=<4 digit pin number>** command. [PR/746400: This issue has been resolved.]
- On SRX240 devices, the routing protocol daemon (rpd) might dump a core file while processing malformed RIP or RIPng message from neighbor during adjacency establishment. [PR/772601: This issue has been resolved.]

- On all branch SRX Series and J Series devices, the RSVP interface state might show as DOWN for point-to-point interfaces, if there is a change in the Maximum Transmission Unit (MTU) after enabling RSVP on the interface. [PR/772807: This issue has been resolved.]
- On SRX650 devices, when auto installation is enabled, the interface-control (dcd) process may stop and interfaces cannot be configured. [PR/ 773616: This issue has been resolved.]
- On SRX210 devices, changes made to the VPI and VCI values of ADSL interfaces do not take effect until the chassis is rebooted. [PR/783992: This issue has been resolved.]

J-Web

- On all branch SRX Series devices, when the flash storage is full, J-Web displays the "Session Expired" message. [PR/569931: This issue has been resolved.]
- On all branch SRX Series devices, the copyright year in the help files is incorrect. [PR/735607: This issue has been resolved.]
- On all branch SRX Series devices, when the installation process fails, you will not be able to repeat the installation because the 'OK' button that is required to exit the installation process is missing. [PR/739682 : This issue has been resolved.]
- On J Series devices, the EZ-Setup (J-Web Initialization setup) fails with the following error:

Fetching setup configuration....Please wait. [PR/748173: This issue has been resolved.]

- On all branch SRX Series devices, the J-Web interface is vulnerable to HTML cross-site-scripting attacks, also called XST, or Cross-Site-Tracing. [PR/752398: This issue has been resolved.]
- On SRX210 devices, the Junos OS 11.4 release fails to import node configurations when you configure chassis cluster setup using J-Web. [PR/753533: This issue has been resolved.]
- On all branch SRX Series devices, using J-Web, when you click on enable log in **Monitor > Security > IDP > Attacks** page, the page gets disabled and then the J-Web page is not accessible. [PR/ 768559: This issue has been resolved.]
- On SRX210 HE, BE, and HE-PoE devices, when you login to J-Web, the following message is displayed:

JWEB is not supported on this platform. [PR/781659: This issue has been resolved.]

- On all branch SRX Series devices, In J-Web when the router is in Cluster mode after RG failover the primary node is displayed as secondary hold in the **Dashboard > System-identification > Cluster** details. This is due to RPC get data error. [PR/786700: This issue has been resolved.]
- On all branch SRX devices, The "Default" radio buttons are not working after configuring on **Configure > Security > UTM > Web Filtering > Add profile > Fallback** options. [PR/794441: This issue has been resolved.]

Network Address Translation (NAT)

- On branch SRX Series and J Series devices, commit of static NAT rules might fail when you commit interfaces, security zone, and NAT at same time in root or LSYS. Problem might also occur that the commit of static NAT rules might fail when commit for security zone, and NAT occur at same time. [PR/756240: This issue has been resolved.]
- On SRX240 devices, static NAT rules are not being enforced when family ethernet-switching is used. [PR/785106: This issue has been resolved.]

Virtual Private Network (VPN)

- On SRX240 devices, a reconnect request using dynamic VPN might trigger a core dump in the the HTTP process (httpd-gk). [PR/753743: This issue has been resolved]
- On SRX240 devices, SIP ALG through Dynamic VPN does not work always as expected. [PR/776883: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1 R2 for Branch SRX Series Services Gateways

Application Layer Gateway

- On SRX240 devices, in certain circumstances, memory objects used by H323 ALG does not release properly once it is used. It is fixed by enforcing the releasing logic when memory object is not needed. [PR/746367: This issue has been resolved.]

AX411

- On SRX550 devices, AX411 was not supported. [PR/750288: This issue has been resolved.]

Chassis Cluster

- On SRX240 devices, the secondary Cluster node was disabled when the chassis daemon was shut down on the secondary node due to a fabric link failure. [PR/696365: This issue has been resolved.]
- On SRX210, SRX220, and SRX240 devices, when you change the reth MTU size on chassis cluster, the system crashes inconsistently. [PR/720927: This issue has been resolved.]
- On SRX650 devices, upgrade validation fails on SRX Clusters when using 'validate' option when attempting to upgrade from 10.3 or earlier versions to 10.4 or higher version. [PR/735150: This issue has been resolved.]

Command-Line Interface (CLI)

- On all branch SRX Series devices, there was no specific CLI command to display the count of sessions allowed, denied, or terminated due to UAC enforcement. [PR/733995: This issue has been resolved.]
- On SRX210 devices, the **show interface at <> extensive** command may not display the correct value when the “at” interface is **UP** on the SHDSL MPIM. [PR/738322: This issue has been resolved.]

- On all branch SRX Series devices, the Application-groups statistics are shown as unassigned and unknown for the **show services application-identification statistics application-groups** command output without displaying the details. [PR/740014: This issue has been resolved.]

DHCP

- On SRX220 devices, only the first three options present in the Options Request of a DHCPv6 Solicit/Request will be correctly populated from the dhcp-attributes specified within a local inet6 pool. [PR/741823: This issue has been resolved.]
- On SRX210 devices, after adding or deleting some vlan configurations, if an interface is made a member of vlan 1 and that interface also serves as a DHCP server, it is sometimes observed that the interface fails to respond to incoming DHCP requests. [PR/746382: This issue has been resolved.]

Flow and Processing

- On SRX100 devices, UTM/AV was not working with only 1 PC search from yahoo.co.jp, it failed to load when content size setting is set to 20. [PR/722652: This issue has been resolved.]
- On SRX240 devices, Flowd core occurred after upgrading to 11.2R5 due to mbuf not being cleaned up. [PR/734885: This issue has been resolved.]
- On SRX240 devices, when you enable "Change password every time the user logs out" on the active directory, the user is unable to change his or her password. [PR/740869: This issue has been resolved.]

Interfaces and Routing

- On SRX550 and SRX650 devices, 2-Port 10-Gigabit Ethernet XPIM was supported only in slot 6. [PR/600699: This issue has been resolved.]
- On J Series devices, with ISO-VPN setup, CLNS traffic outage might occur due to ISIS route loss. Issue might occur when sending a set of prefix via BGP the size of which is higher than the maximum IS-IS fragment size. It could also be seen by redistributing the same set of prefixes from static routes into IS-IS database. [PR/745969: This issue has been resolved.]

IPv6

- On all branch SRX Series and J Series devices, the flow bytes counters tracked on a per-interface basis are incorrect for IPv6 flows. The flow input bytes statistics are correct, but flow output bytes statistics are reversed to the source/destination interfaces. [PR/740911: This issue has been resolved.]
- On all branch SRX Series devices, IPv6 tunnels over IPv4 network does not work in flow mode. [PR/741765 : This issue has been resolved.]

J-Web

- On SRX210 devices, when you enter the destination port range as 1112 to some invalid value, it displays the following error message:

Please enter valid ending port 0-65535

This error message conflicts with the available range (1024 through 64387).

[PR/673840: This issue has been resolved.]

- On all branch SRX Series devices, if the device monitors more than one access point, a packet capture is enabled on one access point. When you try to see the details of other access points on the **Monitor > Wireless LAN** page, you see a Data Refresh Failed error message. [PR/692344: This issue has been resolved.]
- On all branch SRX and J Series devices, when configuring a NAT rule-set in J-Web interface, security zones cannot be displayed if no interfaces are configured for these security zones. [PR/703264: This issue has been resolved.]
- On all branch SRX Series devices, on the Application Tracking monitor page the automatic refresh loads the chart with the data but it does not clear the pop-up window with the message "Data is very minimal to display in PIE Chart". [PR/724995: This issue has been resolved.]
- On SRX550 and SRX650 devices, on the **Configure > Security > UTM > Antivirus > ADD page**, configuring the content-size-limit to maximum range was limited to 20...20000, but the actual content-size-limit to maximum range was 20...40000. [PR/725946: This issue has been resolved.]
- On all branch SRX Series devices, an incorrect confirm message is displayed for Download and Install application package in Appsecure Settings configuration page. [PR/728992: This issue has been resolved.]
- On all branch SRX Series devices, on the **UTM > Web Filtering** page, when adding a profile, the Surf control categories are incorrectly listed for the Web Enhanced category. [PR/729643: This issue has been resolved.]
- On all branch SRX Series devices, the 'Apply' button in the **Configure > Security > Logging** page did not work. [PR/730861: This issue has been resolved.]
- On all branch SRX Series devices, the add and the update buttons are not available in the License page when you install 30days/1day Trial license. [PR/735174: This issue has been resolved.]
- On all branch SRX Series devices, you could not configure Query timeout and Query retries options for the Sophos engine using J-Web. [PR/736810: This issue has been resolved.]
- On SRX240 devices, on the **Monitor > Security > Policy > Activities** page, when you clicked the View logs icon on a logging-enabled security policy, the J-Web session timed out. [PR/738294: This issue has been resolved.]
- On all branch SRX Series devices, the "security screen ids-options" that checks tcp flags may drop IPv6 TCP flows, when the IPv6 fragmentation extension header is present. [PR/740840: This issue has been resolved.]

- On SRX550 and SRX650 devices, reboot did not work. [PR/741014: This issue has been resolved.]
- On SRX550 devices, the **Maintain > Reboot and Snapshot** page displayed the unsupported "External storage" option. [PR/741593: This issue has been resolved.]
- On all branch SRX Series devices, the Application Tracking monitor page does not work for non default interval with Cumulative Statistics. [PR/747271: This issue has been resolved.]

SNMP

- On SRX210 devices, when a default IP address is used as an SNMP engine ID, after the device is rebooted or power cycled, the SNMP local engine ID might be incorrectly set to 80 00 0a 4c 01 00 00 00 00. [PR/613625: This issue has been resolved.]

Switching

- On SRX650 devices, though the MTU value is configured to 9192 bytes, the packets of payload size ranging between 9147 to 9150 were dropped. [PR/742310: This issue has been resolved.]

Unified Threat Management (UTM)

- On SRX650 devices, if AV is enabled, the FTP connection would hang when you are trying to access a folder with "quit" in the name. [PR/739635: This issue has been resolved.]
- On SRX220 devices, Enhanced Web Filtering requests to Websense TSC may go to connectivity/timeout fallback due to TSC only serving 10,000 requests per connection. A new method of juggling multiple connections are introduced in 12.1R1 and 11.4R3, and later releases. [PR/741697: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1 R1 for Branch SRX Series Services Gateways

Application Identification

- On all branch SRX Series devices, the Application-groups statistics for the show command output **show services application-identification statistics application-groups** showed as 'unassigned' and 'unknown' without displaying the details. [PR/740014: This issue has been resolved.]

Application Layer Gateways (ALGs)

- On SRX650 devices, when receiving frames with TCP data that had been split between multiple packets, flowd placed these into multiple internal buffers. [PR/676431: This issue has been resolved.]

Authentication

- On all branch SRX Series devices, firewall authentication supported a maximum of 64 users or groups in policy "client-match". [PR/661587: This issue has been resolved.]

- On SRX210 devices, when pass-through authentication was configured, the initial telnet session was closed with the reason "synproxy failure" instead of "TCP FIN". Subsequent sessions were closed with a correct reason. [PR/680048: This issue has been resolved.]

AX411

- On SRX210 devices, after several days the traffic no longer flowed through radio 2. [PR/579520: This issue has been resolved.]

Chassis Cluster

- On SRX100 and J2320 devices in a chassis cluster, after you configured link-speed for the **reth** interface, the **reth** interface went down and stayed down even after you deleted the configuration. [PR/466409: This issue has been resolved.]
- On SRX650 devices, the antivirus feature caused forwarding slowness at traffic peaks. [PR/610336: This issue has been resolved.]
- On SRX Series devices operating in chassis cluster mode, the xml response to the RPC validate command did not return the closing `</routing-engine>` tag for the second node. [PR/679413: This issue has been resolved.]
- On SRX650 devices, for switching deployment in chassis cluster, multicast traffic would duplicate. [PR/689153: This issue has been resolved.]
- On SRX240 devices, switching did not work. [PR/701539: This issue has been resolved.]

Command-Line Interface (CLI)

- On SRX Series devices, deleting the groups at a certain hierarchy level could erase **junos-defaults** groups. [PR/689912: This issue has been resolved.]
- On all branch SRX and J Series devices, in operation mode, the description information for screens could not be shown. [PR/712242: This issue has been resolved.]

Dynamic Host Configuration Protocol (DHCP)

- On SRX210 and SRX240 devices, when autoinstallation was configured to run on a particular interface and the default static route was set with the options discard, retain, and no-advertise, the DHCP client running on the interface tried to fetch the configuration files from the TFTP server. During this process, the UDP data port on the TFTP server were unreachable. Because the TFTP server was unreachable, the autoinstallation process remained in the configuration acquisition state. When autoinstallation was disabled, the TFTP failed. [PR/454189: This issue has been resolved.]
- On SRX210 devices, memory leak occurred in the DHCP client module when the IP address was renewed. This leak was noticed by monitoring memory usage of the dhcpd process. [PR/504471: This issue has been resolved.]
- On SRX210 devices, DHCP packets passed and were not detected as spoofed packets. [PR/681998: This issue has been resolved.]

Flow and Processing

- On SRX210 devices, the packet size of the remote end ping was less than 1480 because the packets were getting dropped on the **at** interface. The default MTU on **at** interface is 1496 and the default MTU of the remote host Ethernet interface is 1514. [PR/469651: This issue has been resolved.]
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices in chassis cluster mode, a powered device connected to both nodes rebooted when there was RGO failure. This was due to power interruption. [PR/574899: This issue has been resolved.]
- On J4350 devices, the external ICMP redirect packets were dropped. [PR/587948: This issue has been resolved.]
- On SRX210, SRX240, and J6350 devices, OSPF got stuck at the init state over the E1/T1 link when interface configurations were continuously changed. [PR/660264: This issue has been resolved.]
- On SRX240 and J2320 devices, when the LDP-based signaling was used for VPLS neighbor discovery, a restart of routing was required for the LDP-based signaling to work properly (after the configuration was committed). [PR/668555: This issue has been resolved.]
- On SRX220 devices, the throughput dropped to 20 Mbps on a 100 Mbps interface (only in Gigabit Ethernet-to-100 Mbps direction) when a bi-directional traffic was sent between Gigabit Ethernet interface and 100 Mbps interface. [PR/671248: This issue has been resolved.]
- On SRX650 devices, when Layer 2 LAG (4 x 1 Gbps) was configured, traffic was limited to 1 Gbps. [PR/677656: This issue has been resolved.]
- On all branch SRX Series devices, in PPPoE over reth, when the PPPoE session was already established, one node was rebooted, and the underlying **reth** interface would fail over to the newly rebooted node. Multicast traffic did not go through. [PR/678755: This issue has been resolved.]
- On J Series devices, when a certain number of **lt-0/0/0** logical units were configured with encapsulation VLAN, all **lt** logical interfaces were not activated. This was due to a problem with default system interface MAC allocation. [PR/680885: This issue has been resolved.]
- On all branch SRX Series devices, when multicast traffic passed through **pp0 ifl** and its underlying interface was reth, and when you added or deleted a reth member link, multicast route next hops were not updated. [PR/687652: This issue has been resolved.]
- On SRX650 devices, built in ports are not counting the FCS error statistics. [PR/690726: This issue has been resolved.]
- On SRX240 and SRX550 devices, the socket cleanup was not handled properly and caused rtlogd (security log daemon). [PR/691867: This issue has been resolved.]
- On SRX110 devices, the system boot time was high if a USB key was plugged in. [PR/695831: This issue has been resolved.]
- On SRX210 devices, Packet Forwarding Engine core files were observed periodically. [PR/697032: This issue has been resolved.]

- On all branch SRX Series devices, a memory leak occurred during audit event processing. [PR/698907: This issue has been resolved.]
- On all branch SRX Series and J Series devices, the PIM register messages were dropped on the physical interfaces that hosted multiple unnumbered interfaces. [PR/698943: This issue has been resolved.]
- On SRX220 devices, IP monitoring did not remove more than 1 route. [PR/699072: This issue has been resolved.]
- On SRX650 devices, the policer was not policing the configured rates at higher bandwidths. [PR/703003: This issue has been resolved.]
- On SRX240 devices, when you downgraded from Junos OS Release 11.4 or a later image to Junos OS Release 11.1 or an earlier image, the security package was deleted. Although it was automatically installed when the IDP daemon came up, the automatic installation failed sometimes due to application identification (AI) installation error. [PR/705113: This issue has been resolved.]
- On SRX210, SRX220, and SRX240 devices, the DSL line started renegotiation with DSLAM when you configured MTU on VDSL interface in VDSL mode. After configuring MTU on VDSL interface, the line took time to train with DSLAM and to come up. [PR/706795: This issue has been resolved.]
- On SRX240, SRX5600, and SRX5800 devices, using FTP to the st0 interface for data transfer fails occasionally. [PR/706827: This issue has been resolved.]
- On all branch SRX Series devices, the spoofing screen dropped IPv6 packets sourced from the unspecified source address (example: 0:0:0:0:0:0:0:0/128 or ::/128). [PR/721913: This issue has been resolved.]

Hardware

- On SRX100, SRX110, and SRX210 devices, when you powered on or rebooted the device, the Subscriber Identity Module (SIM) was locked. If the SIM Personal Identification Number (PIN) or the unlock code was configured in the **set interfaces cl-0/0/8 cellular-options gsm-options sim-unlock-code configuration** command, then Junos OS made an attempt to unlock the SIM only once. This was to keep the SIM from being blocked. If the SIM was blocked, you had to provide a PIN Unblocking Key (PUK) obtained from the service provider. If the wrong SIM PIN was configured, the SIM remained locked, and the administrator unlocked it by using the remaining two attempts. [PR/711812: This issue has been resolved.]

Interfaces and Routing

- On SRX240 devices, Junos XML Management protocol RPC parsing failed for configuration groups with wildcard output created with "show | display xml". It contained "<name><*></name>", which the parser could not understand. [PR/683682: This issue has been resolved.]
- On all branch SRX and J Series devices, when the peers were connected directly through the T1 link and Ge links, the ping operation to the peer link local address of the T1 link interface did not go through the T1 link; instead it was through the Ge link. [PR/684159: This issue has been resolved.]

- On SRX210 devices, the G.SHDSL line did not appear when CPE was configured with annex-auto in 2-wire or 4-wire mode with ADTRAN DSLAM, and in 2-wire mode with Cisco DSLAM. [PR/686617: This issue has been resolved.]
- On SRX210 and SRX220 devices, plugging off a cable to an interface disabled the interface at the driver level and the interface never linked up. [PR/694484: This issue has been resolved.]
- On all branch SRX Series devices, idle-timeout configuration did not work. [PR/696617: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, custom attacks were not getting detected if there was no IDP signature database license available. [PR/710147: This issue has been resolved.]

Installation

- On SRX210 devices, the Image upgrade using the validate option failed when NTP name resolution failed. [PR/692594: This issue has been resolved.]

J-Web

- On all branch SRX Series devices, the wrong page showed up while editing ppd0 and ppe0 pages. [PR/660575: This issue has been resolved.]
- On all branch SRX Series devices, after performing any configuration change in the CLI or J-Web, the pending commit pop-up window did not appear in the J-Web interface. [PR/670459: This issue has been resolved.]
- On all branch SRX Series devices, on the configure > interface page, the edit SHDSL interface pop-up window was blank. [PR/671487: This issue has been resolved.]
- On all branch SRX Series devices, PPPoE wizard support was not available when the device was operating in a chassis cluster. [PR/681117: This issue has been resolved.]
- On SRX650 devices, PPPoE wizard support was not available for the 2-Port 10-Gigabit Ethernet XPIM. [PR/681224: This issue has been resolved.]
- On SRX100 devices, editing a physical interface (fe-x/x/x) on the Interfaces > Ports page required the availability of a configured MAC address. [PR/682012: This issue has been resolved.]
- On SRX240 devices, an error message was displayed when users committed DNAT pool configuration with an IP address of 0.0.0.0/0 even though the commit executed successfully. [PR/682915: This issue has been resolved.]
- On all branch SRX Series and J Series devices, deleting the configuration using the Configure>CLI Tools>Point and Click option resulted in session expiration. [PR/684532: This issue has been resolved.]
- On all branch SRX Series devices, field validation for IP network address was missing in the PPPoE wizard, which resulted in commit errors. [PR/684663: This issue has been resolved.]

- On all branch SRX Series devices, context-sensitive Help had formatting problems with use of the PPPoE wizard. [PR/686090: This issue has been resolved.]
- On all branch SRX Series devices, the Assign to Zone option in the zone page of PPPoE wizard was not mapping the actual zone configured in the router. [PR/686843: This issue has been resolved.]
- On SRX110 devices using the Internet Explorer browser, configuring the Global option using Configure > UTM > Anti-virus was not supported. [PR/688890: This issue has been resolved.]
- On SRX210, SRX220, SRX240, and SRX650 devices, when you discarded any available MIB profile, file or predefined object from accounting-options on the Point and Click CLI Configuration page (Configure > CLI Tools > Point and Click CLI), the J-Web session timed out. [PR/689261: This issue has been resolved.]
- On SRX110 devices, the J-Web interface did not work on the SRX110H-VB models. [PR/689614: This issue has been resolved.]
- On SRX100, SRX210, SRX220, and SRX240 devices, the dashboard links were not pointing to the correct tabs in J-Web. When links were open in J-Web, they navigated only up to the top tabs. For example, when users clicked the Sessions link in the Security Resource panel, they navigated only to the top Monitor tab. [PR/689682: This issue has been resolved.]
- On all branch SRX Series devices, while creating a PPPoE connection, the wizard failed to create a security policy with a new zone. Without the security policy, the traffic did not flow through the PPPoE interface. [PR/689731: This issue has been resolved.]
- On SRX550 devices, Enhanced Web Filtering was not supported. [PR/690446: This issue has been resolved.]
- On all branch SRX Series devices, when you configured an AX411 Wireless LAN Access Point using Configure > Wireless LAN > Setting, you could not set Bolivia as a country. [PR/691824: This issue has been resolved.]
- On all branch SRX Series devices, upgrading access points using J-Web (Configure > Wireless LAN > Firmware Upgrade) did not work. [PR/694627: This issue has been resolved.]
- On all branch SRX Series devices, the "Privacy settings" values of the "Neighboring Access Points" were not listed in the Wireless LAN Monitor page (J-Web > Monitor > Wireless LAN > Neighboring Access Points). [PR/697360: This issue has been resolved.]
- On all branch SRX Series devices, after a software upgrade of the device, when users committed the configuration changes using the Configure > Wireless Lan > Settings page, the following message was displayed: "warning: requires 'ax411-wlan-ap' license". [PR/697531: This issue has been resolved.]
- On all branch SRX Series devices, when users edited an uncommitted PPPoE connection that had a zone without a firewall policy, the J-Web page blocked the zone page of the PPPoE wizard, and a warning was displayed. [PR/697891: This issue has been resolved.]
- On SRX100 and SRX110 devices, VLAN assignment was not supported under Configure > security > 802.1x > Exclusion list. [PR/698929: This issue has been resolved.]

- On SRX110 devices, when users edited physical interfaces, the OK button in the dialog box did not work. [PR/700501: This issue has been resolved.]
- On all branch SRX Series devices, in Internet Explorer, the dashboard panels did not show any data until they were refreshed. [PR/703958]
- On SRX110 devices, the Chassis View was not visible on the Dashboard page. [PR/717743: This issue has been resolved.]
- On all branch SRX Series devices, the CPU data graph was not working. [PR/719042: This issue has been resolved.]
- On all branch SRX Series devices in a chassis cluster, the primary node did not show security resources details on the Dashboard page and the message log "Security Resources:error: invalid value: firewall-policies" was generated. [PR/720435: This issue has been resolved.]
- On SRX100, SRX210, SRX220, SRX240, and SRX650 devices, after you configured the interface speed and duplex on the J-Web interface, the page showed the speed and duplex values as following:
 - Speed:10Mbps/100Mbps/1Gbps
 - Duplex:Full Duplex/Half Duplex

When you went to another page and then returned back to this page, the display changed as following:

- Speed:10m/100m/1g
- Duplex:full duplex/half duplex

[PR/721679: This issue has been resolved.]

- On SRX210 devices, when the CLI option to delete a DHCP pool was used, the J-Web interface Monitor page (Monitor > Services > DHCP > Pools) did not display the correct value in the "Excluded address" field. It displayed the text "[objec object]" in the table. [PR/723555: This issue has been resolved.]
- On all branch SRX Series devices, changes to URL category value caused value deletion in the Web-filtering profile. [PR/723884: This issue has been resolved.]
- On all branch SRX Series devices, the Application Tracking monitor page displayed a blank table instead of a warning message, when there was no data to display. [PR/724985: This issue has been resolved.]
- On SRX550 and SRX650 devices, on the Configure > Security > UTM > Antivirus > ADD page, configuring the content-size-limit to maximum range was limited to 20...20000, but the actual content-size-limit to maximum range was 20...40000. [PR/725946: This issue has been resolved.]
- On all branch SRX Series devices, the 'Apply' button in the Configure>Security>Logging page did not work. [PR/730861: This issue has been resolved.]
- On all branch SRX Series devices, you could not configure Query timeout and Query retries options for the Sophos engine using J-Web. [PR/736810: This issue has been resolved.]

- On SRX240 devices, on the Monitor>Security>Policy>Activities page, when you clicked the View logs icon on a logging-enabled security policy, the J-Web session timed out. [PR/738294: This issue has been resolved.]
- On SRX210, SRX220, SRX240, and SRX550 devices, the upload configuration file (Maintain>Configuration Management>Upload) generated an error when J-Web was accessed through the HTTP protocol. [PR/739507: This issue has been resolved].
- On SRX550 and SRX650 devices, reboot did not work. [PR/741014: This issue has been resolved.]
- On all branch SRX Series devices, on the Monitor > Application Tracking page, the application groups statistics displayed as 'unknown' and 'unassigned' without proper details even when applications were configured. [PR/741540: This issue has been resolved.]

Network Address Translation (NAT)

- On SRX650 devices, when you configured source NAT with PAT and a static NAT prefix and generated an FTP connection from the IPv4 FTP client to the IPv6 address of the device, the source NAT rule was not applied to the FTP data session, and the source address of the data session was not translated correctly. [PR/558460: This issue has been resolved.]
- On SRX240 devices, while processing an interface state change, nsd encountered an error and printed a log message indicating the error type and the interface in question. However, it was not able to handle that error condition properly and followed a null pointer. This led to a crash and a core file being generated. [PR/687465: This issue has been resolved.]
- On SRX240 devices, using the **show security nat source persistent-nat-table all family inet6** command resulted in display error and MGD core file dump. [PR/711541: This issue has been resolved.]

Switching

- On SRX210 devices, the switch fabric did not work when configured on the ge-0/0/0 and ge-0/0/1 interfaces. [PR/694631: This issue has been resolved.]

Unified Threat Management (UTM)

- On SRX650 devices, there was drop in UTMEAV throughput performance. [PR/583630: This issue has been resolved.]
- On SRX650 devices, a commit error occurred when the policy used UTM "application-services"warning: license not installed for". This was only a display issue. [PR/600941: This issue has been resolved.]
- On SRX220 devices, the maximum number of connections per second supported by the Enhanced Web Filtering solution was less than that supported by the Surf Control solution. [PR/609094: This issue has been resolved.]
- On all branch SRX Series devices, performance drop was observed. [PR/671777: This issue has been resolved.]

- On SRX240 and SRX650 devices, the resident set size memory of the UTM daemon increased during host-inbound traffic when the UTM feature was configured. [PR/675114: This issue has been resolved.]
- On SRX550 devices, USB Modem was not supported. [PR/678657: This issue has been resolved.]
- On all branch SRX Series devices, while configuring UTM on J-Web interface, validation went to an infinite loop when the engine type was not chosen in Anti Virus > Global options. Also, configuring and then clicking the OK button in global options in Configure>UTM> Anti Spam page resulted in inaccessibility of J-Web. [PR/683649: This issue has been resolved.]
- On all branch SRX Series devices, when users configured UTM features, using the Configure > UTM > web filtering > global options page, the interface button did not work. Also, the custom objects page on UTM, sometimes rejected certain values while configuring the URL category name. [PR/685534: This issue has been resolved.]
- On SRX210 PoE devices, whenever the SMTP client issued a pipeline command, it caused errors in tagging spam e-mails. [PR/696109: This issue has been resolved.]
- On SRX550 devices, UTM EWF functionality was not supported. [PR/733527: This issue has been resolved.]

Upgrade and Downgrade

- On SRX550 devices, U-boot upgrade was not supported. Customers were advised not to upgrade u-boot even if a different version was available. [PR/731819: This issue has been resolved.]

Virtual Private Network (VPN)

- On SRX100 devices, the SCEP enrollment request to non-Microsoft CA was rejected due to an invalid RSA signature. [PR/567846: This issue has been resolved.]
- On SRX220 devices with IPsec route-based VPN configured, ST interface input counters did not increment. The traffic worked fine over the VPN, but it was not counted in the interface statistics. [PR/672738: This issue has been resolved.]
- On SRX240 devices, the packets containing more than 1667 bytes passing through the Layer3 VPN on the st0 interface failed. [PR/681057: This issue has been resolved.]

Related Documentation

- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)
- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)

- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

Errata for the Junos OS Software Documentation

This section lists outstanding issues with the software documentation.

Feature Support Reference for SRX Series and J Series Devices

- In Junos OS Feature Support Reference for SRX Series and J Series Devices, the Chassis Cluster table incorrectly indicates that Layer 2 Ethernet switching capability in chassis cluster mode is supported on SRX100 devices. Layer 2 Ethernet switching capability in chassis cluster mode is not supported on SRX100 devices.
- In this guide, the Table 14: DHCP Support of the “Dynamic Host Configuration Protocol” section incorrectly states that DHCPV6 relay agent is supported on SRX100, SRX110, SRX210 SRX220, SRX240, and SRX650 devices. The DHCPV6 relay agent is not supported on Branch SRX Series devices.

IP Monitoring webpage

- In “Example: Configuring IP Monitoring,” the **policy** keyword is missing from the **set** command in both the “CLI Quick Configuration” section and in Step 8 in the “Step-by-Step Procedure” section. The correct command should be:

[edit]

```
user@host# set services ip-monitoring policy Payment-Server-Tracking then  
preferred-route route 1.1.1.0/24 next-hop 1.1.1.99
```

J Series Services Router Advanced WAN Access Configuration Guide

- The example given in the “Configuring Full-Cone NAT” section in the guide available at <http://www.juniper.net/techpubs/software/jseries/junos85/index.html> is incorrect. The correct and updated example is given in the revised guide available at <http://www.juniper.net/techpubs/software/jseries/junos90>).

J2320, J2350, J4350, and J6350 Services Router Getting Started Guide

- The “Connecting to the CLI Locally” section states that the required adapter type is DB-9 female to DB-25 male. This is incorrect; the correct adapter type is DB-9 male to DB-25 male.

J-Web

- **J-Web Security Package Update Help page**—This Help page does not contain information about the download status.

- **J-Web pages for stateless firewall filters**—There is no documentation describing the J-Web pages for stateless firewall filters. To find these pages in J-Web, go to **Configure>Security>Firewall Filters**, and then select **IPv4 Firewall Filters** or **IPv6 Firewall Filters**. After configuring the filters, select **Assign to Interfaces** to assign your configured filters to interfaces.
- **J-Web configuration Instructions**— Because of ongoing J-Web interface enhancements, some of the J-Web configuration example instructions in the Junos administration and configuration guides became obsolete and thus were removed. For examples that are missing J-Web instructions, use the provided CLI instructions.

Junos OS Class of Service Configuration Guide for Security Devices

- In this guide, the following note has been added to the “CoS Device Configuration Overview” section: When the T1/E1 mPIM is oversubscribed, we recommend that you configure its shaping rate for consistent CoS functionality. The shaping rate should be less than the total link speed.

Junos OS CLI Reference Guide

- The [edit applications application *application-name* term *term-name*] hierarchy level for the alg (Applications) configuration statement is incorrect. The correct hierarchy level is [edit applications application *application-name* <term *term-name*>].
- The “sessions-per-client” topic lacks the following support note: The **sessions-per-client limit** CLI command supports the antispam, content filtering, and antivirus UTM features, but it does not support Web filtering.
- This guide incorrectly lists the following topics. These commands are not supported:
 - disable-call-id-hiding
 - show security alg sip transactions
- In the “show security policies” topic, the “show security policies Output Fields” table includes the following incorrect information:

Applications	ALG: If an ALG is associated with the session, the name of the ALG. Otherwise, 0.

The correct information is:

Applications

ALG: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed.

However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications.

- The Description for the limit (Security SCTP) option under the hierarchy [edit security gprs sctp profile profile-name], incorrectly states you can set the global messages rate limit per session. The description should say that you can set the rate limit for local Services Processing Unit (SPU) packets.

Junos OS Layer 2 Bridging and Switching Configuration Guide for Security Devices

- In this guide, the section "Configuring Layer 2 Bridging and Transparent Mode" includes an incorrect example, "Example: Configuring Layer 2 Trunk Interfaces with Multiple Units." The example is in error because the SRX Series devices do not support multiple units.

Junos OS High Availability Configuration Guides

- In this guide, the **vrrp-inherit-from vrrp-group** command has been added to the following topics: "VRRP Configuration Hierarchy" and "vrrp-group".

Junos OS Interfaces Configuration Guide for Security Devices

- In this guide, Table 11, "MTU Values for the SRX Series Services Gateways PIMs," does not specify the maximum MTU and default IPMTU values for the following PIMs:
 - 2-Port 10 Gigabit Ethernet XPIM
 - 16-Port Gigabit Ethernet XPIM
 - 24-Port Gigabit Ethernet XPIM

The following table lists these values:

Table 9: MTU Values for the SRX Series Services Gateways PIMs

PIM	DefaultMediaMTU(Bytes)	MaximumMTU(Bytes)	DefaultIPMTU (Bytes)
2-Port 10 Gigabit Ethernet XPIM	1514	9192	1500
16-PortGigabitEthernetXPIM	1514	9192	1500

Table 9: MTU Values for the SRX Series Services Gateways PIMs (*continued*)

24-Port Gigabit Ethernet XPIM	1514	9192	1500

Junos OS Monitoring and Troubleshooting Guide for Security Devices

- In the “Example: Enabling Packet Capture on a Device” section, the CLI quick configuration shows an incorrect **set** command. The correct set command is **set forwarding-options packet-capture file filename pcap-file files 100 size 1024 world-readable**.
- In this guide, the watchdog disable option was removed from Step 3 in “Recovering the Root Password for SRX Series Devices” as it does not support SRX Series devices.

Junos OS MPLS Configuration Guide for Security Devices

- In this guide, section *VPLS Exceptions on J Series and SRX Series Devices* includes an incorrect information mentioned as follows:

LDP signaling is not supported.

In this guide, section *BGP Signaling* includes an incorrect note mentioned as follows:

LDP signaling is not supported for the VPLS routing instance.

Junos OS Security Basics

- The topic Understanding Policy Application Timeouts Contingencies under **Security Basics > Security Policy Applications for Security Devices > Policy Application Timeout**, contains erroneous information. It should read as follows:

When setting timeouts, be aware of the following contingencies:

- If an application contains several application rule entries, all rule entries share the same timeout. You need to define the application timeout only once. For example, if you create an application with two rules, the following commands will set the timeout to **20** seconds for both rules:

```
user@host# set applications application test protocol tcp destination-port 1035-1035
inactivity-timeout 20
```

```
user@host# set applications application test term test protocol udp
```

```
user@host# set applications application test term test source-port 1-65535
```

```
user@host# set applications application test term test destination-port 1111-1111
```

- If multiple custom applications are configured with custom timeouts, then each application will have its own custom application timeout. For example:

```
user@host# set applications application ftp-1 protocol tcp source-port 0-65535
destination-port 2121-2121 inactivity-timeout 10
```

```
user@host# set applications application telnet-1 protocol tcp source-port 0-65535
destination-port 2300-2348 inactivity-timeout 20
```

With this configuration, Junos OS applies a 10-second timeout for destination port **2121** and a 20-second timeout for destination port **2300** in an application group.

Junos OS Security Configuration Guide

In the “Understanding SIP ALGs and NAT” topic, information in the following sections is incorrect:

- **Call Re-INVITE Messages**

This section incorrectly states:

When one or more media sessions are removed from a call, pinholes are closed and bindings released just as with a BYE message.

The correct information is:

When all the media sessions or media pinholes are removed from a call, the call is removed when a BYE message is received.

- **Call Session Timers**

This section incorrectly states:

The SIP ALG uses the **session-expires** value to time out a session if a Re-INVITE or UPDATE message is not received. The ALG receives the **session-expires** value, if present, from the 200 OK responses to the INVITE and uses this value for signaling timeout. If the ALG receives another INVITE before the session times out, the ALG resets all timeout values to this new INVITE or to default values, and the process is repeated. As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist.

The correct information is (The **session-expires** value is not supported on SRX Series devices):

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist.

- **Table Requesting Messages with NAT Table**

This table incorrectly states:

Outbound Request (from private to public)	Route:	Replace ALG address with local address

The correct information is:

Outbound Request (from private to public)	Route:	Replace local address with ALG address

- The “Unified Threat Management Overview” and “Web Filtering Overview” lack the following support note: The **sessions-per-client limit** CLI command supports the antispam, content filtering, and antivirus UTM features, but it does not support Web filtering.

- In “Example: Configuring a Route-Based VPN,” the **show security zones** output for the SRX Series device erroneously shows host-inbound-traffic configured for the vpn-chicago zone; this configuration is not included in the CLI Quick Configuration and the Step-by-Step Procedure.
- In Example: Configuring a User Role Firewall on an SRX Series Device and Acquiring User Role Information from an Active Directory Authentication Server, the **redirect-url** option in step 2 of the redirection procedure is incorrect. The URL and variables should be enclosed in quotation marks.

```
[edit]
user@host# set services unified-access-control captive-portal acs-device
redirect-url "https://%ic-url%/?target=%dest-url%&enforcer=%enforcer-id%"
```

- In “Example: Configuring AppTrack,” of the *Junos OS Security Configuration Guide for Security Devices*, the **set security log mode stream** statement was omitted from the log configuration statements. The updated log configuration should read:

```
user@host# set security log mode stream
user@host# set security log format sd-syslog
user@host# set security log source-address 5.0.0.254
user@host# set security log stream app-track-logs host 5.0.0.1
```

- In Chapter 48, “AppTrack Application Tracking,” of the *Junos OS Security Configuration Guide for Security Devices*, the operational command **show security application-tracking statistics applications** is incorrect. The operational command to display all application identification statistics, including application tracking statistics, is **show services application-identification statistics applications**.

The operational command **show application-tracking counters** is incorrect. The command should be **show security application-tracking counters**.

- In this guide, the “Understanding Chassis Cluster Redundancy Group Interface Monitoring” section includes a note that states interface monitoring is not supported on Redundancy Group 0 (RG0) for branch SRX Series devices. This note should also state that it is not supported for J Series devices.
- In this guide, the section “Understanding Chassis Cluster Redundancy Group Interface Monitoring” has a misleading statement in the steps describing the “Policy Lookup for a User” figure. It states that the device forwards the packet from its buffer to its destination IP address 1.2.2.2. However, this is true only for HTTP and Telnet traffic. For FTP traffic, after successful authentication the device closes the session and the user must reconnect to the FTP server at IP address 1.2.2.2.

Junos OS Security Network Address Translation

- In Example: Configuring NAT for Multiple ISPs under Network Address Translation for Security Devices > Configuration > NAT for Multiple ISPs the statement **set routing-options rib-groups isp import-rib inet.0** was omitted from the configuration. The updated configuration should read:

```
set routing-options rib-groups isp import-rib inet.0
set routing-options rib-groups isp import-rib isp1.inet.0
set routing-options rib-groups isp import-rib isp2.inet.0
```

In addition, because zone based address-book for NAT rules is unsupported, you should not use the statements provided in the example; use global address book instead.

- The command **show security nat source persistent-nat-table** under **Network Address Translation > Administration > Source NAT Operational Commands** is:

- Missing the option: **summary**—Display persistent NAT bindings summary.
- Contains incomplete sample output. The corrected sample output is as follows:

```
user@host> show security nat source persistent-nat-table internal-ip 9.9.9.1 internal-port 60784
```

Internal	Left_time/	Curr_Sess_Num/	Source	Reflective	Source	Type
In_IP	In_Port	I_Protocol	Ref_IP	Ref_Port	R_Protocol	NAT Pool
9.9.9.1	60784	udp	66.66.66.68	60784	udp	dynamic-customer-source
any-remote-host	254/300	0/30	105			

```
user@host> show security nat source persistent-nat-table all
```

Internal	Left_time/	Curr_Sess_Num/	Source	Reflective	Source	Type
In_IP	In_Port	I_Protocol	Ref_IP	Ref_Port	R_Protocol	NAT Pool
Conf_time	Max_Sess_Num	NAT Rule				
9.9.9.1	63893	tcp	66.66.66.68	63893	tcp	dynamic-customer-source
any-remote-host	192/300	0/30	105			
9.9.9.1	64014	udp	66.66.66.68	64014	udp	dynamic-customer-source
any-remote-host	244/300	0/30	105			
9.9.9.1	60784	udp	66.66.66.68	60784	udp	dynamic-customer-source
any-remote-host	254/300	0/30	105			
9.9.9.1	57022	udp	66.66.66.68	57022	udp	dynamic-customer-source
any-remote-host	264/300	0/30	105			
9.9.9.1	53009	udp	66.66.66.68	53009	udp	dynamic-customer-source
any-remote-host	268/300	0/30	105			
9.9.9.1	49225	udp	66.66.66.68	49225	udp	dynamic-customer-source
any-remote-host	272/300	0/30	105			
9.9.9.1	52150	udp	66.66.66.68	52150	udp	dynamic-customer-source
any-remote-host	274/300	0/30	105			
9.9.9.1	59770	udp	66.66.66.68	59770	udp	dynamic-customer-source
any-remote-host	278/300	0/30	105			
9.9.9.1	61497	udp	66.66.66.68	61497	udp	dynamic-customer-source
any-remote-host	282/300	0/30	105			
9.9.9.1	56843	udp	66.66.66.68	56843	udp	dynamic-customer-source
any-remote-host	-/300	1/30	105			

```
user@host> show security nat source persistent-nat-table summary
```

Persistent NAT Table Statistics on FPC5 PIC0:

binding total : 65536

binding in use : 0

enode total : 524288

enode in use : 0

Junos OS System Log Messages Reference

- The *AV System Log Messages* topic lists incorrect facilities for the systems logs.

On all SRX Series devices, antivirus (AV) system logs are generated with the facility **LOG_USER** or **LOG_DAEMON**.

[Table 10 on page 142](#) shows the correct facilities for the system logs.

Table 10: Antivirus System Logs

System Logs	Incorrect Facility	Correct Facility
AV_PATTERN_GET_FAILED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_KEY_EXPIRED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_KL_CHECK_FAILED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_TOO_BIG	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_UPDATED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_WRITE_FS_FAILED	LOG_FIREWALL	LOG_DAEMON
AV_SCANNER_READY	LOG_FIREWALL	LOG_DAEMON
AV_VIRUS_DETECTED_MT	LOG_PFE	LOG_USER

Junos OS WLAN Configuration and Administration Guide

- This guide is missing information that the AX411 Access Point can be managed from SRX100 and SRX110 devices.
- This guide is missing the information that on all branch SRX devices, managing AX411 WLAN Access Points through an Layer 3 Aggregated Ethernet (ae) interface is not supported.

Standards Supported by the Junos OS

- In this guide, the “Supported MPLS Standards” section has an incorrect statement regarding node protection in facility backup. This has been corrected as follows: Node protection in facility backup is supported for RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*.

Various Guides

- Some Junos OS user, reference, and configuration guides—for example the [Junos Software Routing Protocols Configuration Guide](#), [Junos OS CLI User Guide](#), and [Junos OS System Basics Configuration Guide](#)—mistakenly do not indicate SRX Series device support in the “Supported Platforms” list and other related support information; however, many of those documented Junos OS features are supported on SRX Series devices. For full, confirmed support information about SRX Series devices, please refer to the [Junos OS Feature Support Reference for SRX Series and J Series Devices](#).

Errata for the Junos OS Hardware Documentation

This section lists outstanding issues with the hardware documentation.

AX411 Access Point Hardware Guide

- This guide incorrectly documents the maximum number of supported access points on the SRX Series devices. It should state that on the SRX210, SRX240, and SRX650 devices, you can configure and manage up to four access points (maximum).
- This guide is missing information that the AX411 Access Point can be managed from SRX100 and SRX110 devices.

J Series Services Routers Hardware Guide

- The procedure “Installing a DRAM Module” omits the following condition:

All DRAM modules installed in the router must be the same size (in megabytes), type, and manufacturer. The router might not work properly when DRAM modules of different sizes, types, or manufacturer are installed.
- This guide incorrectly states that only the J2350 Services Router complies with Network Equipment Building System (NEBS) criteria. It should state that the J2350, J4350, and J6350 routers comply with NEBS criteria.
- This guide is missing information about 100Base-LX connector support for 1-port and 6-port Gigabit Ethernet uPIMs.

SRX Series Services Gateways for the Branch Physical Interface Modules Hardware Guide

- In the “SRX Series Services Gateway Interfaces Power and Heat Requirements” section, the PIM Power Consumption Values table contains the power consumption value for the 1-port Gigabit Ethernet Small Form-Factor Pluggable (SFP) Mini-PIM value as 3.18 W.

The correct power consumption value for the 1-port Gigabit Ethernet Small Form-Factor Pluggable (SFP) Mini-PIM is 4.4 W.

- This guide incorrectly documents that slot 3 of the SRX550 Services Gateway can be used to install GPIMs. The correct information is:
 - In Table 10: SRX Series Services Gateway Interface Port Number Examples, for 2-Port 10 Gigabit Ethernet XPIM, you can install the XPIM only in slot 6 of the SRX550 Services Gateway.
 - In Table 44: Slots for 20-Gigabit GPIMs, for 20-Gigabit GPIM slots, you can install the GPIM only in slot 6 of the SRX550 Services Gateway.

SRX100 Services Gateway Hardware Guide

- In the “Connecting an SRX100 Services Gateway to the J-Web Interface” section, the following information is missing in the note:



NOTE: Microsoft Internet Explorer version 6.0 is also supported as backward compatible from Microsoft Internet Explorer version 7.0.

SRX210 Services Gateway Hardware Guide

- In the “Connecting an SRX210 Services Gateway to the J-Web Interface” section, the following information is missing in the note:



NOTE: Microsoft Internet Explorer version 6.0 is also supported as backward compatible from Microsoft Internet Explorer version 7.0.

SRX240 Services Gateway Hardware Guide

- In the “Connecting the SRX240 Services Gateway to the J-Web Interface” section, the following information is missing in the note:



NOTE: Microsoft Internet Explorer version 6.0 is also supported as backward compatible from Microsoft Internet Explorer version 7.0.

SRX110 Services Gateway 3G USB Modem Quick Start

- The SRX110 Services Gateway 3G USB Modem Quick Start is updated with the J-Web procedures and it is available on the Juniper website.

SRX210 Services Gateway 3G ExpressCard Quick Start

- Several tasks are listed in the wrong order. “Task 6: Connect the External Antenna” should appear before “Task 3: Check the 3G ExpressCard Status,” because the user needs to connect the antenna before checking the status of the 3G ExpressCard. The correct order of the tasks is as follows:

1. Install the 3G ExpressCard
 2. Connect the External Antenna
 3. Check the 3G ExpressCard Status
 4. Configure the 3G ExpressCard
 5. Activate the 3G ExpressCard Options
- In “Task 6: Connect the External Antenna,” the following sentence is incorrect and redundant: “The antenna has a magnetic mount, so it must be placed far away from radio frequency noise sources including network components.”
 - In the “Frequently Asked Questions” section, the answer to the following question contains an inaccurate and redundant statement:

Q: Is an antenna required? How much does it cost?

A: The required antenna is packaged with the ExpressCard in the SRX210 Services Gateway 3G ExpressCard kit at no additional charge. The antenna will have a magnetic mount with ceiling and wall mount kits within the package.

In the answer, the sentence “The antenna will have a magnetic mount with ceiling and wall mount kits within the package” is incorrect and redundant.

SRX210 Services Gateway Quick Start Guide

- The section on installing software packages is missing the following information:

On SRX210 devices, the `/var` hierarchy is hosted in a separate partition (instead of the `root` partition). If Junos OS installation fails as a result of insufficient space:

 1. Use the **`request system storage cleanup`** command to delete temporary files.
 2. Delete any user-created files both in the `root` partition and under the `/var` hierarchy.

SRX650 Services Gateway Hardware Guide

- The “Maintaining the SRX650 Services Gateway Power Supply” section incorrectly states that the status of the power supplies on the SRX650 Services Gateway can be checked by issuing the **`show chassis environment pem`** command. The **`show chassis environment pem`** command is not supported on the SRX650 Services Gateway.

Related Documentation

- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)
- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)

- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 146](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers

This section includes the following topics:

- [Upgrading and Downgrading among Junos OS Releases on page 146](#)
- [Upgrading an AppSecure Device on page 148](#)
- [Upgrade and Downgrade Scripts for Address Book Configuration on page 148](#)
- [Hardware Requirements for Junos OS Release 12.1 for SRX Series Services Gateways and J Series Services Routers on page 151](#)

Upgrading and Downgrading among Junos OS Releases

All Junos OS releases are listed in sequence on the JUNOS Software Dates & Milestones webpage:

<http://www.juniper.net/support/eol/junos.html>

To help in understanding the examples that are presented in this section, a portion of that table is replicated here. Note that releases footnoted with a 1 are Extended End-of-Life (EOL) releases.

Product	FRS Date
Junos 12.1	03/28/2012
Junos 11.4 ¹	12/21/2011
Junos 11.3	08/15/2011
Junos 11.2	08/03/2011
Junos 11.1	03/29/2011
Junos 10.4 ¹	12/08/2010
Junos 10.3	08/15/2010
Junos 10.2	05/28/2010
Junos 10.1	02/15/2010
Junos 10.0 ¹	11/04/2009
Junos 9.6	08/06/2009
Junos 9.5	04/14/2009
Junos 9.4	02/11/2009
Junos 9.3 ¹	11/14/2008
Junos 9.2	08/12/2008
Junos 9.1	04/28/2008
Junos 9.0	02/15/2008
Junos 8.5 ¹	11/16/2007

You can directly upgrade or downgrade between any two Junos OS releases that are within three releases of each other.

- Example: Direct release upgrade

Release 10.3 → (*bypassing Releases 10.4 and 11.1*) Release 11.2

To upgrade or downgrade between Junos OS releases that are more than three releases apart, you can upgrade or downgrade first to an intermediate release that is within three releases of the desired release, and then upgrade or downgrade from that release to the desired release.

- Example: Multistep release downgrade

Release 11.3 → (*bypassing Releases 11.2 and 11.1*) Release 10.4 → Release 10.3

Juniper Networks has also provided an even more efficient method of upgrading and downgrading using the Junos OS EEOL releases. EEOL releases generally occur once a calendar year and can be more than three releases apart. For a list of, EEOL releases, go to <http://www.juniper.net/support/eol/junos.html>

You can directly upgrade or downgrade between any two Junos OS EEOL releases that are within three EEOL releases of each other.

- Example: Direct EEOL release upgrade

Release 9.3 (EEOL) → (*bypassing Releases 10.0 [EEOL] and 10.4 [EEOL]*) Release 11.4 (EEOL)

To upgrade or downgrade between Junos OS EEOL releases that are more than three EEOL releases apart, you can upgrade first to an intermediate EEOL release that is within three EEOL releases of the desired EEOL release, and then upgrade from that EEOL release to the desired EEOL release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 8.5 (EEOL) → (*bypassing Releases 9.3 [EEOL] and 10.0 [EEOL]*) Release 10.4 (EEOL) → Release 11.4 (EEOL)

You can even use a Junos OS EEOL release as an intermediate upgrade or downgrade step if your desired release is several releases later than your current release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 9.6 → Release 10.0 (EEOL) → Release 10.2

For additional information about how to upgrade and downgrade, see the *Junos OS Installation and Upgrade Guide*.

Upgrading an AppSecure Device

Use the **no-validate** Option for AppSecure Devices.

For devices implementing AppSecure services, use the no-validate option when upgrading from Junos OS Release 11.2 or earlier to Junos OS 11.4R1 or later. The application signature package used with AppSecure services in previous releases has been moved from the configuration file to a signature database. This change in location can trigger an error during the validation step and interrupt the Junos OS upgrade. The no-validate option bypasses this step.

Upgrade and Downgrade Scripts for Address Book Configuration

Beginning with Junos OS Release 12.1, you can configure address books under the **[security]** hierarchy and attach security zones to them (zone-attached configuration). In Junos OS Release 11.1 and earlier, address books were defined under the **[security zones]** hierarchy (zone-defined configuration).

You can either define all address books under the **[security]** hierarchy in a zone-attached configuration format or under the **[security zones]** hierarchy in a zone-defined configuration format; the CLI displays an error and fails to commit the configuration if you configure both configuration formats on one system.

Juniper Networks provides Junos operation scripts that allow you to work in either of the address book configuration formats (see [Figure 1 on page 150](#)).

- [About Upgrade and Downgrade Scripts on page 149](#)
- [Running Upgrade and Downgrade Scripts on page 150](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 151](#)

About Upgrade and Downgrade Scripts

After downloading Junos OS Release 12.1, you have the following options for configuring the address book feature:

- **Use the default address book configuration**—You can configure address books using the zone-defined configuration format, which is available by default. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.
- **Use the upgrade script**—You can run the upgrade script available on the Juniper Networks support site to configure address books using the new zone-attached configuration format. When upgrading, the system uses the zone names to create address books. For example, addresses in the trust zone are created in an address book named **trust-address-book** and are attached to the trust zone. IP prefixes used in NAT rules remain unaffected.

After upgrading to the zone-attached address book configuration:

- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.


For information on how to configure zone-attached address books, see the Junos OS Release 12.1 documentation.

- **Use the downgrade script**—After upgrading to the zone-attached configuration, if you want to revert to the zone-defined configuration, use the downgrade script available on the Juniper Networks support site. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.



NOTE: Before running the downgrade script, make sure to revert any configuration that uses addresses from the global address book.

Figure 1: Upgrade and Downgrade Scripts for Address Books



Running Upgrade and Downgrade Scripts

The following restrictions apply to the address book upgrade and downgrade scripts:

- The scripts cannot run unless the configuration on your system has been committed. Thus, if the zone-defined address book and zone-attached address book configurations are present on your system at the same time, the scripts will not run.
- The scripts cannot run when the global address book exists on your system.
- If you upgrade your device to Junos OS Release 12.1 and configure logical systems, the master logical system retains any previously-configured zone-defined address book configuration. The master administrator can run the address book upgrade script to convert the existing zone-defined configuration to the zone-attached configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.



NOTE: You cannot run the downgrade script on logical systems.

For information about implementing and executing Junos operation scripts, see the *Junos OS Configuration and Operations Automation 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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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 <http://www.juniper.net/support/eol/junos.html>.

Hardware Requirements for Junos OS Release 12.1 for SRX Series Services Gateways and J Series Services Routers

Transceiver Compatibility for SRX Series and J Series Devices

We strongly recommend that only transceivers provided by Juniper Networks be used on SRX Series and J Series interface modules. Different transceiver types (long-range, short-range, copper, and others) can be used together on multiport SFP interface modules as long as they are provided by Juniper Networks. We cannot guarantee that the interface module will operate correctly if third-party transceivers are used.

Please contact Juniper Networks for the correct transceiver part number for your device.

Power and Heat Dissipation Requirements for J Series PIMs

On J Series Services Routers, the system monitors the PIMs and verifies that the PIMs fall within the power and heat dissipation capacity of the chassis. If power management is enabled and the capacity is exceeded, the system prevents one or more of the PIMs from becoming active.



CAUTION: Disabling the power management can result in hardware damage if you overload the chassis capacities.

You can also use CLI commands to choose which PIMs are disabled. For details about calculating the power and heat dissipation capacity of each PIM and for troubleshooting procedures, see the *J Series Services Routers Hardware Guide*.

Supported Third-Party Hardware

The following third-party hardware is supported for use with J Series Services Routers running Junos OS.

- **USB Modem**

We recommend using a U.S. Robotics USB 56K V.92 Modem, model number USR 5637.

- **Storage Devices**

The USB slots on J Series Services Routers accept a USB storage device or USB storage device adapter with a CompactFlash card installed, as defined in the *CompactFlash Specification* published by the CompactFlash Association. When the USB device is installed and configured, it automatically acts as a secondary boot device if the primary CompactFlash card fails on startup. Depending on the size of the USB storage device, you can also configure it to receive any core files generated during a router failure. The USB device must have a storage capacity of at least 256 MB.

[Table 11 on page 152](#) lists the USB and CompactFlash card devices supported for use with the J Series Services Routers.

Table 11: Supported Storage Devices on the J Series Services Routers

Manufacturer	Storage Capacity	Third-Party Part Number
SanDisk—Cruzer Mini 2.0	256 MB	SDCZ2-256-A10
SanDisk	512 MB	SDCZ3-512-A10
SanDisk	1024 MB	SDCZ7-1024-A10
Kingston	512 MB	DTI/512KR
Kingston	1024 MB	DTI/1GBKR
SanDisk—ImageMate USB 2.0 Reader/Writer for CompactFlash Type I and II	N/A	SDDR-91-A15
SanDisk CompactFlash	512 MB	SDCFB-512-455
SanDisk CompactFlash	1 GB	SDCFB-1000.A10

J Series CompactFlash and Memory Requirements

[Table 12 on page 153](#) lists the CompactFlash card and DRAM requirements for J Series Services Routers.

Table 12: J Series CompactFlash Card and DRAM Requirements

Model	Minimum CompactFlash Card Required	Minimum DRAM Required	Maximum DRAM Supported
J2320	1 GB	1 GB	2 GB
J2350	1 GB	1 GB	2 GB
J4350	1 GB	1 GB	2 GB
J6350	1 GB	1 GB	2 GB

Related Documentation

- [New Features in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 63](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 84](#)
- [Known Limitations in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 89](#)
- [Outstanding Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 110](#)
- [Resolved Issues in Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 111](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for Branch SRX Series Services Gateways and J Series Services Routers on page 135](#)

Junos OS Release Notes for High-End SRX Series Services Gateways

Powered by Junos OS, Juniper Networks high-end SRX Series Services Gateways provide robust networking and security services. High-end SRX Series Services Gateways are designed to secure enterprise infrastructure, data centers, and server farms. The high-end SRX Series Services Gateways include the SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

- [Release Notes Updates on page 154](#)
- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 168](#)
- [Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 174](#)
- [Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 187](#)
- [Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 188](#)

- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways](#) on page 210
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for High-End SRX Series Services Gateways](#) on page 217

Release Notes Updates

The technical documentation for Junos OS Release 12.1R1 for SRX Series and J Series devices has been rereleased in the Juniper Networks standard format. This new tabbed webpage format simplifies browsing by enabling you to navigate to content based on the job you need to do: Overview, Configuration, Administration, or Troubleshooting.

- Access to related information and to PDF versions of each document are available from each tabbed page.
- For Junos OS Release 12.1R1, “legacy” PDFs are also provided for those customers who have already bookmarked the original Junos OS Release 12.1R1 documentation.
- The Junos OS Release Notes for SRX Series and J Series devices will continue to be available in both HTML and PDF format.
- See [Junos OS for SRX Series: Release 12.1](#)

New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways

The following features have been added to Junos OS Release 12.1. Following the description is the title of the manual or manuals to consult for further information.



NOTE: For the latest updates about support and issues on Junos Pulse, see the [Junos Pulse Release Notes](#).

Release 12.1R6 New Features

Chassis Grounding for SRX1400 Through SRX5800 Services Gateways




WARNING:

In order to meet safety and electromagnetic interference (EMI) requirements and to ensure proper operation, you must properly ground the services gateway chassis before connecting power. This requirement applies to the following services gateway models without exception:

- SRX1400 Services Gateway
 - SRX3400 Services Gateway
 - SRX3600 Services Gateway
 - SRX5600 Services Gateway
 - SRX5800 Services Gateway
-

For all services gateway models, the accessory box shipped with the device includes one cable lug that attaches the grounding cable to the services gateway chassis. The cable lug is shown in [Figure 2 on page 155](#).

Figure 2: Grounding Cable Lug



Before services gateway installation begins, a licensed electrician must attach the cable lug to the grounding cable that you supply. A cable with an incorrectly attached lug can damage the services gateway. The grounding cable must be no smaller than specified in [Table 13 on page 155](#), or as required by local electrical codes:

Table 13: Grounding Cable Wire Specification

Services Gateway Type	Grounding Cable Wire Specification
SRX1400 Services Gateway	14-AWG (2.1 mm ²), minimum 60°C wire
SX3400 Services Gateway	10-AWG (5.3 mm ²), minimum 60°C wire
SRX3600 Services Gateway	10-AWG (5.3 mm ²), minimum 60°C wire
SRX5600 Services Gateway	6-AWG (13.3 mm ²), minimum 60°C wire
SRX5800 Services Gateway	6-AWG (13.3 mm ²), minimum 60°C wire



NOTE: For the SRX5800 services gateway models, we previously specified 10-AWG wire for the grounding cable. Where you have installed such grounding cables, you can safely leave them in service. However, all new installations of SRX5800 Services Gateways must have grounding cables sized according to [Table 13 on page 155](#).

If you have lost the grounding cable lug supplied with the services gateway, contact your Juniper Networks representative to obtain a replacement.

[Figure 3 on page 156](#) through [Figure 7 on page 158](#) show the locations of the chassis grounding points on the listed SRX Series Services Gateway models. We recommend that you confirm that your services gateway chassis is properly grounded as soon as practical. For full instructions on grounding the services gateway chassis, see the hardware documentation for your services gateway.

Figure 3: Connecting the Grounding Cable, SRX1400 Services Gateway




Figure 4: Connecting the Grounding Cable, SRX3400 Services Gateway




Figure 5: Connecting the Grounding Cable, SRX3600 Services Gateway




Figure 6: Connecting the Grounding Cable, SRX5600 Services Gateway





Figure 7: Connecting the Grounding Cable, SRX5800 Services Gateway



g030295

Release 12.1R3 New Features

Network Address Translation (NAT)

- **Network Address Translation support for port mapping**—This feature is supported on all high-end SRX Series devices.

Network Address Translation (NAT) now provides **destination-port low to high** and **mapped-port low to high** statements to allow static NAT to map ports as follows:

- To map multiple IP addresses to the same IP addresses on a specified range of ports
- To map a specific IP address and port to a different IP address and port

Examples of configuring static NAT with port mapping:

```
set security nat static rule-set rs from zone trust rule r1 match destination-address 1.1.1.1/32
set security nat static rule-set rs from zone trust rule r1 destination-port 100 to 200
set security nat static rule-set rs from zone trust rule r1 then static-nat prefix 10.1.1.1/32
set security nat static rule-set rs from zone trust rule r1 then static-nat prefix mapped-port 300 to 400
set security nat static rule-set rs from zone trust rule r2 match destination-address 1.1.1.1/32
set security nat static rule-set rs from zone trust rule r2 destination-port 300 to 400
set security nat static rule-set rs from zone trust rule r2 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs from zone trust rule r2 then static-nat prefix mapped-port 300 to 400
set security nat static rule-set rs from zone trust rule r3 match destination-address 1.1.1.3/32
set security nat static rule-set rs from zone trust rule r3 destination-port 300
set security nat static rule-set rs from zone trust rule r3 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs from zone trust rule r3 then static-nat prefix mapped-port 200
```

Release 12.1R2 New Features

Security Policies

- New **show services unified-access-control counters** command—This feature is supported on all high-end SRX Series devices.

The **show services unified-access-control counters** command displays the number of sessions allowed, denied, and terminated by the Unified Access Control (UAC) service when invoked by a firewall policy that specifies the uac-policy action. Counts are reported for each action taken by UAC.

On high-end SRX Series devices, UAC counts are displayed for each PIC on the device. Sessions that were allowed, denied, or terminated by some other firewall policy action are not included in these statistics.

```
user@host> show services unified-access-control counters
PIC: fpc2.pic0
Sessions allowed
Policy action: 0
Timeout action: 0
Sessions denied
Unauthenticated: 0
Policy action: 0
Policy not matched: 0
Timeout action: 0
Sessions terminated
Reevaluation: 0
Signout: 0
```

Software Features

AppSecure

- **AppQoS information added to session output**—This feature is supported on all high-end SRX Series devices.

A new **extensive** option has been added to the **show security flow session** command. This option expands the dynamic application output to include any AppQoS traffic control rules applied to an active session.

[*Junos OS CLI Reference*]

- **Synchronizing application identification package in a chassis cluster**—This feature is supported on all high-end SRX Series devices.

The existing application identification package download feature is modified to download the security package on the primary node and synchronize it on the secondary node.

When you download the application signature package on a device operating in chassis cluster mode, this feature enables you to download and install the package to the primary node, after which the primary and secondary nodes are automatically synchronized. This synchronization helps maintain the same version of the security package on both the primary node and the secondary node.

You can use the command **request services application-identification download** to download the application identification package to the primary node and use the command **request services application-identification download status** to verify the download status of the application package and the status of the synchronization to the secondary node.

[*Junos OS Security Configuration Guide*]

Chassis Cluster

- **Logical interfaces on redundant Ethernet interface scaling**—This feature is supported on all high-end SRX Series devices.

The total number of logical interfaces that you can configure across all the redundant Ethernet (reth) interfaces in a chassis cluster deployment has been increased to 4096.

[*Junos OS Security Configuration Guide*]

Class of Service

- **High-priority queue on SPC**—This feature is supported on all high-end SRX Series devices.

Junos OS Release 12.1 provides a configuration option to enable packets with specific Differentiated Services (DiffServ) code points (DSCP) precedence bits to enter a high-priority queue on the Services Processing Card (SPC) on high-end SRX Series devices. The Services Processing Unit (SPU) draws packets from the high-priority queue and only draws packets from the low-priority queue when the high-priority queue is empty. This feature can reduce overall latency for real-time traffic, such as voice traffic.

To designate packets for the high-priority or low-priority queues, use the **spu-priority** configuration statement at the [edit **class-of-service forwarding-classes class**] hierarchy level. A value of **high** places packets into the high-priority queue, and a value of **low** places packets into the low-priority queue.

[Junos OS Class of Service Configuration Guide for Security Devices]

Denial of Service (DoS)

- **Whitelists**—This feature is supported on all high-end SRX Series devices.

You can configure a whitelist of IP addresses that are to be exempt from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process. To do so, use the new **white-list** statement at the [edit **security screen ids-option screen-name tcp syn-flood**] hierarchy level.

Both IP version 4 (IPv4) and IP version 6 (IPv6) whitelists are supported. Addresses in a whitelist must be all IPv4 or all IPv6. Each whitelist can have up to 32 IP address prefixes.

[Junos OS CLI Reference, Junos OS Security Configuration Guide]

General Packet Radio Service (GPRS)

- **Customized IE removal**—This feature is supported on all high-end SRX Series devices.

You can now configure the GPRS tunneling protocol (GTP) firewall to remove specific information elements (IE) using the user-configured IE number. When you configure the IE-removal, the GTP firewall deletes the corresponding IEs of the GTPv1 messages; updates the length of the GTP, the UDP, and the IP; and then passes the GTPv1 message. The GTP firewall also updates the cyclic redundancy check (CRC) code.

[Junos OS CLI Reference, Junos OS Security Configuration Guide]

Intrusion Detection and Prevention (IDP)

- **Synchronizing IDP security package in a chassis cluster**—This feature is supported on all high-end SRX Series devices.

The existing IDP security package download feature has been modified to download the security package on the primary node and synchronize it on secondary node.

When you download the IDP security package on a device operating in chassis cluster mode, this feature enables you to download and install the IDP signature database to the primary node, after which the primary and secondary nodes are automatically synchronized. This synchronization helps maintain the same version of the security package on both the primary node and the secondary node.

You can use the **request security idp security-package download** command to download the IDP signature database to the primary node.

When you check the security package download status using the **request security idp security-package download status** command, a message is displayed confirming that the downloaded security package is being synchronized on the primary and secondary nodes.

[Junos OS Security Configuration Guide]

J-Web

- **J-Web support for IPv6 stage 3 security features** —This feature is supported on all high-end SRX Series devices.

The following J-Web pages have been added to support the IPv6 feature:

- Proxy ND Configuration page
- DS-Lite Configuration page

The following J-Web pages have been updated to include IPv6 support:

- Source NAT Pool Configuration page
- Static NAT Rule Configuration page
- Source NAT Rule Configuration page
- Destination NAT Pool Configuration page
- Destination NAT Rule Configuration page

[Junos OS Security Configuration Guide]

- **J-Web support for services offloading**—This feature is supported on all high-end SRX Series devices.

Services offloading is a mechanism for processing fast-path packets in the network processor instead of in the Services Processing Unit (SPU). This method reduces the long packet processing latency that arises when packets are forwarded from network processors to SPUs for processing and back to I/O cards (IOCs) for transmission.

This feature can now be enabled or disabled on a chassis through the J-Web user interface. You can use the Chassis configuration page in J-Web to enable or disable services offloading on a chassis.

[*Junos OS Security Configuration Guide*]

Logical Systems

- **IPv6 addresses in logical systems**—This feature is supported on all high-end SRX Series devices.

In Junos OS Release 12.1, IPv6 addresses can be configured in logical systems for the following features:

- Firewall authentication
- BGP routing

[*Junos OS Logical Systems Configuration Guide for Security Devices*]

- **IPv6 dual-stack lite (DS-Lite)**—This feature is supported on all high-end SRX Series devices.

DS-Lite allows migration to an IPv6 access network without changing end user software. IPv4 users can continue to access IPv4 Internet content using their current hardware, while IPv6 users are able to access IPv6 content. A DS-Lite software initiator at the customer edge encapsulates IPv4 packets into IPv6 packets while a software concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4 NAT translations.

The master administrator allocates the number of software initiators that can be configured in a logical system. The master administrator configures maximum and reserved values for the **dslite-software-initiator** configuration statement at the [**edit system security-profile *profile-name***] hierarchy level. The master administrator or user logical system administrator can configure software concentrators at the [**edit security softwires**] hierarchy level.

[*Junos OS Logical Systems Configuration Guide for Security Devices*]

Routing

- **Equal-cost multipath (ECMP) flow-based forwarding**—This feature is supported on SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

An equal-cost multipath (ECMP) set is formed when the routing table contains multiple next-hop addresses for the same destination with equal cost. (Routes of equal cost have the same preference and metric values.) If there is an ECMP set for the active route, the Junos OS software uses a hash algorithm to choose *one* of the next-hop addresses in the ECMP set to install in the forwarding table.

You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. On Juniper Networks devices, per-flow load balancing can be performed to spread traffic across multiple paths between routing devices. On Juniper Networks security devices, source and destination IP addresses and protocols are examined to determine individual traffic flows. Packets for the same flow are forwarded on the same interface; the interface does not change when there are additions or changes to the ECMP set. This is important for features such as source NAT, where the translation is performed only during the first path of session establishment; IDP; ALG; and route-based VPN tunnels. If a packet arrives on a given interface in an ECMP set, the security device ensures that reverse traffic is forwarded through the same interface.



.....

NOTE: ECMP flow-based forwarding on security devices applies only to IPv4 unicast traffic flows. Multicast and IPv6 flows are not supported.

.....

On Juniper Networks security devices, the maximum number of next-hop addresses in an ECMP set that can be installed in the forwarding table is 16. If there are more than 16 next-hop addresses in an ECMP set, only the first 16 addresses are used.

In a chassis cluster deployment, a *local* interface is an interface that is on the same node as the interface on which a packet arrives, and a *remote* interface is an interface that is on the other chassis cluster node. If an ECMP route has both local and remote interfaces in a chassis cluster, then the local interface is favored for the next hop.

To configure ECMP flow-based forwarding on Juniper Networks security devices, first define a load-balancing routing policy by including one or more **policy-statement** configuration statements at the [**edit policy-options**] hierarchy level, with the action **load-balance per-packet**. Then apply the routing policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** configuration statements at the [**edit routing-options**] hierarchy level.

[*Junos OS Routing Protocols and Policies Configuration Guide for Security Devices*]

Security Policies

- **Binary format available for security log messages**—This feature is supported on all high-end SRX Series devices. The feature requires a minimum of 1 gigabyte of memory.

Binary format for security log messages is available on all SRX Series devices. Because security logging can produce large amounts of data, text-based log files can quickly

consume storage and CPU resources. Depending on your implementation of security logging, a log file in a binary-based format can provide more efficient use of on- or off-device storage and improved CPU utilization.

When configured in event mode, security log messages stored in binary format are maintained on the device in a separate log file from system log messages. The **show security log file** operational command supports decoding, converting, and viewing binary log files stored locally on the device. (Events stored in binary format are not accessible with advanced log scripting commands intended for text-based formats.)

When configured in stream mode, security messages generated in the data plane in binary format are streamed directly to external log collection clients in a Juniper Networks-specific binary format. The third-party client handles decoding, converting, and viewing binary log files stored on the remote device.

A programmatic interface provides the tools for developing third-party collection, decoding, conversion, and viewing applications. For information about the programmatic interface, contact your Juniper Networks representative.

[Junos OS Monitoring and Troubleshooting Guide for Security Devices]

- **Hit-count tracking**—This feature is supported on all high-end SRX Series devices.

The new **show security policies hit-count** command displays the utility rate of security policies according to the number of hits they receive. You can use this feature to determine which policies are being used on the device, and how frequently they are used. Depending on the command options that you choose, the number of hits can be listed without order or sorted in either ascending or descending order, and they can be restricted to the number of hits that fall above or below a specific count or within a range. Data is shown for all zones associated with the policies or named zones.

[Junos OS CLI Reference, Junos OS Security Configuration Guide]

SSL Proxy

- **SSL proxy**—This feature is supported on all high-end SRX Series devices.

SSL proxy is a transparent proxy; that is, it performs SSL encryption and decryption between the client and the server, but neither the server nor the client can detect its presence. SSL offload and SSL inspection features require the servers to share their secret keys to be able to decrypt the SSL traffic. However, sharing server keys is sometimes not feasible or might not be available in certain circumstances, in which case, the SSL traffic cannot be decrypted.

SSL proxy addresses this problem by ensuring that it has the keys to encrypt and decrypt the payload. Decryption and encryption take place in each direction (client and server), and the keys are different for both encryption and decryption.

- **Application firewall, IDP, and application tracking with SSL proxy**—With the implementation of SSL proxy, application ID can identify applications encrypted in SSL. SSL proxy can be enabled as an application service in a regular firewall policy rule. IDP, application firewall, and application tracking services can use the decrypted content from SSL proxy. On the SSL payload, IDP can inspect attacks and anomalies; for example, HTTP chunk length overflow on HTTPS. On encrypted applications, such

as Facebook, application firewall can enforce policies and application tracking (when configured in the from and to zones) can report issues, through the logs, based on dynamic application and nested application.



NOTE: If none of the services (application firewall, IDP, or application tracking) are configured, then SSL proxy services are bypassed even if an SSL proxy profile is attached to a firewall policy.

UAC Authentication

- **Security enhancement between a Junos Pulse Access Control Service and an Infranet Enforcer**—This feature is supported on all high-end SRX Series devices with UAC enabled.

To improve security for a UAC configuration, a message appears if the Junos Pulse Access Control Service's certificate cannot be verified on the SRX Series device. To ensure that the Access Control Service and the Enforcer are connected securely, configure the ca-profile on the SRX Series device to verify the Access Control Service's certificate.

[*Junos OS Security Configuration Guide*]

User Role Firewalls

- **Automating the Single Sign-On for User Role FW**—This feature is supported on all high-end SRX Series devices.

To allow users to silently authenticate and access network resources over any protocol (not just HTTP), a hidden browser can be launched to start SPNEGO.

When a user logs on to the AD domain, the authentication server (MS Windows Domain Controller) pushes a log on script to the endpoint. When invoked at the endpoint, the script opens a browser and navigates to the URL of the Junos Pulse Access Control Service.

The following sample Jscript file defines a variable for the URL of the Access Control Service, defines the navigation function, and creates a procedure to open a browser and navigate to the specified URL.



NOTE: This is only a sample script. Your implementation will vary depending on your configuration, browser choice, administrative options, and scripting language.

```
//Variable definitions
var done = false;
var url = WScript.Arguments.Count() ? WScript.Arguments.Item(0) :
"www.your-acs.com" ;

//Event Handlers
function objIE_NavigateComplete2(pdisp, URL) {
    done = true;
}
```

```
//Main section
var objIE = WScript.CreateObject("InternetExplorer.Application", "objIE_");
objIE.Visible = false;
objIE.Navigate(url);

while (!done) {
    WScript.Sleep(1000);
}
```

Save the file using the file extension for the scripting language, and copy it to the Sysvol\Scripts location on the Active Directory server.

Add the file in the Logon Properties window in the Group Policy Operations path. See your Microsoft Active Directory administrators guide for further information.

[*Junos Pulse Administrator Guide*]

- **User role firewall providing flexibility and higher security**—This feature is supported on all high-end SRX Series devices.

Network security enforcement, monitoring, and reporting based solely on IP information soon will not be sufficient for today's dynamic and mobile workforce. By integrating user firewall policies, administrators can permit or restrict network access of employees, contractors, partners, and other users based on the roles they are assigned.

A new match criteria, source-identity, defines applicable roles for each policy. In this way, traffic can be permitted or denied access based on the role of the user, as well as the zone pair, source and destination IP addresses, and application.

The SRX Series device can be configured to:

- Work with existing Junos Pulse Access Control Service features.
- Add direct support for agentless Active Directory domain controller authentication, and use MS single sign-on protocols for dynamic role processing.
- Implement local database support.

[*Junos OS Security Configuration Guide, Unified Access Control Solution Guide for SRX Series Services Gateways*]

- **User role integration into AppTrack logs**—This feature is supported on all high-end SRX Series devices.

User identity details such as user name and user role are added to the AppTrack session create, session close, and volume update logs. This information is obtained from the Unified Access Control (UAC) connection or from the local authentication table. These fields will contain the user name and role associated with the policy match.

You can enable the logging of user name and roles for security policies that provide UAC enforcement or with an entry in the local authentication table. For security policies without Userfw enforcement, the user name and the user role field in the log is displayed as N/A.

The user name is displayed as unauthenticated-user and user role is displayed as N/A, if the device cannot retrieve information for that session because there is no authentication table entry for that session or because logging of this information is

disabled. The user role field in the log will contain the list of all the roles performed by the user if match criteria is specific, authenticated-user, or any and the user name field in the log contains the correct user name. The user role field in the log will contain N/A if the match criteria and the user name field in the log contains unauthenticated user or unknown user.

[*Junos OS Security Configuration Guide*]

**Related
Documentation**

- [Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 187](#)
- [Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 188](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 210](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 168](#)
- [Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 174](#)

Changes in Default Behavior and Syntax in Junos OS Release 12.1 for High-End SRX Series Services Gateways

The following current system behavior, configuration statement usage, and operational mode command usage might not yet be documented in the Junos OS documentation:

AppSecure Application Package Upgrade Changes

- On all high-end SRX Series devices, application tracking is enabled by default. You can disable application tracking with the **set security application-tracking disable** command. This command allows you to disable and re-enable application tracking without modifying your existing zone selections.
- **Application signatures removed after upgrading to Junos OS Release 11.4**—This change applies to all high-end SRX Series devices that use the application identification signature package.

As of Junos OS Release 11.4, the application signature package is downloaded and installed in a separate database, not in the Junos OS configuration file as in previous Junos OS releases.

When you upgrade an SRX Series device from Junos OS Release 11.2 to Junos OS Release 11.4 or later, any predefined application signatures and signature groups from the Junos OS Release 11.2 configuration will be removed when you install the latest predefined signatures and signature groups by using the **request services application-identification install** command. However, the upgrade will not remove custom signatures and signature groups from the Junos OS configuration.

For information about using the **request services application-identification download** and **request services application-identification install** commands, see the *Junos OS CLI Reference*.

Chassis Cluster

- On all SRX Series devices, in Junos OS Release 12.1R5 and earlier releases, when you configured some interfaces using the **interface-range** command, if you later requested the committed configuration using the **<get-configuration>** RPC with the **inherit="inherit"** and **database="committed"** attributes, the device returned **junos:changed-localtime** and **junos:changed-seconds** in the RPC reply instead of **junos:commit-localtime** and **junos:commit-seconds**.
- On all high-end SRX Series devices, the **set chassis fpc <slot-number> pic <slot-number> services-offload** command is not available for configuration and, as a consequence, the devices cannot be used in services-offload mode.

Command-Line Interface (CLI)

- The range in the **maximum-call-duration** command has been changed from 7200 minutes to 720 minutes.
- On all SRX Series devices, the **show security flow session extensive** command has been updated to show the predefined application name.
- On all SRX Series devices, the **reauthentication** command has been introduced to specify the reauthentication interval in seconds.
- On all high-end SRX Series and J Series devices, the **show security gprs gtp counters** command has been updated to show the path-rate-limit counters.
- On all SRX Series devices, the **no-recommended** command has been introduced to specify non recommended attack objects in the dynamic attack group.
- The **client-match match-name** option under security hierarchy [**edit security policies from-zone zone-name to-zone zone-name policy policy-name then permit firewall-authentication**] now supports a maximum of 64 users or user groups in the policy.
- On all high-end SRX Series devices, the **show interface interface-name statistics detail** command was showing incorrect FCS statistics. Additional 4 bytes in the FCS were counted in input statistics but not counted in output statistics. Now the FCS is included in both input and output Ethernet statistics and the **show interface interface-name statistics detail** command displays correct output.
- On all high-end SRX Series devices, a new command **clear security flow statistics**, has been introduced to clear the flow-related system statistics.

Deprecated Items for High-End SRX Series Services Gateways

[Table 14 on page 170](#) lists deprecated items (such as CLI statements, commands, options, and interfaces).

CLI statements and commands are deprecated—rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration. We strongly recommend that you phase out deprecated items and replace them with supported alternatives.

Table 14: Items Deprecated in Release 12.1

Deprecated Item	Replacement	Hierarchy Level or Command Syntax	Additional Information
<code>download-timeout</code>	-	<code>download-timeout timeout</code>	On all high-end SRX Series devices, the download-timeout command is deprecated. If the configuration is present, it will be ignored. The <code>idpd</code> daemon internally triggers the security package to install when an automatic download is completed. There is no need to configure any download timeout.
<code>node</code>	-	<code>request security idp security-package download</code>	On all high-end SRX Series devices operating in a chassis cluster, the following request security idp security-package download commands with the <code>node</code> option is not supported: <code>request security idp security-package download node primary</code> <code>request security idp security-package download node local</code> <code>request security idp security-package download node all</code>

Flow and Processing

- On SRX Series devices, TCP session configuration has been enhanced. In previous releases, the TCP session was only invalidated after 2 seconds of receiving a FIN or ACK packet. As a result, if there were new incoming SYN packets (with the same five tuples) within this 2 seconds, the packets would use the same TCP session, which would be terminated within 2 seconds. Therefore a new session could not be established.

A new configuration attribute, **fin-invalidate-session**, has been introduced at the **[edit security flow tcp-session]** hierarchy level. This command invalidates a TCP session immediately on reception of a FIN or ACK packet. New incoming SYN packets will need to establish a new session.

To end a session immediately on reception of a FIN or ACK packet, use the **set security flow tcp-session fin-invalidate-session** command:



NOTE: The in-service software upgrade (ISSU) procedure is not supported while upgrading from a Junos OS release supporting this command to a Junos OS release that does not support it.

ISSU is supported in the following scenarios:

- While upgrading from a Junos OS release that does not support the command to a Junos OS release that supports the command.
 - While upgrading from a Junos OS release that supports the command to a Junos OS release that also supports the command.
 - If the `fin-invalidate-session` option is not configured.
-
- **TCP Initial Timeout Enhancement**—The minimum value you can configure for TCP session initialization is 4 seconds. The default value is 20 seconds; if required you can set the TCP session initialization value to less than 20 seconds.

In-Service Software Upgrade (ISSU)

On all high-end SRX Series devices, at the beginning of a chassis cluster ISSU, the system automatically fails over all RG-1+ redundancy groups that are not primary on the node from which you start the ISSU. This action ensures that the redundancy groups are all active on only the RG-0 primary node. You no longer need to fail over redundancy groups manually.

After the system fails over all RG-1+ redundancy groups, it sets the manual failover bit and changes all RG-1+ primary node priorities to 255, regardless of whether the redundancy group failed over to the RG-0 primary node.

J-Web

- On all high-end SRX Series devices, the username field does not accept HTML tags or the < and > characters. The following error message appears: A username cannot include certain characters, including < and >.

Logical Systems

- The **logical-systems all** option can now be specified for the **show security screen statistics** operational command.

Management Information Base (MIB)

- On all high-end SRX Series devices, in a chassis cluster environment, the calculation of the primary and secondary node sessions in the **JnxJsSPUMonitoringObjectsTable** object of the SPU monitoring MIB is incorrect because the MIB **jnxJsSPUMonitoringCurrentTotalSession** incorrectly displays total sessions. A doubled session count is displayed because the active and backup nodes are treated as separate sessions, although they are not.

Count only the session numbers on the local node, thereby avoiding a double count, and local total sessions are displayed.

In a chassis cluster environment, the **SPUMonitoringCurrentTotalSession** object of the MIB adds information per each SPU from the local node.

[MIB Reference for SRX1400, SRX3400, and SRX3600 Services Gateways; MIB Reference for SRX5600 and SRX5800 Services Gateways]

Remote ID Check for Site-to-Site VPNs

- As of Junos OS Release 11.4, checks are performed to validate the IKE ID received from the VPN peer device. By default, SRX Series and J Series devices validate the IKE ID received from the peer with the IP address configured for the IKE gateway. In certain network setups, the IKE ID received from the peer (which can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address) does not match the IKE gateway configured on the SRX Series or J Series device. This can lead to a Phase 1 validation failure.

To modify the configuration of the SRX Series or J Series device or the peer device for the IKE ID that is used:

- On the SRX Series or J Series device, configure the **remote-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level to match the IKE ID that is received from the peer. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.



NOTE: If you do not configure **remote-identity**, the device uses the IPv4 or IPv6 address that corresponds to the remote peer by default.

- On the peer device, ensure that the IKE ID is the same as the **remote-identity** configured on the SRX Series or J Series device. If the peer device is an SRX Series or J Series device, configure the **local-identity** statement at the **[edit security ike gateway gateway-name]** hierarchy level. Values can be an IPv4 or IPv6 address, fully qualified domain name, distinguished name, or e-mail address.

Security

- If a *policy* is configured with multiple applications, and more than one of the applications match the traffic, then the *application* that best meets the match criteria is selected.
- Security policies are stored in both the Routing Engine (RE), and the Packet Forwarding Engine (PFE). When you modify the policies on the Routing Engine side, the policies are synchronized to the Packet Forwarding Engine side when you commit the configuration.

The policies in the Routing Engine and Packet Forwarding Engine must always be in sync for the configuration to commit successfully. Under certain circumstances, policies in the Routing Engine and the Packet Forwarding Engine might be out of sync resulting in generation of system core files upon commit completion. As of Junos OS 11.2R6, Junos OS 11.4R3, when the policy configuration is modified and policies are out of sync, the following error message will be displayed when you attempt to commit a configuration:

```
Policy is out of sync between RE and PFE <SPU-name(s)>. Please resync before commit.
```

```
error: configuration check-out failed
```

To synchronize policies between the Routing Engine and the Packet Forwarding Engine, you must:

- Reboot the device (device in standalone mode)
- Reboot both devices (devices in a chassis cluster mode)
- Public key infrastructure (PKI) objects include certificates, key pairs, and certificate revocation lists (CRLs). PKI objects are read from the PKI database when the PKI Daemon starts. The PKI Daemon database loads all certificates into memory at boot time.

When an object is read into memory from the PKI database, the following new log message is created:

```
PKID_PV_OBJECT_READ: A PKI object was read into memory from <location>
```

Related Documentation

- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 187](#)
- [Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 188](#)

- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways](#) on page 210
- [Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways](#) on page 174

Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways

AppSecure

- J-Web pages for AppSecure are preliminary.
- Custom application signatures and custom nested application signatures are not currently supported by J-Web.
- When ALG is enabled, AppID includes the ALG result to identify the application of the control sessions. AppFW permits ALG data sessions whenever control sessions are permitted. If the control session is denied, there will be no data sessions.

When ALG is disabled, AppID relies on its signatures to identify the application of the control and data sessions. If a signature match is not found, the application is considered unknown. AppFW handles them based on the AppID result.

Chassis Cluster

- On all high-end SRX Series devices in a chassis cluster, only four QoS queues are supported per **reth/ae** interface.
- On all high-end SRX Series devices, the first recommended ISSU *from* release is Junos OS Release 10.4R4. If you intend to upgrade from a release earlier than Release 10.4R4, see the release notes for the release that you are upgrading from for information about limitations and issues related to upgrading.
- On all high-end SRX Series devices, ISSUs do not support the following features:
 - DHCP
 - GPRS/GTP/SCTP
 - GRE/IP-IP
 - J-Flow
 - Multicast routing
 - VPN

For the latest ISSU support status, go to the Juniper Networks Knowledge Base (KB): <http://kb.juniper.net/> and search for KB17946.

- In large chassis cluster configurations on SRX3400 or SRX3600 devices, you need to increase the wait time before triggering failover. In a full-capacity implementation, we recommend increasing the wait to 8 seconds by modifying **heartbeat-threshold** and **heartbeat-interval** values in the **[edit chassis cluster]** hierarchy.

The product of the **heartbeat-threshold** and **heartbeat-interval** values defines the time before failover. The default values (**heartbeat-threshold** of 3 beats and **heartbeat-interval** of 1000 milliseconds) produce a wait time of 3 seconds.

To change the wait time, modify the option values so that the product equals the desired setting. For example, setting the **heartbeat-threshold** to 8 and maintaining the default value for the **heartbeat-interval** (1000 milliseconds) yields a wait time of 8 seconds. Likewise, setting the **heartbeat-threshold** to 4 and the **heartbeat-interval** to 2000 milliseconds also yields a wait time of 8 seconds.

- Packet-based forwarding for MPLS and International Organization for Standardization (ISO) protocol families is not supported.
- On SRX5600 and SRX5800 devices, only two of the 10 ports on each PIC of 40-port 1-Gigabit Ethernet I/O cards (IOCs) can simultaneously enable IP address monitoring. Because there are four PICs per IOC, this permits a total of eight ports per IOC to be monitored. If more than two ports per PIC on 40-port 1-Gigabit Ethernet IOCs are configured for IP address monitoring, the commit will succeed but a log entry will be generated, and the accuracy and stability of IP address monitoring cannot be ensured. This limitation does not apply to any other IOCs or devices.
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, IP address monitoring is not supported on redundant Ethernet interface link aggregation groups (LAGs) or on child interfaces of redundant Ethernet interface LAGs.
- On all high-end SRX Series devices in chassis clusters, screen statistics data can be gathered on the primary device only.
- On all high-end SRX Series devices, ISSU does not support version downgrading.
- On all high-end SRX Series devices, only redundant Ethernet interfaces (*reth*) are supported for IKE external interface configuration in IPsec VPN. Other interface types can be configured, but IPsec VPN might not work.

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices DHCP is not supported in a chassis cluster.
- On all high-end SRX Series devices, DHCPv6 client authentication is not supported.

Firewall Filters

- On all SRX Series devices, firewall filters are not supported in Layer 2 transparent mode.

Flow and Processing

- On all high-end SRX Series devices, sessions going through GRE over IPsec are not marked as backup on backup node, instead, they are marked as active on backup node. This behavior causes the real active session to be deleted. Nesting tunnel, like GRE over IPsec, is not supported on chassis cluster.
- On all high-end SRX Series devices, when packet-logging functionality is configured with an improved pre-attack configuration parameter value, the resource usage increases proportionally and might affect the performance.

- On all high-end SRX Series devices, services offloading has the following limitations:
 - Transparent mode is not supported. If transparent mode is configured, a normal session is installed.
 - Link aggregation group (LAG) is not supported. If a LAG is configured, a normal session is installed.
 - Only multicast sessions with one fan-out are supported. If a multicast session with more than one fan-out exists, a normal session is installed.
 - Only active/passive chassis cluster (HA) configuration is supported. Active/active chassis cluster configuration is not supported.
 - Fragmented packets are not supported. If fragmented packets exist, a normal session is installed.
 - Ingress and egress interfaces on different network processors are not supported. If an ingress interface and the related egress interface do not belong to the same network processor, a normal session is installed on the network processor.
 - IP version 6 (IPv6) is not supported. If IPv6 is configured, a normal session is installed.



NOTE: A normal session forwards packets from the network processor to the Services Processing Unit (SPU) for fast-path processing, while a services-offload session processes fast-path packets in the network processor and the packets exit out of the network processor itself.

- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the default authentication table capacity is 45,000; the administrator can increase the capacity to a maximum of 50,000.

On SRX1400 devices, the default authentication table capacity is 10,000; the administrator can increase the capacity to a maximum of 15,000.
- On all high-end SRX Series devices, when devices are operating in flow mode, the Routing Engine side cannot detect the path maximum transmission unit (PMTU) of an IPv6 multicast address (with a large size packet).
- On all high-end SRX Series devices, you cannot configure route policies and route patterns in the same dial plan.
- On all high-end SRX Series devices, high CPU utilization triggered for reasons such as CPU intensive commands and SNMP walks causes the Bidirectional Forwarding Detection protocol (BFD) to flap while processing large BGP updates.
- On all high-end SRX Series devices, downgrading is not supported in low-impact ISSU chassis cluster upgrades (LICU).
- On SRX5800 devices, network processing bundling is not supported in Layer 2 transparent mode.

Hardware

This section covers filter and policing limitations.

- On SRX1400, SRX3400, and SRX3600 devices, the following feature is not supported by a simple filter:
 - Forwarding class as match condition
- On SRX1400, SRX3400 and SRX3600, devices, the following features are not supported by a policer or a three-color-policer:
 - Color-aware mode of a three-color-policer
 - Filter-specific policer
 - Forwarding class as action of a policer
 - Logical interface policer
 - Logical interface three-color policer
 - Logical interface bandwidth policer
 - Packet loss priority as action of a policer
 - Packet loss priority as action of a three-color-policer
- On all high-end SRX Series devices, the following features are not supported by a firewall filter:
 - Policer action
 - Egress filter-based forwarding (FBF)
 - Forwarding table filter (FTF)
- SRX3400 and SRX3600 devices have the following limitations of a simple filter:
 - Forwarding class as match condition
 - In the packet processor on an IOC, up to 400 logical interfaces can be applied with simple filters.
 - In the packet processor on an IOC, the maximum number of terms of all simple filters is 2000.
 - In the packet processor on an IOC, the maximum number of policers is 2000.
 - In the packet processor on an IOC, the maximum number of three-color-policers is 2000.
 - The maximum burst size of a policer or three-color-policer is 16 MB.

- On SRX3400 and SRX3600 devices, when you enable the monitor traffic option using the **monitor traffic** command to monitor the FXP interface traffic, interface bounce occurs. You must use the **monitor traffic interface fxp0 no-promiscuous** command to avoid the issue.

Interfaces and Routing

- On all branch SRX Series devices, IPv6 traffic transiting over IPv4 based IP over IP tunnel (for example, IPv6-over-IPv4 using ip-x/x/x interface) is not supported.
- On all high-end SRX Series devices, the **set protocols bgp family inet flow** and **set routing-options flow** CLI statements are no longer available, because BGP flow spec functionality is not supported on these devices.
- On all high-end SRX Series devices, the Link Aggregation Control Protocol (LACP) is not supported on Layer 2 interfaces.
- On all high-end SRX Series devices, BGP-based virtual private LAN service (VPLS) over aggregated Ethernet (**ae**) interfaces is not supported. It works on child ports and physical interfaces.

Internet Key Exchange Version 2 (IKEv2)

On all high-end SRX Series devices, IKEv2 does not include support for:

- Policy-based tunnels
- Dial-up tunnels
- Network Address Translation-Traversal (NAT-T)
- VPN monitoring
- Next-Hop Tunnel Binding (NHTP) for st0—Reusing the same tunnel interface for multiple tunnels
- Extensible Authentication Protocol (EAP)
- IPv6
- Multiple child SAs for the same traffic selectors for each QoS value
- Proposal enhancement features
- Reuse of Diffie-Hellman (DH) exponentials
- Configuration payloads
- IP Payload Compression Protocol (IPComp)
- Dynamic Endpoint (DEP)

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, from Junos OS Release 11.2 and later, the IDP security package is based on the Berkeley database. Hence, when the Junos OS image is upgraded from Junos OS Release 11.1 or earlier to Junos OS 11.2 or later, a migration of IDP security package files needs to be performed. This is done automatically on

upgrade when the IDP daemon comes up. Similarly, when the image is downgraded, a migration (secDb install) is automatically performed when the IDP daemon comes up, and previously installed database files get deleted.

However, migration is dependent on the XML files for the installed database to be present on the device. For first-time installation, full update files are required. If the last update on the device was an incremental update, migration might fail. In such a case, you have to manually download and install the IDP security package using the **download** or **install** CLI command before using the IDP configuration with predefined attacks or groups.

As a workaround, use the following CLI commands to manually download the individual components of the security package from the Juniper Security Engineering portal and install the full update:

- **request security idp security-package download full-update**
- **request security idp security-package install**
- On all high-end SRX Series devices, the IDP policies for each user logical system are compiled together and stored on the data plane memory. To estimate adequate data plane memory for a configuration, consider these two factors:
 - IDP policies applied to each user logical system are considered unique instances because the ID and zones for each user logical system are different. Estimates need to take into account the combined memory requirements for all user logical systems.
 - As the application database increases, compiled policies will require more memory. Memory usage should be kept below the available data plane memory to allow for database increases.
- On all high-end SRX Series devices, ingress as ge-0/0/2 and egress as ge-0/0/2.100 works with flow showing both source and destination interface as ge-0/0/2.100.
- IDP does not allow header checks for nonpacket contexts.
- On all high-end SRX Series devices, application-level distributed denial-of-service (application-level DDoS) detection does not work if two rules with different application-level DDoS applications process traffic going to a single destination application server. When setting up application-level DDoS rules, make sure that you do not configure rulebase-ddos rules that have two different application-ddos objects when the traffic destined to one application server can process more than one rule. Essentially, for each protected application server, you have to configure the application-level DDoS rules so that traffic destined for one protected server processes only one application-level DDoS rule.



NOTE: Application-level DDoS rules are terminal, which means that once traffic is processed by one rule, it will not be processed by other rules.

The following configuration options can be committed, but they will not work properly:

source-zone	destination-zone	destination-ip	service	application-ddos	Application Server
source-zone-1	dst-1	any	http	http-appddos1	1.1.1.1:80
source-zone-2	dst-1	any	http	http-appddos2	1.1.1.1:80

- On all high-end SRX Series devices, application-level DDoS rule base (rulebase-ddos) does not support port mapping. If you configure an application other than default, and if the application is from either predefined Junos OS applications or a custom application that maps an application service to a nonstandard port, application-level DDoS detection will not work.

When you configure the application setting as default, intrusion detection and prevention (IDP) uses application identification to detect applications running on standard and nonstandard ports; thus, the application-level DDoS detection would work properly.

- IDP deployed in both active/active and active/passive chassis clusters has the following limitations:
 - No inspection of sessions that fail over or fail back.
 - The IP action table is not synchronized across nodes.
 - The Routing Engine on the secondary node might not be able to reach networks that are reachable only through a Packet Forwarding Engine.
 - The SSL session ID cache is not synchronized across nodes. If an SSL session reuses a session ID and it happens to be processed on a node other than the one on which the session ID is cached, the SSL session cannot be decrypted and will be bypassed for IDP inspection.
- IDP deployed in active/active chassis clusters has a limitation that for time-binding scope source traffic, if attacks from a source (with more than one destination) have active sessions distributed across nodes, then the attack might not be detected because time-binding counting has a local-node-only view. Detecting this sort of attack requires an RTO synchronization of the time-binding state that is not currently supported.

IPv6 IPsec

IPv6 IPsec implementation has the following limitations:

- IPv6 routers do not perform fragmentation. IPv6 hosts should either perform path maximum transmission unit (PMTU) discovery or send packets smaller than the IPv6 minimum MTU size of 1280 bytes.
- Because IPv6 addresses are 128 bits long compared to IPv4 addresses, which are 32-bits long, IPv6 IPsec packet processing requires more resources. Therefore, a small performance degradation is observed.

- IPv6 uses more memory to set up the IPsec tunnel. Therefore, the IPsec IPv4 tunnel scalability numbers might drop.
- The addition of IPv6 capability might cause a drop in the IPsec IPv4-in-IPv4 tunnel throughput performance.
- The IPv6 IPsec VPN does not support the following functions:
 - 4in6 and 6in4 policy-based site-to-site VPN, IKE
 - 4in6 and 6in4 route-based site-to-site VPN, IKE
 - 4in6 and 6in4 policy-based site-to-site VPN, Manual Key
 - 4in6 and 6in4 route-based site-to-site VPN, Manual Key
 - 4in4, 6in6, 4in6, and 6in4 policy-based dial-up VPN, IKE
 - 4in4, 6in6, 4in6, and 6in4 policy-based dial-up VPN, Manual Key
 - Remote Access—XAuth, config mode, and shared IKE identity with mandatory XAuth
 - IKE authentication—public key infrastructure/digital signature algorithm (PKI/DSA)
 - IKE peer type—Dynamic IP
 - Chassis cluster for basic VPN features
 - IKE authentication—PKI/RSA
 - Network Address Translation-Traversal (NAT-T)
 - VPN monitoring
 - Hub-and-spoke VPNs
 - Next Hop Tunnel Binding Table (NHTB)
 - Dead Peer Detection (DPD)
 - Simple Network Management Protocol (SNMP) for IPsec VPN MIBs
 - Chassis cluster for advanced VPN features
 - IPv6 link-local address
 - All high-end SRX Series devices dependency

IPv6 Support

- **NSM**—Consult the Network and Security Manager (NSM) release notes for version compatibility, required schema updates, platform limitations, and other specific details regarding NSM support for IPv6 addressing on all high-end SRX Series devices.
- **Security policy**—Only IDP for IPv6 sessions is supported only for all high-end SRX Series devices. UTM for IPv6 sessions is not supported. If your current security policy uses rules with the IP address wildcard any, and UTM features are enabled, you will encounter configuration commit errors because UTM features do not yet support IPv6 addresses. To resolve the errors, modify the rule returning the error so that it uses the any-ipv4 wildcard; and create separate rules for IPv6 traffic that do not include UTM features.

J-Web

• Browser compatibility on high-end SRX Series devices

Device	Application	Supported Browsers	Recommended Browser
SRX1400, SRX3400, SRX3600, SRX5600, SRX5800	J-Web	- Mozilla Firefox version 3.6 or later - Microsoft Internet Explorer version 7.0	Mozilla Firefox version 3.6 or later

- To use the Chassis View, a recent version of Adobe Flash that supports ActionScript and AJAX (Version 9) must be installed. Also note that the Chassis View is displayed by default on the Dashboard page. You can enable or disable it using options in the Dashboard Preference dialog box, but clearing cookies in Internet Explorer also causes the Chassis View to be displayed.
- On all high-end SRX Series devices, if the device is running the worldwide version of the Junos OS and you are using the Microsoft Internet Explorer Web browser, you must disable the Use SSL 3.0 option in the Web browser to access the device.
- On all high-end SRX Series devices, users cannot differentiate between Active and Inactive configurations on the System Identity, Management Access, User Management, and Date & Time pages.

Logical Systems

- The master logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota because traffic loss could occur. When upgrading all high-end SRX Series devices from Junos OS Release 11.2, make sure that the reserved CPU quota in the security profile that is bound to the master logical system is configured for 1 percent or more. After upgrading from Junos OS Release 11.2, the reserved CPU quota is added to the default security profile with a value of 1 percent.
- Starting with Junos OS Release 11.2, address books can be defined under the **[security]** hierarchy level instead of the **[security zones]** hierarchy level. This enhancement makes configuring your network simpler by allowing you to share IP addresses in address books when configuring features such as security policies and NAT. You can attach zones to address books—this is known as zone-attached configuration.

Junos OS Release 12.1 continues to support address book configuration under the **[security zones]** hierarchy level—this is known as zone-defined configuration. However, we recommend that zone-attached address book configuration be used in the master logical system and user logical systems.

If you upgraded your high-end SRX Series devices to this Junos OS Release 12.1, and are configuring logical systems on the device, the master logical system retains any previously configured zone-defined address book configuration. The master administrator can run the address book upgrade script to convert zone-defined configuration to zone-attached configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.

See the section, “Upgrade and Downgrade Scripts for Address Book Configuration” of [“Upgrade and Downgrade Instructions for Junos OS Release 12.1 for High-End SRX Series Services Gateways”](#) on page 217.

- On all high-end SRX Series devices, the logical systems feature does not support ALGs for user logical systems because ALGs are configured globally. If you enable ALGs at the root master logical system level, they are also enabled for user logical systems in Junos OS Release 12.1. In this case, user logical system traffic is processed by the ALGs, and corresponding ALG flow sessions are initiated under the user logical system. You can only enable and disable ALGs at the root master logical system level.
- On all high-end SRX Series devices, quality-of-service (QoS) classification across interconnected logical systems does not work.
- On all high-end SRX Series devices, the number of logical system security profiles you can create is constrained by an internal limit on security profile IDs. The security profile ID range is from 1 through 32 with ID 0 reserved for the internally configured default security profile. When the maximum number of security profiles is reached, if you want to add a new security profile, you must first delete one or more existing security profiles, commit the configuration, and then create the new security profile and commit it. You cannot add a new security profile and remove an existing one within a single configuration commit.

If you want to add more than one new security profile, the same rule is true. You must first delete the equivalent number of existing security profiles, commit the configuration, and then create the new security profiles and commit them.

- **User and administrator configuration for logical systems**—Configuration for users for all logical systems and all user logical systems administrators must be done at the root level by the master administrator. A user logical system administrator cannot create other user logical system administrators or user accounts for their logical systems.
- **Name-space separation**—The same name cannot be used in two logical systems. For example, if logical-system1 includes the username “Bob” then other logical systems on the device cannot include the username “Bob”.
- **Commit rollback**—Commit rollback is supported at the root level only.
- **Trace and debug**—Trace and debug are supported at the root level only.
- **Class of service**—You cannot configure class of service on logical tunnel (lt-0/0/0) interfaces.
- **ALGs**—The master administrator can configure ALGs at the root level. The configuration is inherited by all user logical systems. It cannot be configured discretely for user logical systems.

Network Address Translation (NAT)

- On all high-end SRX Series devices, in case of SSL proxy, sessions are whitelisted based on the actual IP address and not on the translated IP address. Because of this, in the whitelist configuration of the SSL proxy profile, the actual IP address should be provided and not the translated IP addresses.

Example:

Consider a destination NAT rule that translates destination IP address 20.20.20.20 to 5.0.0.1 using the following commands:

- **set security nat destination pool d1 address 5.0.0.1/32**
- **set security nat destination rule-set dst-nat rule r1 match destination-address 20.20.20.20/32**
- **set security nat destination rule-set dst-nat rule r1 then destination-nat pool d1**

In the above scenario, to exempt a session from SSL proxy inspection, the following IP address should be added to the whitelist:

- **set security address-book global address ssl-proxy-exempted-addr 20.20.20.20/32**
- **set services ssl proxy profile ssl-inspect-profile whitelist ssl-proxy-exempted-addr**
- Maximum capacities for source pools and IP addresses have been extended on all high-end SRX Series devices as follows:

Pool/PAT Maximum Address Capacity	SRX1400	SRX3400 SRX3600	SRX5600 SRX5800
Source NAT pools	8192	8192	12288
IP addresses supporting port translation	8192	8192	12288
PAT port number	256M	256M	384M

Increasing the capacity of source NAT pools consumes memory needed for port allocation. When source NAT pool and IP address limits are reached, port ranges should be reassigned. That is, the number of ports for each IP address should be decreased when the number of IP addresses and source NAT pools is increased. This ensures NAT does not consume too much memory. Use the **port-range** statement in configuration mode in the CLI to assign a new port range or the **pool-default-port-range** statement to override the specified default.

Configuring port overloading should also be done carefully when source NAT pools are increased.

For source pool with port address translation (PAT) in range (64,510 through 65,533), two ports are allocated at one time for RTP/RTCP applications, such as SIP, H.323, and RTSP. In these scenarios, each IP address supports PAT, occupying 2048 ports (64,512 through 65,535) for Application Layer Gateway (ALG) module use. On SRX5600 and SRX5800 devices, if all of the 4096 source pool is configured, a port allocation of 8,388,608 is reserved for twin port use.

- **NAT rule capacity change**—To support the use of large-scale NAT (LSN) at the edge of the carrier network, the device-wide NAT rule capacity has been changed.

The number of destination and static NAT rules has been incremented as shown in [Table 15 on page 185](#). The limitation on the number of destination-rule-set and static-rule-set has been increased.

[Table 15 on page 185](#) provides the requirements per device to increase the configuration limitation as well as to scale the capacity for each device.

Table 15: Number of Rules on all High-End SRX Series Devices

NAT Rule Type	SRX1400	SRX3400 SRX3600	SRX5600 SRX5800
Source NAT rule	8192	8192	8192
Destination NAT rule	8192	8192	8192
Static NAT rule	20480	20480	20480

The restriction on the number of rules per rule set has been increased so that there is only a device-wide limitation on how many rules a device can support. This restriction is provided to help you better plan and configure the NAT rules for the device.

Security

- On all high-end SRX Series devices, the current SSL proxy implementation has the following connectivity limitations:
 - The SSLv2 protocol is not supported. SSL sessions using SSLv2 are dropped.
 - SSL sessions where client certificate authentication is mandatory are dropped.
 - SSL sessions where renegotiation is requested are dropped.
- On all high-end SRX Series devices, for a particular session, the SSL proxy is only enabled if a relevant feature related to SSL traffic is also enabled. Features that are related to SSL traffic are Intrusion Detection and Prevention (IDP), application identification, application firewall, and application tracking. If none of the above listed features are active on a session, the SSL proxy bypasses the session and no logs are generated in this scenario.
- On all high-end SRX Series devices, the limitation on the number of addresses in an address-set has been increased to 1024. The default value of address-set is 1024. The number of addresses in an address-set, which depends on the device, is equal to the number of addresses supported by the policy.

Simple Network Management Protocol (SNMP)

- On all high-end SRX Series devices, the **show snmp mib** CLI command will not display the output for security related MIBs. We recommend that you use an SNMP client and prefix **logical-system-name@** to the community name. For example, if the community is **public**, use **default@public** for default root logical system.

Virtual Private Networks (VPNs)

- On all high-end SRX Series devices, when you enable VPN, overlapping of the IP addresses across virtual routers is supported with the following limitations:
 - An IKE external interface address cannot overlap with any other virtual router.
 - An internal/trust interface address can overlap across any other virtual router.
 - An **st0** interface address cannot overlap in route-based VPN in point-to-multipoint tunnels such as NHTB.
 - An **st0** interface address can overlap in route-based VPN in point-to-point tunnels.
- On all high-end SRX Series devices, the DF-bit configuration for VPN only works if the original packet size is smaller than the **STO** interface MTU, and larger than the **external interface - ipsec overhead**.
- The local IP feature is not supported on the following:
 - All SRX Series devices in chassis cluster configuration
 - All high-end SRX Series devices
- On all high-end SRX Series devices, the IPsec NAT-T tunnel scaling and sustaining issues are as follows:
 - For a given private IP address, the NAT device should translate both 500 and 4500 private ports to the same public IP address.
 - The total number of tunnels from a given public translated IP cannot exceed 1000 tunnels.

Related Documentation

- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 188](#)
- [Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 187](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 210](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 168](#)

Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways

The following problems currently exist in Juniper Networks high-end SRX Series Services Gateways. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Outstanding Issues in Junos OS Release 12.1R7 for High-End SRX Series Services Gateways

Chassis Cluster

- The counter for incoming traffic on a fabric interface (used for chassis cluster) always shows zero (0). [PR/520962]
- When the **multicast forwarding-cache threshold suppress** CLI option is set to a particular value, an incorrect suppressed value is obtained. [PR/741011]

Command-Line Interface (CLI)

- The SRX1400, SRX3600, and SRX5400 devices become unresponsive and then crash when you configure the **request security pki ca-certificate ca-profile-group** command. [PR/896161]

Flow and Processing

- On SRX3600 devices, a change bit is set for a gencfg client after the client closes. A change bit is set on an ifstate before the client changes to the next state. The function `rts_ifstate_client_close` moves the client from the next location to the END of the chain and clears ALL the bits. As a workaround, clear all the bits from beginning of the chain. [PR/786080]

Interfaces and Routing

- When you use the aggregated redundant Ethernet (reth) interface, traffic loss might be observed when the link member is down. [PR/858519]

Related Documentation

- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 174](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 210](#)

Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways

The following are the issues that have been resolved in Junos OS Release 12.1 for Juniper Networks SRX Series Services Gateways. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.



NOTE: For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://www.juniper.net/prsearch>.

Resolved Issues in Junos OS Release 12.1R7 for High-End SRX Series Services Gateways

Chassis Cluster

- On devices in a chassis cluster, massive amounts of MAC addresses were generated on the fabric link switch port. [PR/833609: This issue has been resolved.]
- On devices in a chassis cluster, Juniper Services Redundancy Protocol (jsrpd) process log messages were displayed even though the cluster was stable with no failover events. [PR/861704: This issue has been resolved.]

Command-Line Interface (CLI)

- On high-end SRX Series devices, fragment packets were dropped if the **tcp-no-flag** option was configured in the screen options. [PR/790916: This issue has been resolved.]
- On all high-end SRX Series devices, running the **show security screen statistics logical-system all zone X** command generated core files, if the X zone did not have screens enabled and if it was part of a logical system. [PR/866559: This issue has been resolved.]

Flow and Processing

- If an SRX Series device was configured as an IPv4 to IPv6 translator, the device used the next header as "IPv6 Fragment" even for packets smaller than 1280 bytes. [PR/754823: This issue has been resolved.]
- The SRX Series services gateways SRX1400, SRX3400, and SRX3600 devices did not pad ethernet packets with zeros. Therefore, some packets contained fragments of data from previous packets handled by the interface. [PR/829536: This issue has been resolved.]
- On SRX series device with Protocol Independent Multicast (PIM) protocol enabled, the flowd process crashed and generated coredump file when there was a unicast PIM register message received with encapsulated multicast data and NAT process was involved in the session for the received PIM packet. This issue occurred both on SRX standalone and High Availability (HA) chassis cluster mode. When the issue occurred on HA chassis cluster mode, the flowd process crashed on both node 0 and node 1. [PR/842253: This issue has been resolved.]

- When you commit security policy changes, under certain load conditions (based on the Services Processing Unit (SPU) usage and number of active sessions) and in situations where policy rematch must be performed (either when policy rematch is configured or new policies are added, or the order is changed), SPU usage increase and partial packet drops were observed. [PR/854412: This issue has been resolved.]
- On devices enabled with SYN cookie protection, after the SYN cookie function was triggered, the SYN cookie did not send ACK to the client to update the TCP window size after a handshake with the server. When the client sent ACK with a PSH flag to the device as the third TCP ACK during the TCP three-way hand shake, the device did not recognize the ACK. This resulted in TCP connection failure. [PR/859222: This issue has been resolved.]
- On all high-end SRX Series devices, when AI is handling SSL encrypted sessions with SSLFP enabled, a core file is generated if the client sends large amount of data to the server in a single transaction. [PR/859951: This issue has been resolved.]
- On devices with increased ALG or proxy traffic, memory leaks in global data plane memory were observed, and traffic (FTP, MSRPC, AppID, and so on) dropped. [PR/859956: This issue has been resolved.]

Infrastructure

- When the backup Routing Engine kernel failed, some devices sent a message to the master Routing Engine to generate a core file. This is caused problems. [PR/854501: This issue has been resolved.]

Interfaces and Routing

- For Physical Interface Module (PIM) traffic, the device used its embedded IP packet for session lookup and policy lookup. [PR/806912: This issue has been resolved.]
- In an invalid subnet configuration on a multicast group, when you perform a commit or commit check, the routing protocol process (rpd) crashes and generated a core file. [PR/856925: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices with the IDP application-level distributed denial-of-service (DDoS) feature enabled, if the binary analysis report function was enabled, the device generating IDP application-level DDoS attack logs crashed the flowd process and core files were generated. [PR/865469: This issue has been resolved.]
- When the **no-reset-on-policy** option was set and there were two active policies in a dataplane, and only one session referred to the older policy, flowd core files were generated, if the application identification indicated a change in application (from the default one, for example, FTP running on Telnet port), because of policy re-lookup. [PR/880408: This issue has been resolved.]

J-Web

- In J-Web, information on routes was not listed under the **Configure > Routing > Static Routing** section. [PR/864324: This issue has been resolved.]

System Logs

- On SRX5800 devices, when configuration messages exceeded the interprocess communication message (IPC) maximum transmission unit (MTU), occasionally the following error message was displayed:

ipc_msg_write: %PFE-3: IPC message type: 27, subtype: 2 exceeds MTU, mtu 3216, length 3504[

[PR/612757: This issue has been resolved.]

Unified Threat Management (UTM)

- On devices in a chassis cluster, the antivirus database was not synchronized on both the cluster nodes. [PR/863181: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R6 for High-End SRX Series Services Gateways

Application Layer Gateways (ALGs)

- On SRX1400 devices, after you upgraded to Junos OS Release 11.4R2, the RTSP ALG did not open a pinhole for IXIA, because the /r/n characters were added to the packet. [PR/842470: This issue has been resolved.]

Chassis Cluster

- After multiple node failovers, the chassis cluster LEDs showed as unlit even if the cluster was stable. [PR/789190: This issue has been resolved.]
- On SRX3600 devices, in certain circumstances one of the Services Processing Cards (SPCs) was stuck due to a hardware fault, and the following error message was displayed in the jsrpd log:

Jan 17 23:07:22 Index: 16 PFE Id: 16, Error_code: 0x01 - Loopback

[PR/851317: This issue has been resolved.]

Dynamic Host Configuration Protocol (DHCP)

- On all high-end SRX Series devices, the DHCPV6 server did not create any server binding. [PR/799829: This issue has been resolved.]

Dynamic VPN

- Dynamic VPN on Windows 7, 64-bit operating system (OS) did not work in some environments. [PR/842607: This issue has been resolved.]

Flow and Processing

- During configuration and maintenance of an SRX Series device, occasionally the security match policies were out of sync between the Packet Forwarding Engine and the Routing Engine. In most cases, an error message was displayed during an attempt to commit the configuration. [PR/836489: This issue has been resolved.]

Interfaces and Routing

- Transmit (Tx) and receive (Rx) lockup of the tsec1 (em0) controller caused the em0 interface to go down and all the field-replaceable units (FRUs) to go offline. [PR/820210: This issue has been resolved.]
- On SRX1400 devices, commit on configuration with the lt-0/0/0 interface failed. [PR/845837: This issue has been resolved.]

In-Service Software Upgrade (ISSU)

- On SRX5600 devices, in-service software upgrade (ISSU) failed. [PR/843708: This issue has been resolved.]

J-Web

- When you tried to commit from J-Web for logical systems configurations, the following error message was received even if configuration changes were made: "You have pending changes from previous commit" [PR/812896: This issue has been resolved.]
- In J-Web, the session expired when the idle-timeout value was set too low. [PR/830644: This issue has been resolved.]
- In J-Web, when more than one security policy was configured on a device, the first policy was not listed in the Apply-Policy section. [PR/837799: This issue has been resolved.]
- When the policy name was "0", the preprocessor hypertext function treated it as empty, so the log output could not be viewed. [PR/853093: This issue has been resolved.]

Security Policy

- During configuration and maintenance of a device, occasionally the security match policies did not synchronize between the Packet Forwarding Engine and the Routing Engine. In most cases, an error message was displayed during the attempt to commit the configuration. [PR/836489: This issue has been resolved.]

- On SRX3600 devices, in the flow session output, if the security policy of the flow session referenced a custom application with **application-protocol ignore** set, the application field was set incorrectly. [PR/852081: This issue has been resolved.]

SNMP

- SNMP query for maximum total session (jnxJsSPUMonitoringMaxTotalSession) was taking the maximum value, that is, max-cp-session value. [PR/838214: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R5 for High-End SRX Series Services Gateways

Application Layer Gateways (ALGs)

- The MSRPC ALG dropped some large packets under the Kerberos authentication environment because the Kerberos ticket token size was too large and MSRPC bind packet was too big for the ALG to handle. [PR/817453: This issue has been resolved.]

Chassis Cluster

- On SRX5600 and SRX5800 devices in a chassis cluster, when the cluster was up and if an SPC with control link was pulled out (either one in a dual control link), the node on which SPC was pulled out would cause all the cards to go offline. Other nodes would become primary for both RG0 and RG1+ because it was the split brain case, and traffic loss was expected for 18-20 seconds. [PR/819124: This issue has been resolved.]
- On all high-end SRX Series devices, the performance monitor message format was changed in Junos OS release 11.1 and later releases. In the earlier releases (prior to Junos OS Release 11.1), the message format generated a rtlogd core file and rtlogd restarted automatically after 1 or 2 seconds. [PR/819700: This issue has been resolved.]

Command-Line Interface (CLI)

- In certain configurations, the following message was displayed in the logs: "PFEMAN: Sent Resync request to Master". This message was not an indication of a problem, and did not require any action from the customer. This message has been removed. [PR/802355: This issue has been resolved.]
- A new filter was added in dynamic attack groups in the CLI. The two flags under filters were recommended (which means true) and no-recommended (which means false). Previously, only 'recommended=true' flag was supported. [PR/828494: This issue has been resolved.]

Dynamic Host Configuration Protocol (DHCP)

- On SRX3600 devices, **dhcp-option82** did not work. [PR/794522: This issue has been resolved.]

Flow and Processing

- The content filter for SMTP block extension did not work when the name of the attached file was in Japanese language. [PR/724960: This issue has been resolved.]
- An event core file was generated after heavy amount of logs were archived on a remote site. [PR/771228: This issue has been resolved.]
- Generation of a flowd core file was triggered by cache errors. [PR/805975: This issue has been resolved.]
- On SRX5600 and SRX5800 devices, the lchip correctable error correction code (ECC) had been categorized as a **cmalarm** major error. When the lchip correctable ECC error was reported by the system, the related FPC was identified as a hardware issue, causing traffic interruption. If the device is configured in a chassis cluster, data plane Redundancy Group (RG) failover occurred. [PR/809380: This issue has been resolved.]
- On all high-end SRX Series devices, the “XLR ingress paused” condition might cause traffic drop and latency to process traffic. [PR/818621: This issue has been resolved.]
- The SMTP parser session would hang when TLS was enabled. [PR/824027: This issue has been resolved.]

Interfaces and Routing

- The ge interface went down and did not go up until the configuration was removed and readded. [PR/813616: This issue has been resolved.]

J-Web

- On all high-end SRX Series devices, configuring the global address book in J-Web was not supported. [PR/807115: This issue has been resolved.]

Logical Systems

- Logical systems with policy count option displayed the statistics after a while following a **show** command, or the counters stopped to increment when both redundant groups were not on the same node as a result of a failover. [PR/782546: This issue has been resolved.]
- When you committed any changes under logical systems configuration, the security policy failed to resolve DNS objects in the security address book. Hence, the traffic might hit another unexpected security policy, or default-deny instead, and then cause traffic outage. [PR/810723: This issue has been resolved.]

Management Information Base (MIB)

- The `jnxJsldp LastSignatureUpdateTime.0` MIB value never changed. [PR/691785: This issue has been resolved.]

Upgrade and Downgrade

- After an upgrade to Junos OS Release 11.4R5, if OSPF was enabled for any of the st0 interfaces, an internal processing error prevented the default route from being advertised. [PR/822352: This issue has been resolved.]

Virtual Private Network

- The pulse client showed the following error message “Failed to connect to server” when multiple clients were connected to the device through dynamic VPN and a configuration related to IKE negotiation changed on the device. [PR/737787: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R4 for High-End SRX Series Services Gateways

Application Layer Gateways (ALGs)

- On all high-end SRX Series devices, when using the IKE-ESP-NAT ALG to pass through for the Cisco EZ-VPN client, the IKE handshake might not be successful, because the IKE packet coming from the VPN server got dropped. [PR/791549: This issue has been resolved.]
- On all high-end SRX Series devices, the SQL ALG did not work properly when TNS RESEND (type 11) was 8 bytes long. [PR/806893: This issue has been resolved.]

Command-Line Interface (CLI)

- On all high-end SRX Series devices, the `set chassis fpc pic services-offload` command did not work. [PR/787526: This issue has been resolved.]

Flow and Processing

- On all high-end SRX Series devices, when the syn-cookie feature was enabled along with syn-flood screen with a low timeout value, high latency TCP sessions may fail to establish successfully. The clients sessions receive unresponsive connections because the SRX device has timed out the flow for this session. The device will also drop subsequent packets from the client due to no state found. [PR/692484: This issue has been resolved.]
- Typing the Alt-break sequence on the console placed the console interface in debugger mode (the `db>` prompt). You can now configure the “`system no-debugger-on-alt-break`” statement to disable the Alt-break sequence on the serial console. [PR/717491: This issue has been resolved.]
- On all high-end SRX Series devices, in some cases an SPU message **Nexthop XXXX on ifl XXX. Ignoring...** was logged. [PR/726580: This issue has been resolved.]

- On all high-end SRX Series devices, high CPU use due to mgd process might result when the **run show config** command was specified during configuration mode. In addition, the httpd process go high. [PR/729617: This issue has been resolved.]
- The Junos OS kernel might crash because of a timing issue in the ttymodem() internal I/O processing routine. The crash can be triggered by simple remote access (such as Telnet or SSH) to the device. [PR/755448: This issue has been resolved.]
- On all high-end SRX Series devices, the interface did not respond due to heavy traffic on the FIOC interface. [PR/776179: This issue has been resolved.]
- On all high-end SRX Series devices, if security policies were configured with a large number of applications using the same source and destination ports, then policy configuration updates might not work as expected. [PR/793151: This issue has been resolved.]

Interfaces and Routing

- On SRX3600, SRX3400, and SRX1400 devices, the aggregated Ethernet interface might go down after users configured “Active LACP” on the back-to-back connected AE bundles. [PR/770998: This issue has been resolved.]
- After a Routing Engine switchover, LACP and MIB process (mib2d) core files might be created. [PR/790966: This issue has been resolved.]
- On SRX Series devices, after a reboot, the VLAN interface was down while it's physical interface member was up. [PR/795363: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On SRX Series devices, the detector was not updated in the control plane when the update-attack-database-only flag was used during security-package installation. [PR/778816: This issue has been resolved.]
- On SRX650 devices, the messages “idpd: Failed to add license for feature” error appears during commit when IDP is configured. [PR/793545: This issue has been resolved.]

Logical Systems

- On all SRX Series devices running logical systems, you cannot use snmpwalk for SNMP polling. [PR/791859: This issue has been resolved.]

Network Address Translation (NAT)

- On SRX Series devices, commit of static NAT rules might fail when you commit interfaces, security zone, and NAT at same time in root or LSYS. Problem might also occur that the commit of static NAT rules might fail when commit for security zone, and NAT occur at same time. [PR/756240: This issue has been resolved.]

SNMP MIBs

- On all high-end SRX Series devices in a chassis cluster, long pauses and timeout were seen during SNMP walk/query of the device. A delay occurred in the kernel's query of the gr-0/0/0 (GRE) interface. [PR/800735: This issue has been resolved.]

- On high-end SRX Series devices, Routing Engine failover occurs due to possible out-of-sync information about already allocated SNMP interface index values, duplicate SNMP interface index values might be allocated. As a result, the mib2d process might crash or the SNMP interface index value of zero might be allocated for newly created interfaces. [PR/806098: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R3 for High-End SRX Series Services Gateways

Application Layer Gateways (ALGs)

- On SRX5800 devices, the fragmented packets with the DF bit set (do not fragment) might be dropped by the device when processed by ALG. This problem might occur when fragmented packet is set to DF when it should not be fragmented anymore. [PR/754504: This issue has been resolved.]
- On SRX3600 devices, SIP ALG is dropping SIP ACK messages, if messages are using the folding format. [PR/787879: This issue has been resolved.]
- On SRX Series devices, when using the IKE-ESP-NAT ALG to pass through for Cisco EZ-VPN client, the IKE handshake might not be successful, because IKE packet coming from VPN server gets dropped. [PR/791549: This issue has been resolved.]

AppSecure

- On all SRX Series devices, in the output of the **show services application-identification application-system-cache** command, the **application-system-cache** table for P2P encrypted traffic is incorrectly marked as **Encrypted: No** instead of **Encrypted:Yes**. [PR/704023: This issue has been resolved.]

Chassis Cluster

- On SRX3400 devices, flow sessions on backup node could not be cleared previously. [PR/729967: This issue has been resolved.]
- On all high-end SRX Series devices in chassis cluster, when the secondary node is reboot or shutdown, there could be a transient traffic drop on the primary node. The amount of drop depends on the number of active sessions. The drop is caused by a route change and an RTO cold sync during a secondary reboot. After the route change and RTO cold synchronization is complete, the traffic returns to normal state and the drop time window might be a few seconds. [PR/734966: This issue has been resolved.]
- On SRX3600 device in chassis cluster, when the configuration is saved on a remote file server, you must mention absolute/relative path for storing the file. If the absolute/relative path is not used, it fails to save the file. However this issue is not seen when device is operating in stand-alone mode. [PR/752363: This issue has been resolved.]
- On all high-end SRX Series devices, an RG-0 failover after ISSU completion (both nodes have been upgraded and have joined the cluster with priorities restored) might cause the new RG-0 secondary node to restart all FPCs. This might lead to an RG-1+ failover. [PR/770710: This issue has been resolved.]

Flow and Processing

- On SRX5600 virtual chassis, when you swap the members of a link aggregation group (LAG), a vmcore or ksyncd core file might be created on the backup Routing Engine. [PR/711679: This issue has been resolved.]
- On SRX5800 and SRX5600 devices, when you attempt configure more than 2000 zones, the system displays the incorrect message. This is only a display issue and has no impact on the functionality. [PR/734378: This issue has been resolved.]
- On SRX Series devices, after upgrading to Junos OS release 12.1, if a commit is tried after executing a commit confirmed, the following error message is displayed:
error: problem checking file: No such file or directory.
[PR/741239: This issue has been resolved.]
- On all SRX Series devices, the captive portal redirect does not work with strict SYN checking option enabled in the firewall. [PR/743466: This issue has been resolved.]
- On SRX Series devices, commands after STARTTLS are encrypted, and can not be understood by the SMTP parser, are causing the session to hang till TCP session is closed, with no packets forwarded. [PR/750047: This issue has been resolved.]
- On SRX3600 devices, it was observed that the graceful restart mechanism was not getting aborted even if the link to the upstream neighbor was detected as down. This was leading to higher routing protocol convergence times as the route would not fail over to an alternate path until the graceful restart timer expired. [PR/751640: This issue has been resolved.]
- On SRX3600 devices, the inbound “to-self” SSH traffic is accepted by the device even though 'ssh' is not explicitly included in the 'host-inbound-traffic' configuration for the ingress interface within the security-zone. [PR/754392: This issue has been resolved.]
- On SRX Series devices, when SYN flood packets per second (pps) over SCREEN attack-threshold, SYN cookie will be triggered by default. As design, the SRX Series device sends SYN ACK to client with ISN, and once the correct ACK is received, the device sends SYN to server. But, the ACK packet(from client) might create a session and forwarded to server. Due to no SYN packet received instead of ACK from client, server will send RST and the RST will be forwarded to client, and connection would be reset. [PR/755727: This issue has been resolved.]
- When an FPC restart is performed, some of the PICs and IFDs were unable to be created by chassisd due to EBUSY error returned by the kernel. Kernel was unable to process the new requests until the previous states of the same object (PIC, IFD in our case) was consumed by all peers interested in this. The enhancement addresses the design which makes sure new state changes which could have been processed by the faster peers are not blocked due to these slower peers. [PR/769632: This issue has been resolved.]
- On SRX3400 devices, RPD memory leak occurs when SNMP polls BGP and BGP is not configured. [PR/776637: This issue has been resolved.]

Hardware

- On SRX3400 devices, the message log is too granular, indicating blower speed changes frequently from normal to intermediate speed which is filling up logs and causing inability to troubleshoot using logs. [PR/776254: This issue has been resolved.]

Interfaces and Routing

- On SRX5800 devices, when multiple interfaces are bound to the same security zone, if the first fragmented packet and the second fragmented packet arrive in different interfaces, the second fragmented packet is dropped. [PR/777343: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, when two or more IDP policies are configured in the root LSYS and one policy is active in the root LSYS and a different policy is active in the custom LSYS, the referenced LSYS policy might not get compiled properly after a signature update. [PR/749126: This issue has been resolved.]
- On all high-end SRX Series devices, detector does not get updated in the control plane when update-attack-database-only flag is used during security-package install.
[PR/778816: This issue has been resolved.]

J-Web

- On SRX5800 devices, if multiple J-Web clients are connected to a device, it could cause high CPU usage on the routing engine. [PR/741432: This issue has been resolved.]
- On all high-end SRX Series devices, the source interface for IP monitoring must be a logical interface. However, the corresponding configuration screen on J-Web does not list logical interface and only list physical interface. [PR/754523: This issue has been resolved.]
- The J-Web interface is vulnerable to HTML cross-site-scripting attacks, also called XST, or cross-site-tracing. [PR/752398: This issue has been resolved.]

Logical Systems (LSYS)

- On high-end SRX Series devices, the flowd process (the process responsible for traffic forwarding in SRX devices) might crash when running on logical system (LSYS). [PR/780019: This issue has been resolved.]

Network Address Translation (NAT)

- On SRX Series devices, commit of static NAT rules might fail when you commit interfaces, security zone, and NAT at same time in root or LSYS. Problem might also occur that the commit of static NAT rules might fail when commit for security zone, and NAT occur at same time. [PR/756240: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R2 for High-End SRX Series Services Gateways

AppSecure

- On all high-end SRX Series devices, automatic refresh on the Application Tracking monitor page does not clear the pop-up window with the message **Data is very minimal to display in PIE Chart**, even when the chart gets loaded with data on refresh. [PR/724995: This issue has been resolved.]

Command-line Interface (CLI)

- On all high-end SRX Series devices, the Application-groups statistics are shown as **unassigned** and **unknown** for the **show services application-identification statistics application-groups** command output without displaying the details. [PR/740014: This issue has been resolved.]

Flow and processing

- On SRX3600 devices, abnormal SQL traffic might cause the flowd process to crash when SQL ALG is enabled. [PR/737468: This issue has been resolved.]
- On SRX Series devices, when you enable "Change password every time the user logs out" on the active directory, the user is unable to change his or her password. [PR/740869: This issue has been resolved.]

Interfaces and Routing

- On SRX5600 and SRX5800 devices, regarding the K2-RE (64-bit Routing Engine) when speed or link mode is statically configured on the router for the fxp0 interface, the driver for fxp0 accepts the configuration from DCD process, but does not propagate the setting to the hardware driver. Instead, the driver setting is forced to auto-negotiate. Thus, as the fxp0 interface is auto-negotiating, and the far end device is forced to 100/full, the auto-negotiation on fxp0 will detect the speed but not the duplex and defaults that duplex to half-duplex. [PR/704740: This issue has been resolved.]
- On SRX3400 devices, inserting a defective 16-Port SFP Gigabit Ethernet IOC on the device causes all other SFP cards to be unrecognized. [PR/711461: This issue has been resolved.]

- On SRX3600 devices, under certain circumstance, reverse static NAT might not work correctly. [PR/739142: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On SRX3600 devices, IDP configured on Junos OS release 11.4R1, the output of the command “show security idp policy-commit-status” incorrectly displays “Active policy not configured” even when active policy is configured. [PR/741736: This issue has been resolved.]
- On all high-end SRX Series devices, after 24 hours of stress run, with a huge number of sessions generated, IDP sessions were not increasing along with the flow sessions. [PR/742882: This issue has been resolved.]
- On all high-end SRX Series devices, when IDP signature package automatic update is configured, sometimes signature package (secDb) installation fails due to timing issue. [PR/743915: This issue has been resolved.]

IPv6

- On all high-end SRX Series devices, the flow bytes counters tracked on a per-interface basis are incorrect for IPv6 flows. The flow input bytes statistics are correct, but flow output bytes statistics are reversed to the source/destination interfaces. [PR/740911: This issue has been resolved.]

J-Web

- On SRX1400 devices, on the Configure>Security>Policy Elements>Address Book>Address Sets page in the J-Web interface, the nested address-sets are not displayed and cannot be configured. [PR/726227: This issue has been resolved.]
- On all high-end SRX Series devices, on the Configure>Security>Logging page in the J-Web interface, the **Apply** button will remain disabled even after changing the configuration. [PR/730861: This issue has been resolved.]
- On SRX1400 devices, when you access J-Web -> Dashboard -> Chassis view, the SYSIO slot is blacked out or grayed out. [PR/738772: This issue has been resolved.]

Network Address Translation (NAT)

- On all high-end SRX Series devices, when configuring a NAT rule set in the J-Web interface, security zones cannot be displayed if no interfaces are configured for those security zones. [PR/703264: This issue has been resolved.]
- On SRX5600 devices, the flowd_xlr core files are generated occasionally resulting in a system crash. The reason for the system crash might be changing or deleting the IP address of the interface while the interface configuration involves enabling persistent source NAT. [PR/744687: This issue has been resolved.]
- On SRX3400 and SRX3600 devices, you can configure a security zone in the format “a.b.c.d”. However, a configuration error was displayed when the same zone name was referenced while configuring the NAT. [PR/748621: This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R1 for High-End SRX Series Services Gateways

Application Layer Gateways (ALGs)

- On SRX3400 and SRX5600 devices, calls per second (CPS) of RTSP ALG traffic in both Layer 2 and Layer 3 mode were dropped to 300 per SPU. [PR/676053: This issue has been resolved.]
- On all high-end SRX Series devices, the tcp-proxy module used by the ALG had a deficiency when handling a TCP stream with large packets. [PR/727649: This issue has been resolved.]

Chassis Cluster

- On all high-end SRX Series devices operating in a chassis cluster, the XML response to the RPC validate command did not return the closing `</routing-engine>` tag for the second node. [PR/679413: This issue has been resolved.]
- On SRX3600 devices, in rare instances, the primary Routing Engine might not send the deletion of the destination route pointing to the decoupled next hop to the secondary Routing Engine. This caused a system panic, and a VMcore file was created on the secondary Routing Engine. [PR/684981: This issue has been resolved.]
- On SRX3400 devices, when a service-offload session was set up in a multicast session on the primary node, it acted as a normal session in the backup node. This occurred because the services-offload flag was not correctly synchronized to the backup node. During a chassis cluster failover, a route change was triggered, and the service-offload session was reinstalled. [PR/696819: This issue has been resolved.]
- On SRX3400 devices, during failover, there was a small window of time in which the SPU did not detect whether an NP was in services-offload mode or not. This might cause a small number of the services-offload sessions to change to normal sessions. [PR/697426: This issue has been resolved.]
- On SRX1400 devices in a chassis cluster, unwanted timed out data path trace messages were seen. [PR/703272: This issue has been resolved.]
- On SRX5800 devices in a chassis cluster, if ALG traffic was too high and both the ports were used, some single ports on backup caused traffic outflow. [PR/705799: This issue has been resolved.]
- On SRX1400 devices in a chassis cluster, the primary control link heartbeats were not seen and were causing inconsistent system behavior. This happened when a 16-Port SFP Gigabit Ethernet I/O card or a 16-Port Copper Ethernet/Fast Ethernet/Gigabit Ethernet I/O card was on the device. [PR/718054: This issue has been resolved.]
- On all high-end SRX Series devices operating in a chassis cluster, if an ECMP (equal-cost multipath) route had both local and remote interfaces, then the local interface was favored for the next hop to avoid the performance-related issues that involved forwarding the traffic across the fabric link. [PR/718807: This issue has been resolved.]
- On SRX3600 device in a chassis cluster mode, when an LACP is configured on a remote switch which is connected to a Reth interface, and the control link between nodes is

lost, it causes both nodes to become primary. [PR/733335: This issue has been resolved.]

Command-Line Interface (CLI)

- On SRX5800 devices, the configuration command **set security pki ca-profile profile-name revocation-check crl disable on-download-failure** could not prevent a revocation check even when the PKI server was unreachable. [PR/605042: This issue has been resolved]
- On all high-end SRX Series devices, the **clear services application-identification application-system-cache** command did not include the following options:
 - node— Clears the cache entry from both the nodes
 - node local— Clears the cache entry from the local device
 - node primary— Clears the cache entry from the primary node[PR/678624: This issue has been resolved.]
- On all high-end SRX Series devices, deleting the groups at a certain hierarchy level could erase junos-defaults groups. [PR/689912: This issue has been resolved.]
- On all high-end SRX Series devices, under certain conditions, session numbers displayed by the **show security flow session summary** and **show security flow session services-offload** commands might not be consistent. [PR/700461: This issue has been resolved.]
- On all high-end SRX Series devices, the **set security zones security-zone zone tcp-rst** command had an incorrect configuration description as **Send RST for TCP SYN packet not matching session**. The correct description is **Send RST for TCP non-SYN packet not matching existing session**. [PR/703196: This issue has been resolved.]
- On all high-end SRX Series devices, the description of the screen was displayed in operational mode. [PR/712242: This issue has been resolved.]
- On all high-end SRX Series devices, the show command for ntp (**show ntp status**) was not working. [PR/722537: This issue has been resolved.]
- On all high-end SRX Series devices, output for the **show security ike security-associations** and **show security ike security-associations detail** commands is presented in a disorderly manner. This was a display issue and had no impact on performance. [PR/729804: This issue has been resolved.]

Dynamic Host Configuration Protocol (DHCP)

- On SRX3600 devices, the DHCP relay server was required to have DNS configuration for DHCP clients. The DHCP daemon read the DNS configuration and the daemon was stopped when there was an invalid dns-server value (for example, 0.0.0.0). [PR/706615: This issue has been resolved.]

Flow and Processing

- On all high-end SRX Series devices, when packet-logging functionality was configured with an improved pre-attack configuration parameter value, the resource usage increased proportionally and affected the performance. [PR/526155: This issue has been resolved.]
- On all high-end SRX Series devices, if the maximum number of leaves on a multicast distribution tree was exceeded, the multicast sessions were created up to the maximum number of leaves, and any multicast sessions that exceeded the maximum number of leaves were ignored. In previous releases, no multicast traffic was forwarded if the maximum number of leaves on the multicast distribution tree was exceeded. The maximum number of leaves on a multicast distribution tree is device specific. [PR/561442: This issue has been resolved.]
- On SRX5600 devices, in certain cases, loading an IDP detector caused a flowd crash, showing memcpy as the top of the stack. [PR/570361: This issue has been resolved.]
- On SRX5800 devices, the message log **ipc_msg_write: %PFE-3: IPC message type: 27, subtype: 2 exceeds MTU, mtu 3216, length 3504** appeared occasionally due to internal communication. [PR/612757: This issue has been resolved.]
- On all high-end SRX Series devices, the BFD session on routing protocols was not working. [PR/671444: This issue has been resolved.]
- On all high-end SRX Series devices, additional neighbor solicitation packets sent were corrupted. The packet had zeroed destination MAC address and the entire packet was prefixed by two random bytes. [PR/671658: This issue has been resolved.]
- On SRX3400 devices, failover with the RG0 and RG1 generated a VM core file. [PR/675860: This issue has been resolved.]
- On all high-end SRX Series devices, the **Link failure happened for DPC%d PFE%d** log message displayed an incorrect FPC number. [PR/683371: This issue has been resolved.]
- On SRX5600 devices, SIP calls were not getting cleared for memory leak. [PR/684000: This issue has been resolved.]
- On SRX3400 and SRX3600 devices, over 4 Gbps EF(TOS 101) queue traffic caused SPU crash. [PR/686133: This issue has been resolved.]
- On SRX3400 and SRX3600 devices, adding and deleting IKE gateways multiple times when SAs were active resulted in a KMD core file. [PR/689470: This issue has been resolved.]
- On all high-end SRX Series devices, a data plane failover resulted in a traffic loss for a short period when the device was active with a large number of routes. [PR/690004: This issue has been resolved.]

- On SRX3400 and SRX3600 devices, the incorrect value for the **error mbuf statistics** option was displayed in the **show xlr pkt_mbuf use** vty command on the SPU.
[PR/693200: This issue has been resolved.]
- On SRX5800 devices, under certain conditions, deleting an **fxp0** configuration and rolling it back set the **fxp0** forcefully to nonnegotiated, 100m/full duplex mode.
[PR/696733: This issue has been resolved.]
- On SRX5600 devices, when IKE and IPsec configuration or security configuration was removed and added back frequently, KMD core files were generated. [PR/698718: This issue has been resolved.]
- On SRX3400 and SRX3600 devices, the **reth** interface did not work properly when you changed the **reth** member interfaces to other interfaces using different speeds and the link speeds were mismatched. [PR/698837: This issue has been resolved]
- When certificate key size was 2048 bytes and total combined certificate request size was greater than 4096 bytes, the PKID service might crash on certificate enrollment.
[PR/698846: This issue has been resolved.]
- On all high-end SRX Series devices, a memory leak occurred during the audit event processing. [PR/698907: This issue has been resolved.]
- On all high-end SRX Series devices, the PIM register messages were dropped on the physical interfaces that hosted multiple unnumbered interfaces. [PR/698943: This issue has been resolved.]
- On SRX5600 devices, the message vector `create_pdp_rsp` changed tunnel state from half to active and inserted it into active timer wheel. These actions resulted in no tunnel lock protection. When `del_pdp_rsp` deleted the user tunnel or clear path, this led to a change in both the aging flag and the timer wheel entry pointer. [PR/699147: This issue has been resolved.]
- On SRX3600 devices, when a client started a new TCP session using the same 5 tuples as the existing TCP session (which was terminated within 2 seconds) and FIN/ACK, the new TCP session was allowed for a short time and then the device dropped all subsequent packets from client or server. When this issue occurred, the device allowed only the new TCP handshake packets and then dropped all the subsequent traffic that belonged to the newly created session. [PR/700301: This issue has been resolved.]
- On SRX3600 devices, a session drop due to not being authenticated was logged as `RT_FLOW_SESSION_CLOSE` and had the wrong close reason other than **unset**.
[PR/701971: This issue has been resolved.]
- On SRX5600 devices, sending the ALG traffic and changing the configuration to delete the related policy resulted in flowd core files. [PR/702418: This issue has been resolved.]
- On SRX3600 devices, when you configured services-offload and when IOC cards were plugged in slot 4 or above, the following warning message was not displayed:

In services offload mode, IOC cards must be in the first four slots (slot0 - slot3).

[PR/703114: This issue has been resolved.]

- On SRX1400 devices, RTSP interleave data packets could not be passed when the RTP length was greater than 3000 bytes. [PR/703663: This issue has been resolved.]
 - On SRX5600 devices, the JSRPD used the IIC device to set HA LED. The JSRPD, whose function is `jsrpd_a2a10_set_led_color()` function, was triggered when the **ge-0/0/2** interface was down. The `jsrpd_a2a10_set_led_color()` function did not close when the LED color set function was done. This behavior caused the FD leak and eventually resulted in the generation of the JSRPD core file. [PR/703799: This issue has been resolved.]
 - On all high-end SRX Series devices, flowd core files were generated during stress test. [PR/704482: This issue has been resolved.]
 - On all high-end SRX Series devices, the following warning messages were not displayed:
 - When heap and/or arena memory utilization reached a certain critical level:
 - `Warning: Heap utilization reaches critical level!`
 - `Warning: Arena utilization reaches critical level!`
 - When the memory utilization fell back below the critical level:
 - `Heap utilization falls back to normal level`
 - `Arena utilization falls back to normal level`
- [PR/705118: This issue has been resolved.]
- On SRX5600 devices, packets with packet length of 500 bytes were corrupted when only packet capture of data path debug was present without configuration of the **record-pic-history** and event **np-egress**. [PR/706858: This issue has been resolved.]
 - On all high-end SRX Series devices, the spoofing screen was dropping IPv6 packets sourced from the unspecified source address (example: 0:0:0:0:0:0:0:128 or ::128). [PR/721913: This issue has been resolved.]
 - On all high-end SRX Series devices, on rare occasions, the central point (CP) session counter got an incorrect value, which triggered the **CP/Flow session table full** alarm. This occurred because the flow table was not full and the alarm was incorrectly set due to the wrong session counter value. [PR/725099: This issue has been resolved.]
 - On all high-end SRX Series devices, configuring the **tcp-options** under individual policy resulted in commit failing with the following error message:
`error reading services-offload attribute`
[PR/726059: This issue has been resolved.]
 - On SRX5600 devices, committing the configuration for security policies reordering causes a core dump in FWDD. [PR/726791: This issue has been resolved.]
 - On SRX3600 devices, when a large volume of firewall-authentication requests and changes of security policy occurred at the same time, the devices might crash due to a race condition. [PR/726940: This issue has been resolved.]
 - On SRX5600 devices, flowd core files were generated. [PR/734308: This issue has been resolved.]

- On all high-end SRX Series devices, under certain circumstances, when an anti-spoofing screen was enabled, the neighbor discovery messages were dropped after a neighbor become unreachable and then reachable again. [PR/734423: This issue has been resolved.]
- On SRX3400, SRX3600, SRX5600, SRX5800 devices, sometimes, traffic sent by authenticated users (authenticated by web authentication) was not passing through the firewall. This is because, authentication entry created in CP was not passed along to SPU. [PR/734869: This issue has been resolved.]
- On SRX5600 devices, NSD core files were generated. [PR/735152: This issue has been resolved.]
- On SRX3600 devices, flowd core files were generated. [PR/735924: This issue has been resolved.]

Hardware

- On all high-end SRX Series devices, during cyclic redundancy check (CRC), an alarm were not generated, when a dip4 hardware error occurred on the central point (CP). [PR/683212: This issue has been resolved.]

Installation

- On all high-end SRX Series devices, while downgrading from Junos Release 11.4 or a later image to Junos Release 11.1 or an earlier image, the security package was deleted. Although it was automatically installed when the IDP daemon came up, this automatic installation failed sometimes due to application identification (AI) installation error. [PR/705113: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On SRX3400 devices, heap memory usage increased slightly after you performed multiple pushes of IDP policies with multiple rules. [PR/684630: This issue has been resolved.]

Interfaces and Routing

- On SRX1400 devices, when the system was rebooted for the first time, if the physical connections for the VPN gateways were down, the **st0** interfaces were up even though VPNs were never established. This situation led to “black holing” traffic when static routes were used over those **st0** interfaces. [PR/669833: This issue has been resolved.]
- On SRX3400 devices, the TX lockup of the **em0** interface caused the **em0** interface to go down and caused all field-replaceable units (FRU) to go offline. [PR/685451: This issue has been resolved.]
- On SRX5600 and SRX5800 devices, using ftp to **st0** interface for data transfer fails occasionally. [PR/706827: This issue has been resolved.]

Internet Protocol Security (IPsec)

- On SRX5800 devices, when the device processed a large number of transit IPsec sessions through ike-esp-nat ALG, after sometimes new sessions failed. [PR/671074: This issue has been resolved.]
- On SRX3400 and SRX3600 devices, initial contact notification sent from peer for the IKEv1 protocol was not getting processed. Because of this, stale IKE and IPsec security association were displayed. However, functionality was not affected. [PR/705123: This issue has been resolved.]
- On all high-end SRX Series devices, occasionally, when keys are generated incorrectly for IPsec, VPN occasionally went down due to encapsulation security payload (ESP) failures. [PR/717969: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- On all high-end SRX Series devices, when IDP signature package automatic update is configured, some times signature package (secDb) installation fails due to timing issue. [PR/743915: This issue has been resolved.]

J-Web

- On all high-end SRX Series devices, on the Configure>Security>Policy>Define AppFW Policy> Add Rule Set page in the J-Web interface , the Default Rule was getting displayed, leading to incorrect configurations. [PR/593551: This issue has been resolved.]
 - On SRX3400 devices in a chassis cluster, while monitoring the device through J-Web, the dashboard incorrectly displayed the session utilization value. The displayed value was above 100 percent and activated or maximum sessions were displayed as 3M or 2.5M. The dashboard incorrectly displayed the combined values of both primary and secondary nodes. [PR/666489: This issue has been resolved.]
 - On SRX3600 devices, the policy validating pop-up window in the J-Web was not cleared after it was dragged. [PR/675554: This issue has been resolved.]
 - On all high-end SRX Series devices, when using the CLI you could configure only an AppQoS rule set without configuring any other diff-services. However, in J-Web, you could configure at least one diff-service for a new AppQoS rule set configuration. [PR/686462: This issue has been resolved.]
 - On SRX1400 devices, when you modified an existing policy or created a new policy with junos-host zone, there was no junos-host zone available in the from-zone or to-zone list. [PR/697863: This issue has been resolved.]
 - On all high-end SRX Series devices, the Application Tracking monitor page did not display the warning **no data to display** when there was no data to display in the table. [PR/724985: This issue has been resolved.]
 - On SRX1400 devices, authentication through J-Web failed if the password contained the special characters: ' , " , + , \
- [PR/725901: This issue has been resolved.]

- On SRX1400 devices, on Configure>Security>Policy Elements>Address Book>Address Sets page, the nested address-sets are not displayed and can not be configured. [PR/726227: This issue has been resolved.]
- On all high-end SRX Series devices, on the Configure > Security > Logging page, the **Apply** button remained disabled even after changing the configuration. [PR/730861: This issue has been resolved.]
- On all high-end SRX Series devices, on the Monitor > Application Tracking page, the application groups statistics are displayed as **unknown** and **unassigned** without proper details even if applications are configured. This problem occurs because the data supplied by the CLI is not correct. [PR/741540: This issue has been resolved.]

Management Information Base (MIB)

- On all high-end SRX Series devices, when you polled the device with five SPCs and three SPCs, the device reported the wrong number of sessions for the object ID jnxJsSPUMonitoringMaxTotalSession (1.3.6.1.4.1.2636.3.39.1.12.1.3.0). [PR/488653: This issue has been resolved.]

Logical Systems

- The **logical-systems all** option was not available in some of the operational commands such as **show security screen statistics**. [PR/598032: This issue has been resolved.]
- On SRX3400 devices, when you configured LSYS and then loaded an override configuration with a different LSYS configuration (both having the old and new LSYS), the proxy-ndp route might fail to push to the Packet Forwarding Engine. If you deleted old LSYS and added new LSYS in one commit, the proxy-ndp route failed to push to the Packet Forwarding Engine. [PR/673930: This issue has been resolved.]
- On SRX3400 devices, when NAT64 was used in the logical system (LSYS), the binding did not age out and the reversed binding did not match successfully. The NAT64 in LSYS function did not work. [PR/675052: This issue has been resolved.]
- On SRX5600 devices in a chassis cluster, some NAT sessions kept invalidated status after multiple failovers. [PR/676385: This issue has been resolved.]
- On SRX3400 devices, in J-Web, you could not edit the **lt** interface for LSYS. [PR/700354: This issue has been resolved.]
- On all high-end SRX Series devices, after a logical systems high availability ISSU, the application firewall rule matched session counter was set to an incorrect value (-1) when a new session triggered an application firewall rule lookup. [PR/704979: This issue has been resolved.]
- On SRX1400 devices, the LSYS capacity number for nat-rule-referenced-prefix lsys profile was displayed incorrectly. [PR/707108: This issue has been resolved.]
- On all high-end SRX Series devices running Junos OS Release 11.2 when a logical system feature was added, diagnostic information was sent to a specific file without rotation control, causing core files to be generated. [PR/721104: This issue has been resolved.]

Network Address Translation (NAT)

- On SRX3600 devices, when there was heavy SIP traffic and a share gate was involved, NAT translation-context might leak. [PR/675869: This issue has been resolved.]
- On all high-end SRX Series devices, the static NAT with a default routing instance was not working. [PR/706183: This issue has been resolved.]
- On all high-end SRX Series devices, when two static NAT rules were configured with the same prefix and all the interfaces were in the default routing-instance, and only one rule was configured with the **static-nat prefix routing-instance default** option, the commit completed successfully without prompting about the overlapping information. [PR/708433: This issue has been resolved.]
- On SRX5600 devices, during high-rate NAT session creation, high CPU usage might be seen on the central point on the backup node. [PR/720010: This issue has been resolved.]
- On SRX5600 and SRX5800 devices, using a same source NAT pool by both NAT and NAT-PT traffic causes a core file to be generated. [PR/736374: This issue has been resolved.]

Security

- On SRX3400 devices, occasionally, the user name and the role name in the AppTrack was displayed as "N/A". [PR/742377: This issue has been resolved.]

SNMP

- On SRX3400, SRX3600, SRX5400, and SRX5600 devices, the health monitor feature under SNMP might send out a false alarm indicating high CPU utilization. [PR/610450: This issue has been resolved.]

Upgrade and Downgrade

- On all high-end SRX Series devices, upgrade failed for Junos OS Release 11.1R6.4 to 11.4R2.3 build due to AI installation failure without xcommit result. [PR/729625: This issue has been resolved.]

Virtual Private Network (VPN)

- On SRX5600 devices in a large configuration with heavy traffic, after reboot failover in chassis cluster, the Routing Engine on the new primary node became very busy. As a result, the FPC was detached, causing traffic to fail when passing through the firewall device. [PR/698150: This issue has been resolved.]

Related Documentation

- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 174](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 210](#)

Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways

Errata for the Junos OS Software Documentation

This section lists outstanding issues with the software documentation.

Feature Support Reference for SRX Series and J Series Devices

- In this guide, the Table 14: DHCP Support of the “Dynamic Host Configuration Protocol” section incorrectly states that DHCPV6 relay agent is supported on SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800 devices. The DHCPV6 relay agent is not supported on high-end SRX Series devices.

J-Web

- J-Web pages for stateless firewall filters**—There is no documentation describing the J-Web pages for stateless firewall filters. To find these pages in J-Web, go to **Configure>Security>Firewall Filters**, and then select **IPv4 Firewall Filters** or **IPv6 Firewall Filters**. After configuring the filters, select **Assign to Interfaces** to assign your configured filters to interfaces.

Junos OS CLI Reference

- The [edit applications application *application-name* term *term-name*] hierarchy level for the alg (Applications) configuration statement is incorrect. The correct hierarchy level is [edit applications application *application-name* <term *term-name*>].
- This guide incorrectly lists the following topics. These commands are not supported:
 - disable-call-id-hiding
 - show security alg sip transactions
- In the “show security policies” topic, the “show security policies Output Fields” table includes the following incorrect information:

Applications	ALG: If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
The correct information is:	
Applications	ALG: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If application-protocol ignore is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when application-protocol ignore is not configured for custom applications.

- In this guide, the **source-threshold** statement incorrectly shows a default value of 1024 per second for number in the Options section. The correct default value is 4000 per second.
- The Description for the **limit** (Security SCTP) option under the hierarchy [edit security gprs sctp profile profile-name], incorrectly states you can set the global messages rate limit per session. The description should say that you can set the rate limit for local Services Processing Unit (SPU) packets.
- In the "Class-of-Service Configuration Statement Hierarchy" section, the syntax for the forwarding-classes statement is incorrect. It states:

```
forwarding-classes {  
  class queue-num queue-number priority (high | low);  
  queue queue-number class-name priority (high | low) < policing-priority (premium  
  | normal) > < spu-priority (high | low) >;  
}
```

This syntax has been corrected as follows:

```
forwarding-classes {  
  class class-name queue-num queue-number <priority (high | low) > <spu-priority  
  (high | low) >;  
  queue queue-number <priority (high | low) >;  
}
```

- In the Security Hierarchy and Statements section, the edit security forwarding-options statements are missing the following:
 - A note indicating that packet-based processing is not supported on the following SRX Series devices: SRX1400, SRX3400, SRX3600, SRX5600, and SRX5800.
 - A link in the Related Topics section to the Junos OS Feature Support Reference for SRX Series and J Series Devices.

Junos OS Layer 2 Bridging and Switching Configuration Guide for Security Devices

- In this guide, the section "Configuring Layer 2 Bridging and Transparent Mode" includes an incorrect example, "Example: Configuring Layer 2 Trunk Interfaces with Multiple Units." The example is in error because the SRX Series devices do not support multiple units.

Junos OS Monitoring and Troubleshooting Guide for Security Devices

- In this guide, the watchdog disable option was removed from Step 3 in "Recovering the Root Password for SRX Series Devices" as it does not support SRX Series devices.
- In the "Example: Enabling Packet Capture on a Device" section, the CLI quick configuration shows an incorrect **set** command. The correct set command is **set forwarding-options packet-capture file filename pcap-file files 100 size 1024 world-readable**.

Junos OS MPLS Configuration Guide for Security Devices

- In this guide, section *VPLS Exceptions on J Series and SRX Series Devices* includes an incorrect information mentioned as follows:

LDP signaling is not supported.

In this guide, section *BGP Signaling* includes an incorrect note mentioned as follows:

LDP signaling is not supported for the VPLS routing instance.

Junos OS Routing Protocols and Policies Guide for Security Services

- Topic “Best Practices for Defining Policies on High-End SRX Series Devices”, shows SRX1400 devices supporting a maximum of 10,000 policies. However, SRX1400 devices actually support a maximum of 40,000 policies.

Junos OS Security Basics

- The topic Understanding Policy Application Timeouts Contingencies under **Security Basics > Security Policy Applications for Security Devices > Policy Application Timeout**, contains erroneous information. It should read as follows:

When setting timeouts, be aware of the following contingencies:

- If an application contains several application rule entries, all rule entries share the same timeout. You need to define the application timeout only once. For example, if you create an application with two rules, the following commands will set the timeout to 20 seconds for both rules:

```
user@host# set applications application test protocol tcp destination-port 1035-1035  
inactivity-timeout 20
```

```
user@host# set applications application test term test protocol udp
```

```
user@host# set applications application test term test source-port 1-65535
```

```
user@host# set applications application test term test destination-port 1111-1111
```

- If multiple custom applications are configured with custom timeouts, then each application will have its own custom application timeout. For example:

```
user@host# set applications application ftp-1 protocol tcp source-port 0-65535  
destination-port 2121-2121 inactivity-timeout 10
```

```
user@host# set applications application telnet-1 protocol tcp source-port 0-65535  
destination-port 2300-2348 inactivity-timeout 20
```

With this configuration, Junos OS applies a 10-second timeout for destination port **2121** and a 20-second timeout for destination port **2300** in an application group.

Junos OS Security Configuration Guide

In the “Understanding SIP ALGs and NAT” topic, information in the following sections is incorrect:

- **Call Re-INVITE Messages**

This section incorrectly states:

When one or more media sessions are removed from a call, pinholes are closed and bindings released just as with a BYE message.

The correct information is:

When all the media sessions or media pinholes are removed from a call, the call is removed when a BYE message is received.

- **Call Session Timers**

This section incorrectly states:

The SIP ALG uses the **session-expires** value to time out a session if a Re-INVITE or UPDATE message is not received. The ALG receives the **session-expires** value, if present, from the 200 OK responses to the INVITE and uses this value for signaling timeout. If the ALG receives another INVITE before the session times out, the ALG resets all timeout values to this new INVITE or to default values, and the process is repeated. As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist.

The correct information is (The **session-expires** value is not supported on SRX Series devices):

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist.

- **Table Requesting Messages with NAT Table**

This table incorrectly states:

Outbound Request (from private to public)	Route:	Replace ALG address with local address

The correct information is:

Outbound Request (from private to public)	Route:	Replace local address with ALG address

- In “Example: Configuring a Route-Based VPN,” the **show security zones** output for the SRX Series device erroneously shows host-inbound-traffic configured for the vpn-chicago zone; this configuration is not included in the CLI Quick Configuration and the Step-by-Step Procedure.
- In the “Example: Configuring an SCTP Profile” the following information is incorrect:

If the SCTP profile is removed from the policy:

- The firewall drops the data packet and allows only the control packet.

The correct information is:

If no profile is attached to a policy the SCTP data packets will be forwarded directly without any inspection by the SCTP module.

- In Example: Configuring a User Role Firewall on an SRX Series Device and Acquiring User Role Information from an Active Directory Authentication Server, the **redirect-url** option in step 2 of the redirection procedure is incorrect. The URL and variables should be enclosed in quotation marks.

[edit]

```
user@host# set services unified-access-control captive-portal acs-device
redirect-url "https://%ic-url%/?target=%dest-url%&enforcer=%enforcer-id%"
```

- In topic "Best Practices for Defining Policies on High-End SRX Series Devices", shows SRX1400 devices supporting a maximum of 10,000 policies. However, SRX1400 devices actually support a maximum of 40,000 policies.
- In "Example: Configuring AppTrack," of the *Junos OS Security Configuration Guide for Security Devices*, the **set security log mode stream** statement was omitted from the log configuration statements. The updated log configuration should read:

```
user@host# set security log mode stream
user@host# set security log format sd-syslog
user@host# set security log source-address 5.0.0.254
user@host# set security log stream app-track-logs host 5.0.0.1
```

- In Chapter 48, "AppTrack Application Tracking," of the *Junos OS Security Configuration Guide for Security Devices*, the operational command **show security application-tracking statistics applications** is incorrect. The operational command to display all application identification statistics, including application tracking statistics, is **show services application-identification statistics applications**.

The operational command **show application-tracking counters** is incorrect. The command should be **show security application-tracking counters**.

Junos OS Security Network Address Translation

- In Example: Configuring NAT for Multiple ISPs under Network Address Translation for Security Devices > Configuration > NAT for Multiple ISPs the statement **set routing-options rib-groups isp import-rib inet.0** was omitted from the configuration. The updated configuration should read:

```
set routing-options rib-groups isp import-rib inet.0
set routing-options rib-groups isp import-rib isp1.inet.0
set routing-options rib-groups isp import-rib isp2.inet.0
```

In addition, because zone based address-book for NAT rules is unsupported, you should not use the statements provided in the example; use global address book instead.

- The command **show security nat source persistent-nat-table** under **Network Address Translation > Administration > Source NAT Operational Commands** is:
 - Missing the option: **summary**—Display persistent NAT bindings summary.

- Contains incomplete sample output. The corrected sample output is as follows:

```
user@host> show security nat source persistent-nat-table internal-ip 9.9.9.1 internal-port 60784
```

Internal	Left_time/	Curr_Sess_Num/	Source	Reflective	Source	Type
In_IP	In_Port	I_Proto	Ref_IP	Ref_Port	R_Proto	NAT Pool
9.9.9.1	60784	udp	66.66.66.68	60784	udp	dynamic-customer-source
any-remote-host	254/300	0/30	105			

```
user@host> show security nat source persistent-nat-table all
```

Internal	Left_time/	Curr_Sess_Num/	Source	Reflective	Source	Type
In_IP	In_Port	I_Proto	Ref_IP	Ref_Port	R_Proto	NAT Pool
9.9.9.1	63893	tcp	66.66.66.68	63893	tcp	dynamic-customer-source
any-remote-host	192/300	0/30	105			
9.9.9.1	64014	udp	66.66.66.68	64014	udp	dynamic-customer-source
any-remote-host	244/300	0/30	105			
9.9.9.1	60784	udp	66.66.66.68	60784	udp	dynamic-customer-source
any-remote-host	254/300	0/30	105			
9.9.9.1	57022	udp	66.66.66.68	57022	udp	dynamic-customer-source
any-remote-host	264/300	0/30	105			
9.9.9.1	53009	udp	66.66.66.68	53009	udp	dynamic-customer-source
any-remote-host	268/300	0/30	105			
9.9.9.1	49225	udp	66.66.66.68	49225	udp	dynamic-customer-source
any-remote-host	272/300	0/30	105			
9.9.9.1	52150	udp	66.66.66.68	52150	udp	dynamic-customer-source
any-remote-host	274/300	0/30	105			
9.9.9.1	59770	udp	66.66.66.68	59770	udp	dynamic-customer-source
any-remote-host	278/300	0/30	105			
9.9.9.1	61497	udp	66.66.66.68	61497	udp	dynamic-customer-source
any-remote-host	282/300	0/30	105			
9.9.9.1	56843	udp	66.66.66.68	56843	udp	dynamic-customer-source
any-remote-host	-/300	1/30	105			

```
user@host> show security nat source persistent-nat-table summary
```

```
Persistent NAT Table Statistics on FPC5 PIC0:
binding total : 65536
binding in use : 0
enode total : 524288
enode in use : 0
```

Junos OS System Log Messages Reference

- The *AV System Log Messages* topic lists incorrect facilities for the systems logs.

On all SRX Series devices, antivirus (AV) system logs are generated with the facility **LOG_USER** or **LOG_DAEMON**.

[Table 16 on page 216](#) shows the correct facilities for the system logs.

Table 16: Antivirus System Logs

System Logs	Incorrect Facility	Correct Facility
AV_PATTERN_GET_FAILED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_KEY_EXPIRED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_KL_CHECK_FAILED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_TOO_BIG	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_UPDATED	LOG_FIREWALL	LOG_DAEMON
AV_PATTERN_WRITE_FS_FAILED	LOG_FIREWALL	LOG_DAEMON
AV_SCANNER_READY	LOG_FIREWALL	LOG_DAEMON
AV_VIRUS_DETECTED_MT	LOG_PFE	LOG_USER

Various Guides

- Some Junos OS user, reference, and configuration guides—for example the [Junos Software Routing Protocols Configuration Guide](#), [Junos OS CLI User Guide](#), and [Junos OS System Basics Configuration Guide](#)—mistakenly do not indicate SRX Series device support in the “Supported Platforms” list and other related support information; however, many of those documented Junos OS features are supported on SRX Series devices. For full, confirmed support information about SRX Series devices, please refer to the [Junos OS Feature Support Reference for SRX Series and J Series Devices](#).

Related Documentation

- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Known Limitations in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 174](#)
- [Outstanding Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 187](#)
- [Resolved Issues in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 188](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.1 for High-End SRX Series Services Gateways

This section includes the following topics:

- [Upgrading and Downgrading among Junos OS Releases on page 217](#)
- [Upgrading an AppSecure Device on page 219](#)
- [Upgrade and Downgrade Scripts for Address Book Configuration on page 219](#)
- [Upgrade Policy for Junos OS Extended End-Of-Life Releases on page 222](#)
- [Hardware Requirements for Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 222](#)

Upgrading and Downgrading among Junos OS Releases

All Junos OS releases are listed in sequence on the JUNOS Software Dates & Milestones webpage:

<http://www.juniper.net/support/eol/junos.html>

To help in understanding the examples that are presented in this section, a portion of that table is replicated here. Note that releases footnoted with a 1 are Extended End-of-Life (EOL) releases.

Product	FRS Date
Junos 12.1	03/28/2012
Junos 11.4 ¹	12/21/2011
Junos 11.3	08/15/2011
Junos 11.2	08/03/2011
Junos 11.1	03/29/2011
Junos 10.4 ¹	12/08/2010
Junos 10.3	08/15/2010
Junos 10.2	05/28/2010
Junos 10.1	02/15/2010
Junos 10.0 ¹	11/04/2009
Junos 9.6	08/06/2009
Junos 9.5	04/14/2009
Junos 9.4	02/11/2009
Junos 9.3 ¹	11/14/2008
Junos 9.2	08/12/2008
Junos 9.1	04/28/2008
Junos 9.0	02/15/2008
Junos 8.5 ¹	11/16/2007

You can directly upgrade or downgrade between any two Junos OS releases that are within three releases of each other.

- Example: Direct release upgrade

Release 10.3 → (*bypassing Releases 10.4 and 11.1*) Release 11.2

To upgrade or downgrade between Junos OS releases that are more than three releases apart, you can upgrade or downgrade first to an intermediate release that is within three releases of the desired release, and then upgrade or downgrade from that release to the desired release.

- Example: Multistep release downgrade

Release 11.3 → (*bypassing Releases 11.2 and 11.1*) Release 10.4 → Release 10.3

Juniper Networks has also provided an even more efficient method of upgrading and downgrading using the Junos OS EEOL releases. EEOL releases generally occur once a calendar year and can be more than three releases apart. For a list of, EEOL releases, go to <http://www.juniper.net/support/eol/junos.html>

You can directly upgrade or downgrade between any two Junos OS EEOL releases that are within three EEOL releases of each other.

- Example: Direct EEOL release upgrade

Release 9.3 (EEOL) → (*bypassing Releases 10.0 [EEOL] and 10.4 [EEOL]*) Release 11.4 (EEOL)

To upgrade or downgrade between Junos OS EEOL releases that are more than three EEOL releases apart, you can upgrade first to an intermediate EEOL release that is within three EEOL releases of the desired EEOL release, and then upgrade from that EEOL release to the desired EEOL release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 8.5 (EEOL) → (*bypassing Releases 9.3 [EEOL] and 10.0 [EEOL]*) Release 10.4 (EEOL) → Release 11.4 (EEOL)

You can even use a Junos OS EEOL release as an intermediate upgrade or downgrade step if your desired release is several releases later than your current release.

- Example: Multistep release upgrade using intermediate EEOL release

Release 9.6 → Release 10.0 (EEOL) → Release 10.2

For additional information about how to upgrade and downgrade, see the *Junos OS Installation and Upgrade Guide*.

Upgrading an AppSecure Device

Use the **no-validate** Option for AppSecure Devices.

For devices implementing AppSecure services, use the **no-validate** option when upgrading from Junos OS Release 11.2 or earlier to Junos OS 11.4R1 or later. The application signature package used with AppSecure services in previous releases has been moved from the configuration file to a signature database. This change in location can trigger an error during the validation step and interrupt the Junos OS upgrade. The no-validate option bypasses this step.

Upgrade and Downgrade Scripts for Address Book Configuration

Beginning with Junos OS Release 11.4 or later, you can configure address books under the **[security]** hierarchy and attach security zones to them (zone-attached configuration). In Junos OS Release 11.1 and earlier, address books were defined under the **[security zones]** hierarchy (zone-defined configuration).

You can either define all address books under the **[security]** hierarchy in a zone-attached configuration format or under the **[security zones]** hierarchy in a zone-defined configuration format; the CLI displays an error and fails to commit the configuration if you configure both configuration formats on one system.

Juniper Networks provides Junos operation scripts that allow you to work in either of the address book configuration formats (see [Figure 8 on page 221](#)).

- [About Upgrade and Downgrade Scripts on page 220](#)
- [Running Upgrade and Downgrade Scripts on page 221](#)

About Upgrade and Downgrade Scripts

After downloading the Junos OS Release 12.1, you have the following options for configuring the address book feature:

- **Use the default address book configuration**—You can configure address books using the zone-defined configuration format, which is available by default. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.
- **Use the upgrade script**—You can run the upgrade script available on the Juniper Networks support site to configure address books using the new zone-attached configuration format. When upgrading, the system uses the zone names to create address books. For example, addresses in the trust zone are created in an address book named **trust-address-book** and are attached to the trust zone. IP prefixes used in NAT rules remain unaffected.

After upgrading to the zone-attached address book configuration:

- You cannot configure address books using the zone-defined address book configuration format; the CLI displays an error and fails to commit.
- You cannot configure address books using the J-Web interface.


For information on how to configure zone-attached address books, see the Junos OS Release 11.4 documentation.

- **Use the downgrade script**—After upgrading to the zone-attached configuration, if you want to revert to the zone-defined configuration, use the downgrade script available on the Juniper Networks support site. For information on how to configure zone-defined address books, see the Junos OS Release 11.1 documentation.



NOTE: Before running the downgrade script, make sure to revert any configuration that uses addresses from the global address book.

Figure 8: Upgrade and Downgrade Scripts for Address Books



Running Upgrade and Downgrade Scripts

The following restrictions apply to the address book upgrade and downgrade scripts:

- The scripts cannot run unless the configuration on your system has been committed. Thus, if the zone-defined address book and zone-attached address book configurations are present on your system at the same time, the scripts will not run.
- The scripts cannot run when the global address book exists on your system.
- If you upgrade your device to Junos OS Release 11.4 or later and configure logical systems, the master logical system retains any previously configured zone-defined address book configuration. The master administrator can run the address book upgrade script to convert the existing zone-defined configuration to the zone-attached configuration. The upgrade script converts all zone-defined configurations in the master logical system and user logical systems.



NOTE: You cannot run the downgrade script on logical systems.

For information about implementing and executing Junos operation scripts, see the *Junos OS Configuration and Operations Automation Guide*.

Upgrade Policy for Junos OS 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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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 <http://www.juniper.net/support/eol/junos.html>.

Hardware Requirements for Junos OS Release 12.1 for High-End SRX Series Services Gateways

Transceiver Compatibility for SRX Series Devices

We strongly recommend that only transceivers provided by Juniper Networks be used on high-end SRX Series Services Gateways interface modules. Different transceiver types (long-range, short-range, copper, and others) can be used together on multiport SFP interface modules as long as they are provided by Juniper Networks. We cannot guarantee that the interface module will operate correctly if third-party transceivers are used.

Please contact Juniper Networks for the correct transceiver part number for your device.

Related Documentation

- [New Features in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 154](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 210](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.1 for High-End SRX Series Services Gateways on page 168](#)

Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers

- [New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 223](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 284](#)
- [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 298](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 404](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 432](#)

New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers

The following features have been added to Junos OS Release 12.1. Following the description is the title of the manual or manuals to consult for further information.

- [Class of Service on page 223](#)
- [High Availability on page 231](#)
- [Interfaces and Chassis on page 231](#)
- [Junos OS XML API and Scripting on page 248](#)
- [Layer 2 Ethernet Services on page 249](#)
- [MPLS Applications on page 250](#)
- [Multicast on page 253](#)
- [Network Management on page 254](#)
- [Routing Protocols on page 256](#)
- [Subscriber Access Management on page 262](#)
- [System Logging on page 273](#)
- [User Interface and Configuration on page 280](#)
- [VPNs on page 282](#)

Class of Service

- **Support for set forwarding class and DSCP value (MX Series routers with MPC/MIC interfaces)**—The set forwarding class and DSCP value for Routing Engine generated traffic is supported on MX Series routers with MPC/MIC interfaces. For example, use **b100110** instead of **100110**. This notation is applicable in all places where a binary DSCP value is specified under the **[edit firewall]** hierarchy level.

[Class of Service]

- **Class-of-service features on ATM MICs (MX Series routers)**—The following class-of-service features are supported on an ATM MIC:

- Traffic shaping and scheduling—Traffic shaping determines the maximum amount of traffic that can be transmitted on an interface. To configure traffic shaping and scheduling profile for ATM MICs, you must configure the service category by including the **atm-service** statement at the following hierarchy level:

[edit class-of-service traffic-control-profiles *traffic-control-profile-name*]

You can configure three different categories of ATM service: constant bit rate (CBR), non-real-time variable bit rate (NRVBR), and real-time variable bit rate (RTVBR). The service category works in conjunction with ATM cell parameters **peak-rate**, **sustained-rate**, and **max burst-size** to impose traffic shaping, transmit-rate, shaping-rate, and default excess-rate for an ATM queue.

- Policing—Policing, or rate limiting, enables you to limit the amount of traffic that passes into or out of the interface. It works with firewall filters to thwart denial-of-service (DoS) attacks. You can enable the input or output transmission rate of ATM traffic by including the **atm-policer** statement at the following hierarchy level:

[edit firewall]

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. The ATM policer controls the maximum rate of traffic sent or received on the interface on which it is applied. To apply the policer at the interface level, you must include the **atm-policer** statement at the following hierarchy level:

[edit interface *at-fpc/pic/port unit unit number*]

To apply limits to the traffic flow, configure the **cdvt** and **peak-rate** parameters within the policer. Define the **policing-action** parameter as **discard**, **discard-tag**, and **count** to set a consequence for the packets that exceed these limits. The consequence is usually a higher loss priority so that if the packets encounter downstream congestion, they are discarded first.

[Class of Service]

- **Policer support for aggregated Ethernet bundles (MX Series routers with MPC/MIC interfaces)**—Aggregated interfaces support single-rate policers, single-rate three-color marking policers, two-rate three-color marking policers, and hierarchical policers. By default, policer bandwidth and burst size applied on aggregated bundles are not matched to the user-configured bandwidth and burst size. Because an aggregated Ethernet interface is a bundle of Ethernet links of the same speed, if the user-configured bandwidth on aggregated bundles is 40 Mbps, each link has 40 Mbps. As a result, the effective bandwidth and burst size available to the aggregated interface are a lot higher than the value configured.

You can configure interface-specific policers applied on an aggregated Ethernet bundle to match the effective bandwidth and burst size to user-configured values.

To configure this feature, include the **shared-bandwidth-policer** statement at the following hierarchy levels:

[edit firewall policer *policer-name*]

[edit firewall three-color-policer *policer-name*]

[edit firewall hierarchical-policer *policer-name*]

This capability applies to all interface-specific policers of the following types: single-rate policers, single-rate three-color marking policers, two-rate three-color marking policers, and hierarchical policers. This capability does not apply to policers, hierarchical policers, and three-color policers used inside filters that are not interface-specific. It also does not apply to implicit policers and prefix-specific action policers.

[Class of Service]

- **Class-of-service features supported on T4000 Core Router**—Starting with Junos OS Release 12.1R1, the following class-of-service (CoS) features are supported on the T4000 Core Router:
 - Behavior aggregate (BA) classifiers:
 - Differentiated Services code point (DSCP) for IP DiffServ
 - DSCP for IPv6 DiffServ
 - IP precedence
 - MPLS EXP
 - IEEE 802.1p CoS
 - IEEE 802.1ad drop eligibility indicator (DEI)
 - Fixed classification—You can configure fixed classification on a logical interface by specifying a forwarding class to be applied to all packets received by the logical interface, regardless of the packet contents.



NOTE: On the T4000 Core Router, BA classification and fixed classification are mutually exclusive. That is, you can configure either BA classification or fixed classification.

- Tricolor marking (TCM)—By default, TCM is enabled. Therefore, you can configure only the **any** option for the drop profile at the [edit class-of-service schedulers *scheduler-name*] hierarchy level.
- Layer 2 rewrite rules:
 - IEEE 802.1p CoS bit rewrite
 - IEEE 802.1ad DEI rewrite
- A maximum of 16 forwarding classes and 4 types of packet loss priorities: **low**, **high**, **medium-low**, and **medium-high**.
- Low and high levels of fabric queuing priorities.
- Default scheduler—By default, the best-effort forwarding class receives 95 percent of the bandwidth and buffer space for the output link, and the network control forwarding class receives 5 percent. The default drop profile causes the buffer to fill and then discard all packets until it has space.

- The lowest of the scaling numbers (classifiers, rewrite rules, and weighted random early detection (WRED)) are associated with MX Series and T Series routers.
- On the T4000 Type 5 FPC, excess bandwidth is shared in the ratio of the transmit rates. This distribution can be updated by configuring the **excess-rate** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level. You can specify the excess rate sharing by percentage or by proportion.



NOTE: In Junos OS Release 12.1R1, the following features are not supported on the T4000 Type 5 FPC :

- Shaping at the physical interface level
- Scheduling and shaping at the logical interface level
- Queue-level rate limiting applied at the logical interface, physical interface, and logical interface set levels
- Layer 3 rewrite rules
- CoS on aggregated Ethernet interface

[Class of Service]

- **Extends filter and policer feature support to T4000 Type 5 FPC (T4000-FPC5-3D)**—The following filter and policer features supported on the T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES) are also supported on the T4000 Type 5 FPC (T4000-FPC5-3D):
 - Label-switched path (LSP) policers
 - Address Resolution Protocol (ARP) policers
 - Tricolor marking policers
 - Forwarding table filters
 - Filter-based forwarding
 - Prefix-specific actions
 - Sampling and port mirroring features

The following filter and policer features supported on the T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES) are not supported on the T4000 Type 5 FPC (T4000-FPC5-3D):

- Service PIC–related filters.
- Logical interface policer as filter action.
- Physical interface policers.
- Applying a policer at the logical interface level.
- Hierarchical policers.
- Filter actions such as **ipsec-sa**, **service-accounting**, and **service-filter-hit**.

- The **dscp 0** action during the interoperation between a T1600 Enhanced Scaling Type 4 FPC and a T4000 Type 5 FPC.
- Shared bandwidth policer.
- A filter attached at the Layer 2 application point (that is, at the logical interface level) is unable to match with the forwarding class of a packet that is set by a Layer 3 classifier such as DSCP, DSCP V6, **inet-precedence**, or **mpls-exp**.
- Using **interface-group** and **interface-group-except** as match conditions for the VPLS family filter.
- The ability to filter MPLS-tagged IPv4 packets based on IP parameters.
- Applying filters at **set interfaces lo0 unit 0 family any filter input filter-name**.
- For a three-color policer operating in color-aware mode and when the PLP of the input packet is medium-low, the color of the input packet to the policer is mapped to the color yellow.

In such a scenario, if the color of the input packet remains unchanged, the policer operates in the following way:

- On a T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES), the packet loss priority (PLP) of the output packet remains medium-low.
- On a T4000 Type 5 FPC (T4000-FPC5-3D), the PLP of the output packet is marked as medium-high.

Because of this difference, for any applications (such as rewrite and WRED selection on egress interface) that use PLP, the packets are treated differently for the same flow depending on the FPC type (T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES) or T4000 FPC5 (T4000-FPC5-3D)) on which the policer is applied.

[Firewall Filters and Traffic Policers]

- **Extends support for tunnel services features on T4000 Type 5 FPC (T4000-FPC5-3D)**—Starting with Junos OS Release 12.1, all the tunnel services features supported on Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs) on MX Series routers are supported on the T4000 Type 5 FPC.

[Services Interfaces, System Basics]

- **Set IPv6 DiffServ code point (DSCP) and MPLS EXP independently (MX Series routers with MPC/MIC interfaces)**—You can set the packet DSCP and MPLS EXP bits independently on IPv6 packets with MPC/MIC interfaces. To enable this feature, include the **protocol mpls** statement at the **[edit class-of-service interfaces interface-name unit logical-unit rewrite-rules dscp-ipv6 rule-name]** hierarchy level.

You can set DSCP IPv6 values only at the ingress MPLS node.

[Class of Service]

- **Class-of-service features supported** — Starting with Junos OS Release 12.1R2, the following class-of-service (CoS) features are supported on the T4000 Core Routers with Type 5 FPCs:

Layer 3 rewrite:

- IPv4 DSCP rewrite
- IPv4 INET precedence rewrite
- IPv6 DSCP rewrite
- MPLS EXP rewrite
- Simultaneous MPLS EXP and IPv4 precedence rewrite

In the case of Layer 3 VPN, Layer 2 VPN, or VPLS, the following rules apply to simultaneous MPLS EXP and IPv4 precedence rewrite operation, under the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules *exp rewrite-rule-name* protocol *protocol-types*]** hierarchy.

The protocol statement defines the types of MPLS packets and possible configurations for the following options:

- mpls
- mpls-inet-both
- mpls-inet-both-non-vpn



NOTE: For Layer 3 VPN, Layer 2 VPN, and VPLS, mpls-inet-both is not supported on T4000 routers.

Aggregated Ethernet:

- All CoS mechanisms that are supported on regular interfaces are supported on bundles.
- CoS with member links can be on different Packet Forwarding Engines and line cards.

Scheduling:

- Physical interface scheduling (eight queues per port)
- Four packet loss priority levels
- Unused bandwidth sharing among queues
- **Switch fabric fault management for T4000 routers**—The T4000 router consists of a Switch Interface Board (SIB) with fabric bandwidth double the capacity of the T1600 router. The fabric fault management functionality is similar to that in T1600 routers.

The fabric fault management functionality involves monitoring all high-speed links connected to the fabric and the ones within the fabric core for link failures and link errors. Action is taken based on the fault and its location. The actions include:

- Reporting link errors in system log files and sending this information to the Routing Engine.
- Reporting link failures at the Flexible Port Concentrator (FPC) or at the SIB and sending this information to the Routing Engine.

- Marking a SIB in **Check** state.
- Moving a SIB into **Fault** state.

The following are the high-level indications of fabric faults that are monitored by Junos OS:

- An SNMP trap is generated whenever a SIB is reported as **Check** or **Fault**.
- **show chassis alarms**—Indicates that a SIB is in **Check** or **Fault** state.
- **show chassis sibs**—Indicates that a SIB is in **Check** or **Fault** state or that a SIB is in **Offline** state when the SIB initializes (this occurs when the SIB does not power on fully).
- **show chassis fabric fpcs**—Indicates whether any fabric links are in error on the FPCs' side.
- **show chassis fabric sibs**—Indicates whether any fabric links are in error on the SIBs' side.
- The `/var/log/messages` system log messages file at the Routing Engine has error messages with the prefix **CHASSISD_FM_ERROR**.
- The SIBs display the **FAIL** LED.

[*System Basics*]

- **Interoperability of Type 3 FPCs and Type 4 FPCs with Type 5 FPCs (T4000 routers)**—Support for interoperability of T640 Enhanced Scaling FPC3, T1600 Enhanced Scaling FPC4, and T640 Enhanced Scaling FPC4-IP with T4000 FPC5 is possible with fabric notification translation. This feature is supported on T4000 routers.

Basic packet forwarding, IPv4, IPv6, MPLS, and multicast (dataplane) are supported through this feature.

- **Support for class-of-service features to ensure quality of service for real-time traffic that is sensitive to latency on a network (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**—The Application Services Modular Line Card (AS MLC) supports the following CoS features on MX240, MX480, and MX960 routers:
 - **Code-point aliases**—A code-point alias is a meaningful name that can be associated with CoS values such as Differentiated Services code points (DSCPs), DSCP IPv6, IP precedence, IEEE 802.1, and MPLS experimental (EXP) bits that can then be used while configuring CoS components.
 - **Classification**—Packet classification associates the packet with a particular CoS servicing level. In Junos OS, classifiers associate incoming packets with a forwarding class and loss priority and, based on the associated forwarding class, assign packets to output queues.
 - **Behavior aggregate**—A method of classification that operates on a packet as it enters the router.
 - **Multifield classification**—A method of classification that can examine multiple fields in the packet.

- Fixed classification—A method of classification that refers to the association of a forwarding class with a packet regardless of its packet contents.
- Scheduling—Schedulers are used to define the properties of output queues. On the AS Modular Carrier Card (AS MCC), the following scheduling features are supported (physical interfaces only):
 - Buffer sizes
 - Delay buffer size
 - Drop profile map
 - Excess priority
 - Excess rate percentage
 - Output-traffic-control profile
 - Priority
 - Scheduler-map
 - Shaping rate
 - Transmit rate
 - WRED rules

[*Class-of-Service*]

- **Ingress CoS on MPC/MIC interfaces (MX Series routers)**—You can apply CoS or hierarchical schedulers on the ingress side of MPC/MIC interfaces. The input and output CoS parameters are independent in most cases.
- **Enhancements to scheduler configuration on FRF.16 physical interfaces**—Starting with Release 12.1R4, Junos OS extends the class-of-service scheduler support on FRF.16 physical interfaces to the **excess-rate**, **excess-priority**, and **drop-profile-map** configurations. The **excess-rate**, **excess-priority**, and **drop-profile-map** statements are configured at the [**edit class-of-service schedulers scheduler-name**] hierarchy level.
 - Support for the **drop-profile-map** configuration enables you to configure random early detection (RED) on FRF.16 bundle physical interfaces.
 - Support for the **excess-rate** configuration enables you to specify the percentage of the excess bandwidth traffic to share.
 - Support for the **excess-priority** configuration enables you to specify the priority for excess bandwidth traffic on a scheduler.

This feature is supported only on multiservices PICs installed on MX Series routers.

- **Accurate reporting of output counters for MLFR UNI NNI bundles**—Starting with Release 12.1R4, Junos OS reports the actual output counters in the multilink frame relay (MLFR) UNI NNI bundle statistics section of the **show interfaces lsq-interface statistics** command output. From this release on, Junos OS also provides per-DLCI counters for logical interfaces. In earlier releases, there was a discrepancy between the actual output counters and the reported value because of errors in calculating the output

counters at the logical interface level. That is, at the logical interface level, the output counter was calculated as the sum of frames egressing at the member links instead of providing the output counter as the sum of per-DLCI output frames.

High Availability

- **Support extended for Layer 3 features (MX Series routers with MPC/MIC interfaces)**—Junos OS Release 12.1 extends support for the following Layer 3 features on MX Series routers with MPC/MIC interfaces:
 - **Pseudowire redundancy**—Enables you to configure redundant Layer 2 circuit pseudowires between devices. Layer 2 circuit and VPLS services can be maintained between devices connected using pseudowires in the network even after certain failures in the control or data plane. Backup pseudowires can be configured between Layer 2 devices in the customer's network and PE routers within the service provider's network.
 - **Distributed PPM support for LACP**—Enables you to switch between distributed and centralized periodic packet management (PPM). By default, distributed PPM is active.
 - **External/Internal BGP VPN load balancing and egress filtering support**—Enables you to load-balance traffic across external and internal BGP paths and simultaneously configure egress filters and policers on the VRF interfaces.
 - **Egress filtering of PIMv4/v6 messages**—Enables you to filter PIM join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or a combination of these values. This is useful when the core of your network is using a mix of IP and MPLS. You can use this feature to selectively filter PIM join and prune messages and forward them to PIM neighbors.

These features can interoperate between an MPC and a DPC when both are present on the same MX Series router.

[Services Interfaces, High Availability, Network Interfaces, Multicast]

- **Support for graceful Routing Engine switchover (GRES) on T4000 routers**—Starting with Junos OS Release 12.1R2, GRES is supported on T4000 routers. GRES enables a T4000 router with redundant Routing Engines to continue forwarding packets even if one Routing Engine fails.

[High Availability]

Interfaces and Chassis

- **T4000 Core Router**—The T4000 Core Router has a capacity of up to 2000 gigabits per second (Gbps), full duplex (4000 Gbps of any-to-any, nonblocking, half-duplex) switching.

The T4000 router features the following new hardware components:

- Redundant, load-sharing, six-input DC power supplies (PWR-T-6-60-DC)
- Routing Engine (RE-DUO-C1800-16G)

- T4000 SIBs (SIB-I-T4000)
- T Series craft interface (CRAFT-T-SERIES)
- Front fan trays (FANTRAY-T4000)
- T4000 FPC5 (T4000-FPC5-3D)
- 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
- 100-Gigabit Ethernet PIC with CFP

The T4000 router supports the following legacy hardware components:

- Control Board (CB-LCC)
- Routing Engine (RE-DUO-C1800-8G)
- Rear fan tray (FAN-REAR-TXP-LCC)
- Enhanced Scaling FPC3 (T640-FPC3-ES)
- Enhanced Scaling FPC4-1P (T640-FPC4-1P-ES)
- T1600 Enhanced Scaling FPC4 (T1600-FPC4)
- PICs
 - SONET/SDH OC192/STM64 PIC with XFP (PD-4OC192-SON-XFP)
 - SONET/SDH OC768c/STM256 PIC (PD-1OC768-SON-SR)
 - 10-port Gigabit Ethernet PIC with SFP (PC-10GE-SFP)
 - Gigabit Ethernet IQ2 PIC with SFP (PC-8GE-TYPE3-SFP-IQ2)
 - Gigabit Ethernet IQ2E PIC with SFP (PC-8GE-TYPE3-SFP-IQ2)
 - 10-Gigabit Ethernet DWDM OTN PIC (PC-1XGE-DWDM-OTN)
 - 10-Gigabit Ethernet IQ2 PIC with SFP (PC-1XGE-TYPE3-XFP-IQ2)
 - 10-Gigabit Ethernet IQ2E PIC with SFP (PC-1XGE-TYPE3-XFP-IQ2E)
 - 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)
 - 10-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP)
 - 100-Gigabit Ethernet PIC (PD-1CE-CFP)
 - Multiservices 500 PIC (PB-MS-500-3)
 - Tunnel Services PIC (PC-TUNNEL)
 - 4-port SONET/SDH OC48C/STM16 PIC with SFP (PC-4OC48-SON-SFP)
 - SONET/SDH OC48C/STM16 EOL PIC (PC-4OC48-SON-SMSR)
 - SONET/SDH OC192/STM64 EOL PIC (PC-1OC192-SON-SR2)

See the *T4000 PIC Guide* for a list of the supported PICs. See the *T4000 Hardware Guide* for a list of the supported hardware, and the procedure to upgrade to a T4000 router.

You can perform an upgrade from a T640 router or T1600 router. The upgrade procedure requires that you power off the T640 or T1600 router during the upgrade.

The T4000 upgrade kit includes five SIBs, two power supplies, two front fan trays, one rear fan tray, and one craft interface. The components can also be ordered individually. The T4000 FPC5 is not included in the upgrade kit, but can be installed after all components of the upgrade kit are installed and operational.



NOTE: Upgrading to a T4000 router is not currently supported from T640 routers connected to a TX Matrix platform or T1600 routers connected to a TX Matrix Plus platform.

[*T4000 Hardware Guide, T4000 PIC Guide*]

- **Optical transceiver support for T640 and T1600 routers**—Starting in Junos OS Release 12.1, the 10-Gigabit Ethernet LAN/WAN PICs with SFP+ in T640 and T1600 routers support the SFPP-10GE-ER transceivers.
- **Sanity polling for FPCs on T Series routers**—Sanity polling is supported for FPCs on T Series routers. You can configure the **sanity-poll** statement for a particular FPC to start a periodic sanity check for error conditions in the FPC.



NOTE: Currently, periodic sanity check is performed only on the routing chip register.

Sanity polling is not supported on FPC5.

You can configure the **sanity-poll** statement for the FPC at the [**edit chassis fpc slot-number**] hierarchy level. On a TX Matrix or TX Matrix Plus router, you can configure the statement at the [**edit chassis lcc number fpc number**] hierarchy level.

The **sanity-poll** statement detects an error condition and generates an emergency system log message in the FPC. You can configure the **retry-count** statement to perform rechecks for a specified number of times after detecting an error. If you do not configure the **retry-count** statement, then by default, the **sanity-poll** statement checks the detected error 10 times for a particular FPC.

If an error persists after all rechecks, sanity polling reports an error and takes appropriate actions. You can configure the **on-error** statement to perform the appropriate actions:

- **raise-alarm** generates the chassis alarm.
- **power cycle** reboots the FPC after generating a core file.
- **power off** halts the FPC, which is useful in case of permanent hardware failure.
- **write-coredump** triggers the core file.

[*System Basics*]

- **Operational mode command to display system-wide memory usage**—The **show system memory** command displays system-wide memory distribution and usage

including the Junos OS kernel, software processes, and memory disks. Use the **show system memory** command for troubleshooting with Juniper Networks Customer Support.

[*System Basics and Services Command Reference*]

- **Display IPv6 statistics for MLPPP bundles**—Starting with Junos OS Release 12.1, the **show interfaces lsq-fpc/pic/port** command displays the packet and byte counters for IPv6 data for Multilink Point-to-Point Protocol (MLPPP) bundles on link services intelligent queuing (LSQ) interfaces.

[*Interfaces Command Reference*]

- **100-Gigabit Ethernet MIC (MIC3-3D-1X100GE-CFP) supports interoperability with the 100-Gigabit Ethernet PIC (Type 4 1X100GE PIC for Enhanced Scaling FPC4) using SA multicast mode (MX Series routers with MPC/MIC interfaces)**—You can configure the 100-Gigabit Ethernet MIC (MIC3-3D-1X100GE-CFP) to interoperate with routers that use the 100-Gigabit Ethernet PIC (Type 4 1X100GE PIC for STFPC4 FPC) by using the **forwarding-mode** statement with the **sa-multicast** option at the [**edit chassis fpc slot pic slot**] hierarchy level. On egress, the router sets the SA multicast bit on the outgoing packets. At the other end, the 100-Gigabit Ethernet PIC detects the multicast bit and steers the packets to the appropriate Packet Forwarding Engine, PFE0 or PFE1. The ingress packet flow is the traffic flowing from the 100-Gigabit Ethernet PIC to the 100-Gigabit Ethernet MIC. There is no explicit CLI configuration required on the 100-Gigabit Ethernet PIC to enable this mode.

[*MX Series 3D Universal Edge Router Line Card Guide, Network Interfaces, System Basics*]

- **IPv6 support for inline sampling**—Starting with Junos OS Release 12.1, all MX Series routers with Modular Port Concentrators (MPCs) support monitoring and sampling services inline for IPv6 packets. To configure inline sampling for IPv6, include the **inline-jflow** statement at the [**edit forwarding-options sampling instance instance-name family inet6 output**] hierarchy level. Inline sampling exclusively supports a new format called IP_FIX that uses UDP as the transport protocol. When you configure inline sampling, include the **version-ipfix** statement at the [**edit forwarding-options sampling instance instance-name family inet6 output flow-server address**] hierarchy level.

New CLI options have been introduced to improve sampling for both IPv4 and IPv6 by enabling customers to set the size for the IPv4 and IPv6 hash tables based on their requirements. The hash table sizes are configured in units of 256K (256*1024). The range of acceptable values is 1 through 15 units.



NOTE: The maximum value of the *sum* of IPv4 and IPv6 hash tables sizes is 15 units.

To configure hash table sizes for IPv4 and IPv6, include one or both of the following statements at the [**edit chassis fpc fpc inline-services flow-table-size**] hierarchy level:

- **ipv4-flow-table-size table-size**
- **ipv6-flow-table-size table-size**

If you do not configure the hash table sizes, the following default values are used: 3840k (15 units) for IPv4 and 6k for IPv6.



NOTE: Changes to configured table sizes trigger a reboot of the FPC because flow table sizes are configured during FPC initialization.

[*Services Interfaces*]

- **Support for multilink-based protocols on channelized MICs (MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.1, multilink-based protocols are extended to the following channelized Modular Interface Cards (MICs) on MX240, MX480, and MX960 routers:
 - 4-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-4CHOC3-2CHOC12)
 - 8-port Channelized SONET/SDH OC3/STM1 (Multi-Rate) MIC with SFP (MIC-3D-8CHOC3-4CHOC12)
 - 8-port Channelized DS3/E3 MIC (MIC-3D-8CHDS3-E3-B)

The following encapsulations and protocols are also supported on the aforementioned MICs:

- Multilink Point-to-Point Protocol (MLPPP)
- Multiclass MLPPP
- Multilink Frame Relay (MLFR) end-to-end (FRF.15)
- Multilink Frame Relay (MLFR) UNI NNI (FRF.16) (also referred to as MFR)
- Compressed Real-Time Transport Protocol (CRTP)

[*Services Interfaces*]

- **Support for FlowTapLite**—Starting with Junos OS Release 11.2R3, all MX Series routers with MPC/MIC interfaces support FlowTapLite for IPv4 and IPv6.
- **ATM PWE3 support on ATM MICs with SFP (MX Series routers)**—The new ATM MIC (model number: MIC-3D-8OC3-2OC12-ATM) enables support for ATM Pseudowire Emulation Edge to Edge (PWE3) on MX Series routers. The MIC is rate-selectable at the following rates: 2 OC12 ports or 8 OC3 ports.

The ATM MIC with SFP is supported on the following MPCs:

- MPC1 Q (MX-MPC1-3D-Q)
- MPC2 Q (MX-MPC2-3D-Q)
- MPC2 EQ (MX-MPC2-3D-EQ)

The following features are supported on the ATM MIC (model number: MIC-3D-8OC3-2OC12-ATM) with SFP:

- Default framing mode on all ports is SONET. The MIC supports both SONET and SDH framing mode. The mode can be set at the MIC level or at the port level. To

enable SONET or SDH framing at the port level, you need to set the framing statement at the **[chassis fpc MPC-slot-number pic MIC-slot-number port port-number]** hierarchy level. To enable SONET or SDH framing at the MIC level, you must set the framing statement at the **[chassis fpc MPC-slot-number pic MIC-slot-number]** hierarchy level.

- ATM pseudowire encapsulation. The pseudowire encapsulation can be either cell-relay or AAL5 transport mode. Both modes enable sending of ATM cells between the MIC and a Layer 2 network.
- Cell relay VPI/VCI swapping. The ATM MIC can overwrite the values for VPI and VCI on both ingress and egress. The ATM MIC can also pass the value transparently (no-rewrite).

To configure the ATM MIC to modify both the VPI and VCI header values on both ingress and egress, you must specify the **psn-vci** statement at the following hierarchy level:

[edit interface at-interface-name/pic/port unit logical-unit-number]



NOTE: Cell relay VPI/VCI swapping on both ingress and egress is not compatible with the ATM policing feature.

To configure the ATM MIC to modify only the VPI values on both ingress and egress, you must specify the **psn-vpi** statement at the following hierarchy level:

[edit interface at-interface-name/pic/port unit logical-unit-number]



NOTE: Cell relay VPI swapping on both ingress and egress is not compatible with the ATM policing feature.

To configure the ATM MIC to pass the value transparently, you must specify the **no-vpivci-swapping** statement at the following hierarchy level:

[edit interface at-interface-name/pic/port unit logical-unit-number]

If none of the configuration statements mentioned earlier are included, for VP pseudowires, VPI values are modified on egress. For VC pseudowires, both VPI and VCI values are modified on egress. The ATM policing feature is compatible with cell relay VPI/VCI swapping on egress.

[Network Interfaces]

- **Support for unicast RPF loose mode (T Series routers)**—Extends support for unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard interface, on the Type 1, Type 2, and Type 3 FPCs on T Series routers. This feature, in conjunction with Remote Triggered Black Hole (RTBH) filtering, provides a mechanism to discard packets from untrusted sources. BGP policies in edge routers ensure that packets with untrusted source addresses have their next hop set to a discard route. When a packet arrives at the router with an untrusted source address, unicast RPF performs a route lookup of the

source address. Because the source address route points to a discard next hop, the packet is dropped. This feature is supported only on the IPv4 (**inet**) address family.

To configure unicast RPF loose mode with the ability to discard packets, you can use the **rpf-loose-mode-discard inet** statement at the **[edit forwarding options]** hierarchy level. Use the **show interfaces extensive** operational mode command to view the packet drops.

[*Network Interfaces*]

- **Static mapping for port forwarding**—In Junos OS Release 12.1, you can configure port forwarding without translation of destination addresses.



NOTE: You can configure port forwarding without translating destination ports when you are using address translation by specifying that the translated port and destination port are the same.

Port forwarding provides translation of the address, or port, or both address and port of a packet to a new destination. The translation facilitates reaching a host within a masqueraded, typically private, network, based on the port number on which it was received from the originating host. Port forwarding allows remote computers, such as public machines on the Internet, to connect to a nonstandard port (port other than 80) of a specific computer within a private network. An example of this type of destination is the host of a public HTTP server within a private network. Port forwarding supports only IPv4 addresses.

To configure port forwarding with port translation only:

- Include the **destined-port port-id translated-port port-id** statement at the **[edit services nat port-forwarding map-id]** hierarchy level. You can specify up to 32 port mappings under a single map ID.
- Include the following statements at the **[edit services nat rule rule-name term term-name then]** hierarchy level:
 - **port-mappings map-name**
 - **no-translation**

[*Services Interfaces, Next-Generation Network Addressing, Systems Basics and Services Command Reference*]

- **Configuring parameters for offloading flows**—Starting with Junos OS Release 12.1, you can set the parameters for flow offloading by configuring the **set trio-flow-offload minimum bytes** and **set trio-flow-offload minimum-age** statements under the **[edit interfaces ms-fpc/pic/port service-options]** hierarchy level. Offloading is supported on all MX Series routers with Modular Port Concentrator (MPCs)/Modular Interface Cards (MICs). The configuration allows any plug-in or daemon on a PIC to generate a flow offload request and offload the flows to the Packet Forwarding Engine.

The **show services sessions** command displays flow offload status for each session.

[*Services Interfaces*]

- **MPC3E (Ethernet 3D Modular Port Concentrator) with two separate Modular Interface Card (MIC) slots that support two MICs (new MIC-3D-1X100GE-CFP and legacy MIC-3D-20GE-SFP) on MX Series routers**—MX960, MX480, and MX240 routers support the MPC3E with two MIC slots. The supported MICs are MIC-3D-1X100GE-CFP and MIC-3D-20GE-SFP, which are field-replaceable units (FRUs). The MPC contains the Packet Forwarding Engine that steers traffic entering through the customer's Ethernet interfaces toward the egress interfaces using fabric on the Switch Control Board. The MPC is inserted into a slot in a router. MICs provide the physical interface and are installed into the MPCs. The MPC3E requires the Enhanced MX Switch Control Board for fabric redundancy. You can also continue to use existing SCBs without fabric redundancy. The MPC interoperates with existing MX Series line cards, including Dense Port Concentrators (DPCs) and Modular Port Concentrators (MPCs).

The MPC3E is based on a new Junos OS chipset for increased scalability for bandwidth, subscribers, and service capabilities of the routers.

The following are the key features of the MPC3E:

- Supports 100-Gigabit Ethernet interfaces
- Supports two separate slots for MICs (MIC-3D-1X100GE-CFP or MIC-3D-20GE-SFP)
- Supports one 100-Gigabit Ethernet port per MIC
- Supports up to 200 Gbps aggregate WAN bandwidth connectivity for the two MIC slots; the line card is oversubscribed in the ratio of 1.5:1
- Supports up to four full-duplex tunnel interfaces for each Packet Forwarding Engine
- Supports intelligent oversubscription services

The MPC3E supports feature parity with the following Junos OS Release 10.4 software features:

- Basic Layer 2 features and virtual private LAN service (VPLS) functionality
- Layer 3 routing protocols
- MPLS
- Multicast forwarding
- Firewall filters and policers
- Class-of-service (CoS) support
- Tunnel support
- Interoperability with existing DPCs and MPCs

The following features are not supported on the MPC3E:

- Fine-grained queuing and input queuing
- Unified in-service software upgrade (ISSU)
- Multilink services

- Internet Group Management Protocol (IGMP) snooping with bridging, integrated routing and bridging (IRB), or VPLS
- Intelligent hierarchical policers
- Layer 2 trunk port
- MPLS fast reroute (FRR) VPLS instance prioritization
- Precision Time Protocol (IEEE 1588)
- Synchronous Ethernet
- J-Flow monitoring and services
- Virtual Chassis support

For more information about the supported and unsupported Junos OS software features for this MPC, see "Protocols and Applications Supported by MX Series MPCs" in the *MX Series Line Card Guide*.

[*MX Series Line Card*]

- **LAG/LACP for 10-port 10-Gigabit Ethernet PICs**—10-port 10-Gigabit Ethernet PICs support link aggregation from the following Type 3 10-Gigabit Ethernet PICs: 1x10-Gigabit Ethernet IQ2, 1x10-Gigabit Ethernet IQ2E, and 10-Gigabit Ethernet-XENPAK. For bandwidth aggregation, load sharing, and link protection, LAG can be enabled. After aggregated Ethernet is enabled, LACP protocol forms an aggregated bundle of member links.

[*Network Interfaces*]

- **SIB to support fabric bandwidth for T4000 router**—The T4000 router is an upgraded version of the T1600 router or the T640 router. The T4000 router consists of a Switch Interface Board (SIB) with fabric bandwidth double the capacity of the T1600 router. As a result of the SIB to support the fabric bandwidth for the T4000 router, the output of the **show chassis fabric topology sib-slot** command is changed. This command displays the connectivity and the link status between the Packet Forwarding Engine and the SIB. The following commands related to switch fabric management are also supported on T4000 routers:
 - **show chassis fabric sibs**—Displays the state of the electrical switch fabric links between the SIB and the Packet Forwarding Engine.
 - **show chassis fabric fpcs**—Displays the state of the electrical switch fabric links between the Flexible PIC Concentrators (FPCs) and the SIBs.

[*System Basics and Services Command Reference*]

- **Command output changes for Type 5 FPC (T4000 routers)**—Starting with Release 12.1, Junos OS supports the Type 5 FPCs, thereby resulting in the following command output changes:
 - **show chassis environment fpc slot** command displays the temperatures and voltages on various sensors on the FPC. The actual information displayed might differ from the existing output format.

- **show chassis hardware** command additionally displays the Type 5 FPC output with the existing output fields.

[*T4000 Hardware Guide, System Basics and Services Command Reference*]

- **Support for 100-Gigabit Ethernet PIC on Type 5 FPC (T4000 Routers)**—Starting with Junos OS Release 12.1R1, the T4000 Core Router supports the 1-port 100-Gigabit Ethernet PIC on Type 5 FPC.

The 100-Gigabit Ethernet PIC is a 1-port 100-Gigabit Ethernet Type 5 PIC with 100-Gigabit C form-factor pluggable transceiver (CFP) (model number PF-1CGE-CFP).

The 100-Gigabit Ethernet PIC on Type 5 FPC supports the following major software features:

- Access to all 100-Gigabit Ethernet port counters through SNMP.
- Juniper Networks enterprise-specific Ethernet Media Access Control (MAC) MIB.

For detailed feature support and exceptions, see the *Ethernet Interfaces Configuration Guide*.



NOTE: Graceful Routing Engine switchover (GRES) and unified in-service software upgrade (unified ISSU) are not supported on T4000 Core Routers in Junos OS Release 12.1.

[*Ethernet Interfaces*]

- **Support for 10-Gigabit Ethernet LAN/WAN PIC with SFP+ on Type 5 FPC (T4000 Router)**—Starting with Junos OS Release 12.1R1, the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-12XGE-SFP) is supported on T4000 routers.

The following major software features are supported on the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- Access to all 10-Gigabit Ethernet port counters through SNMP
- LAN PHY mode

For detailed feature support and exceptions, see the *Ethernet Interfaces Configuration Guide*.

[*Ethernet Interfaces*]

- **VJX1000 Virtual Router**—Introducing the VJX Series family of Junos OS-based virtual routers that run within the Junosphere environment. Junosphere is a cloud-based, on-demand networking environment that enables network design, testing, and training using routers running Junos OS and security systems. The VJX Series delivers the software functionality of Juniper Networks routers including command-line interfaces (CLIs), control plane behavior, protocol operation, and forwarding functions.

For more information about Junosphere and VJX1000, see

http://www.juniper.net/techpubs/en_US/release-independent/junosphere/information-products/pathway-pages/junosphere/product/index.html.

- **Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles**—Starting with Junos OS Release 12.1, the **show interfaces lsq-fpc/pic/port** command has an **l2-statistics** option that displays Layer 2 queue statistics for the link services intelligent queuing (LSQ) and redundant LSQ interfaces. The queue statistics are displayed for Multilink Point-to-Point Protocol (MLPPP), FRF.15, and FRF.16 bundles on Multiservices PICs.

[*Interfaces Command Reference*]

- **Aggregated Ethernet interfaces support hierarchical queuing and shaping (MX Series routers with MPC/MIC interfaces)**—Extends support for aggregated Ethernet interfaces in non-link-protect mode through Junos OS Release 10.2 on MX Series routers with MPC/MIC interfaces. The scheduler functions supported are per-unit scheduler, hierarchical scheduler, and shaping at the physical and logical interface (aggregated interface) level.

[*Network Interfaces*]

- **Support for managing HTTP subscriber sessions based on request URI or domain name of a site**—You can configure a service set to set up a URL rule or a collection of URL rules that enable clients logging in to the router using HTTP sessions to be allowed or denied access. This functionality uses the capabilities of the High-Availability Chassis Manager (HCM) component. The URL rule or rule set contains a sequence of parameters that enable or discard access to HTTP clients for the website or server to which they want to access. You can configure the hostname or the domain name of websites for which you want to monitor and manage access by HTTP clients. When an HTTP client sends a GET request to the router, if the host portion or the uniform resource identifier (URI) portion of the HTTP request header matches the configured values in the defined URL rule in the service set, an action is taken as specified in the URL rule. You can specify an action to perform one of the following tasks when a match is found for the incoming HTTP request:

- **Accept**—Causes the HTTP requests to be processed and enables access to the requested site.
- **Accept and Log Requests**—Causes the HTTP requests to be processed and stores a system logging entry for each client session that was established.
- **Accept and Count Requests**—Enables access for the client that sends the HTTP GET request and saves each request from the client in a counter. This counter displays the cumulative value of all service sets that contain URL rules with the matching domain name or the IP address of the client's HTTP request that enabled access.
- **Discard**—Causes the HTTP requests to be dropped and disables access to the requested site.
- **Discard and Log Requests**—Causes the HTTP requests to be processed and stores a system logging entry for each client session that was terminated.

To configure the URL rule for processing of HTTP requests from clients, include the **url-rule** statement at the [**edit services hcm**] hierarchy level. To group a set of URL rules in a set, include the **url-rule-set** statement at the [**edit services hcm**] hierarchy level.

Each Multiservices DPC network processing unit (NPU) can service up to 50,000 requests or transactions per second from HTTP clients if the only operation that is running on the service plane is the URL monitoring process.

If more than one request URL or hostname is configured in a URL rule of a service set, the client that sends the HTTP request is enabled access when the first match is found for the domain name or URI portion of the HTTP request. The remaining hostnames or URIs are disregarded in the URL rule or rule set. If a hostname is not specified, "any" hostname is assumed, which is specified by an asterisk (*) in the URL rule. Similarly, if a request URI is not specified, "any" URI is assumed.

The HTTP URL management and monitoring feature for client requests works only for Services SDK applications.

[*Services Interfaces*]

- **Support for L-PDF and AACL on aggregated Multiservices interfaces**—Aggregated Multiservices PICs (ams interfaces) enable multiple Multiservices interfaces grouped together in a single bundle and cause the traffic destined for this ams group to be distributed over the member service PICs of the group. This capability enables load-balancing of traffic across various service PICs in an ams group.

You can configure the application identification (APPID) service and the intrusion detection and prevention (IDP) functionality on M120 or M320 routers equipped with Aggregated Multiservices PICs. Ams interfaces enable an N:1 redundancy mechanism to cluster together N number of ms- interfaces in an ams group that supports load sharing. Flows to be handled by APPID and IDP are distributed dynamically to all Multiservices PICs in an ams group using the Packet Forwarding Engine. This method of dynamic dispersion of packet flows avoids the limitations of throughput and scaling that might occur with a single Multiservices PIC.

[*Services Interfaces*]

- **Support for L-PDF and AACL on aggregated Multiservices interfaces**—Aggregated Multiservices PICs (ams interfaces) enable multiple Multiservices interfaces grouped together in a single bundle and cause the traffic destined for this ams group to be distributed over the member services PICs of the group. Junos Trio chipsets enable the calculation of a symmetric hash for the forward and reverse flows, and support a microcode map in the forwarding plane. This capability enables load-balancing of traffic across various services PICs in an ams group.

You can configure the application-aware access list (AACL) service and the local policy decision functionality (L-PDF) on M120 or M320 routers equipped with Aggregated Multiservices PICs. Ams interfaces enable an N:1 redundancy mechanism to cluster together N number of Multiservices interfaces in an ams group that supports load sharing.

Traffic policers are instantiated on a per-service PIC basis. As a result, if the traffic for one L-PDF subscriber is distributed over multiple Multiservices interfaces, the traffic policer functionality does not operate consistently.

The N:1 load sharing on ams- interfaces is stateless. After a failover from one Multiservices interface that encounters a fault to another Multiservices interface in the ams group, the state is rebuilt on the Multiservices interface that assumes the role of

a primary PIC for the flows that are routed through it. For ms- and rms- interfaces, the collection of statistics entries in the bulk statistics file is performed using the statistics reports received from the Multiservices PICs. For the ams- interfaces, this method of retrieval and storage of statistics is not possible because of multiple PICs containing statistics for the same subscriber. For interfaces in an ams group, statistics from the different Multiservices PICs in the ams group are collected and aggregated on the Routing Engine. On the Routing Engine, a timer control is activated and statistics are saved in the bulkstats file based on this timer. This method of collection causes the statistics records in the bulkstats file to be displayed with a small delay period.



NOTE: L-PDF uses the Berkeley database for management of configured settings. Because Junos OS Release 12.1 uses a database format that is different from the Berkeley database, when you upgrade from a Junos OS release that uses the Berkeley database to a Junos OS release in which L-PDF is supported on aggregated Multiservices interfaces, the previously stored settings in the database are deleted along with the statistical details.

[Services Interfaces]

- **NAT with deterministic port block allocation**—You can configure NAT algorithm-based allocation of blocks of destination ports. By specifying this method, you ensure that an IP address and port are always mapped to the same IP address in a pool and the same block of ports, thus eliminating the need for the logging of address translations. Other benefits include:
 - Configuration of source IP address prefix matching in the **from** clause of a NAT rule provides a high degree of scalability.
 - All Junos OS-supported ALGs are supported.

To configure deterministic port block allocation, include the **deterministic-port-block-allocation block-size *block-size*** statement at the [edit services nat pool *pool-name* port] hierarchy level and include the **translation-type deterministic-napt44** statement at the [edit services nat rule *rule-name* term *term-name* then translated] hierarchy level.

You can use the following commands to display information:

- **show services nat deterministic-nat nat-port-block internal-host *ip-address***
- **show services nat deterministic-nat internal-host nat-address *ip-address* nat-port *port-number***

[Next-Generation Network Addressing Solutions]

- **Distributed PPM support for LACP (T Series and M320 routers)**—Enables you to switch between distributed and centralized periodic packet management (PPM). By default distributed PPM is active. To enable centralized PPM, include the **ppm centralized** statement at the [edit interfaces *interface-name* fastether-options 802.3ad lacp] hierarchy level or the [edit interfaces *interface-name* gigether-options 802.3ad lacp] hierarchy level. To reenabling distributed PPM, include the **ppm distributed** statement at the [edit interfaces *interface-name* fastether-options 802.3ad lacp] hierarchy level

or the `[edit interfaces interface-name gigether-options 802.3ad lacp]` hierarchy level. You can use the `show lacp interfaces` command to display LACP Statistics output.

[Ethernet Interfaces, Interfaces Command Reference]

- **Support for reducing APS switchover time in Layer 2 Circuits (M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP)**—Starting in Junos OS Release 12.1, you can configure the `fast-aps-switch` statement to reduce the Automatic Protection Switching (APS) switchover time in Layer 2 circuits. You can configure the `fast-aps-switch` statement at the `[edit interfaces interface-name sonet-options aps]` hierarchy level. The `fast-aps-switch` statement can be configured on M320 routers with Channelized OC3/STM1 Circuit Emulation PIC with SFP only. Additionally, to achieve reduction in the APS switchover time, the Layer 2 circuit encapsulation type for the interface receiving traffic from a Layer 2 circuit neighbor must be Structure-Agnostic TDM over Packet (SAToP).



NOTE: The `fast-aps-switch` statement must be configured on both working and protect circuits.

The output of the `show l2circuit connections` operational command includes the **APS-active** and **APS-inactive** flags. These flags indicate the APS state of the interface. The **APS-active** flag indicates that the interface belongs to the working path. Similarly, the **APS-inactive** flag indicates that the interface belongs to the protective path.

[Interfaces Command Reference, VPNs, SONET/SDH Interfaces, Layer 2 Circuits]

- **Single-core Routing Engine for M7i and M10i routers**—Starting with Junos OS Release 12.1R2, a single-core Routing Engine is added to the M7i and M10i routers. This Routing Engine is based on the single-core Intel Xeon CPU, operating at 1.73 GHz with 2 MB cache. It has two DDR3 DIMM slots operating at 800 MHz that support 4 GB memory with error checking and correction (ECC). The new Routing Engine also supports:
 - 82574 Gigabit Ethernet Controller
 - 4 GB CompactFlash card
 - USB 2.0
 - Front accessible 64 GB solid-state drive (SSD)

All CLI commands supported on the older Routing Engine are supported on the new Routing Engine.

- **Support for WAN PHY mode on 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 routers)**—Starting with Junos OS Release 12.1R2, WAN PHY mode is supported on the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP), which is plugged into the Type 5 FPC of T4000 routers.

The following WAN PHY features are supported on the 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- WAN PHY mode on a per-port basis.
- Insertion and detection of path trace messages.

- Ethernet WAN Interface Sublayer (WIS) object.

To configure WAN PHY mode on a per-port basis, set the **wan-phy** option for the **framing** statement at the **[edit interface *interface-name*]** hierarchy level.



NOTE: When PHY mode changes, interface traffic is disrupted because of port reinitialization.

When WAN PHY mode is configured on an interface, the following SONET options are supported:

- Loopback (local and remote)
- Path trace
- Trigger options

[Ethernet Interfaces]

- **Support for interoperability between the 100-Gigabit Ethernet PIC on Type 4 FPC (T1600 routers) and the 100-Gigabit Ethernet PIC on Type 5 FPC (T4000 routers)**—Enables the interoperability between the 100-Gigabit Ethernet PIC on Type 4 FPC (on T1600 routers) and the 100-Gigabit Ethernet PIC on Type 5 FPC (on T4000 routers) by enabling a source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC on Type 5 FPC. The SA multicast bit steering mode uses the multicast bit in the source MAC address for packet steering.

By default, the SA multicast bit steering mode is not enabled on the 100-Gigabit Ethernet PIC on Type 5 FPC. To enable the SA multicast bit steering mode on the 100-Gigabit Ethernet PIC on Type 5 FPC, include the **forwarding-mode sa-multicast** statement at the **[edit chassis fpc *fpc-slot-number* pic *pic-slot-number*]** hierarchy level.



NOTE: The configuration of the **forwarding-mode sa-multicast** statement results in a PIC bounce—that is, the 100-Gigabit Ethernet PIC on Type 5 FPC goes offline and comes back online.

[Ethernet Interfaces, System Basics]

- **Support for flow monitoring services using Multiservices 500 PIC on T4000 routers**—Starting with Junos OS Release 12.1R2, the support for the following flow monitoring services is extended to T4000 routers using the Multiservices 500 PIC on Enhanced Scaling FPC3:
 - Active flow monitoring
 - Flow aggregation
 - Traffic sampling

[Services Interfaces, T4000 Router PIC Guide]

- **Improvements to Interface Transmit Statistics Reporting (MX Series devices)**—On MX Series devices, the logical interface-level statistics show only the offered load,

which is often different from the actual transmitted load. To address this limitation, Junos OS introduces a new configuration option in Releases 11.4R3, 12.1R4 and later. The new configuration option, **interface-transmit-statistics** at the **[edit interface *interface-name*]** hierarchy level, enables you to configure Junos OS to accurately capture and report the transmitted load on interfaces.

When the **interface-transmit-statistics** statement is included at the **[edit interface *interface-name*]** hierarchy level, the following operational mode commands report the actual transmitted load:

- **show interface *interface-name* <detail | extensive>**
- **monitor interface *interface-name***
- **show snmp mib get *objectID.ifIndex***



NOTE: This configuration is not supported on Enhanced IQ (IQE) and Enhanced IQ2 (IQ2E) PICs.

The **show interface *interface-name*** command also shows whether the interface-transmit-statistics configuration is enabled or disabled on the interface.

- **Support for disabling an FPC with degraded fabric bandwidth**—An FPC working with degraded fabric bandwidth can affect the re-routing process and can cause partial traffic black holes. On an MX960, MX480, or MX240 router, you can configure the option to bring down an FPC whose fabric bandwidth has degraded because of link errors or bad fabric planes. This configuration is particularly useful in partial black hole scenarios where bringing the FPC offline results in faster re-routing.

To configure this option on an FPC, use the **offline-on-fabric-bandwidth-reduction** statement at the **[edit chassis fpc *slot-number*]** hierarchy level.

Configuring this feature does not affect the system. You can configure this feature without restarting the FPC or restarting the system.

[*System Basics*]

- **Limiting traffic black-hole time by detecting Packet Forwarding Engine destinations that are unreachable over the fabric (MX240, MX480, and MX960 routers)**—Enables the MX240, MX480, and MX960 routers to limit traffic black-hole time by detecting unreachable destination Packet Forwarding Engines. The router signals neighboring routers when it cannot carry traffic because of the inability of some or all source Packet Forwarding Engines to forward traffic to some or all destination Packet Forwarding Engines on any fabric plane, after interfaces have been created. This inability to forward traffic results in a traffic black hole.

Packet Forwarding Engine destinations can become unreachable for the following reasons:

- The control boards go offline as a result of a CLI command or a pressed physical button.
- The fabric control boards are turned offline because of high temperature.

- Voltage or polled I/O errors in the SIBs detected by the SPMB.
- All Packet Forwarding Engines receive destination errors on all planes from remote Packet Forwarding Engines, even when the SIBs are online.
- Complete fabric loss caused by destination timeouts, even when the SIBs are online.

When the system detects any unreachable Packet Forwarding Engine destinations, healing from a traffic black hole is attempted. If the healing fails, the system turns off the interfaces, thereby stopping the traffic black hole.

The recovery process consists of the following phases:

1. Fabric plane restart phase: Healing is attempted by restarting the fabric planes one by one. This phase does not start if the fabric plane is functioning properly and a single Flexible PIC Concentrator (FPC) is bad. An error message is generated to specify that a black hole is the reason for the fabric plane being turned offline. This phase is performed for fabric plane errors only.
2. Fabric plane and FPC restart phase: The system waits for the first phase to be completed before examining the system state again. If the black hole condition still persists after the first phase is performed or if the problem occurs again within a duration of 10 minutes, healing is attempted by restarting both the fabric planes and the FPCs. If you configured the **action-fpc-restart-disable** statement at the **[edit chassis fabric degraded]** hierarchy level to disable restart of the FPCs when a recovery is attempted, an alarm is triggered to indicate that a traffic black hole has occurred. In this second phase, three steps are taken:
 1. All the FPCs that have destination errors on a PFE are turned offline.
 2. The fabric planes are turned offline and brought back online, one by one, starting with the spare plane.
 3. The FPCs that were turned offline are brought back online.
3. FPC offline phase: The system waits for the second phase to be completed before examining the system state again. Traffic black hole is limited by turning the FPCs offline and by turning off interfaces because previous attempts at recovery have failed. If the problem is not resolved by restarting the FPCs or if the problem recurs within 10 minutes after restarting the FPCs, this phase is performed.

By default, the system limits black-hole time by detecting severely degraded fabric. You do not need to configure anything to enable this feature. However, you can limit recovery actions to fabric plane restart only. You need to fix the traffic black hole by performing steps 2 and 3 manually.

In Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later, new alarms are added to indicate which FPCs are creating a traffic black hole in the system and to provide information about FPCs that are turned offline to stop the black hole in the recovery process.

In Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later, new error messages are added to indicate whether the traffic black hole is detected by unreachable FPCs in the system, or it is due to all planes being offline. These messages

also indicate the actions taken on FPCs and planes to stop the black hole—for example, FPC online, FPC offline, FPC restart, FPC power off, plane online, and plane offline.

Two new CLI commands are introduced for this feature:

- The **show chassis fabric unreachable-destinations** command shows the list of destinations that have changed from reachable to unreachable.
- The **show chassis fabric reachability** command shows the current state of fabric destination reachability, based on periodic reachability checks.

[*System Basics*]

Junos OS XML API and Scripting

- **Event policy support for configuration changes using Junos OS configuration mode commands**—Starting with Junos OS Release 12.1, you can configure an event policy to modify the configuration using Junos OS configuration mode commands and then commit the updated configuration. To configure an event policy to modify the configuration, include the **change-configuration** statement at the [**edit event-options policy policy-name then**] hierarchy level, and specify the configuration mode commands that are executed upon receipt of one or more configured events. The commands are executed in the order in which they appear in the event policy configuration. The commands update the candidate configuration, which is then committed, provided that no commit errors occur.

Configure the **commit-options** child statement to customize the event policy commit operation. Configure the **retry** statement to have the system attempt the change configuration event policy action a specified number of times if the first attempt fails. The **user-name** statement specifies the user under whose privileges the configuration changes and commit are made.

[*Junos OS Configuration and Operations Automation Guide*]

- **Event policy support to override the system log priority of the triggering event**—Starting with Junos OS Release 12.1, you can configure an event policy to override the default system log priority of a triggering event so that the system logs the event with a different facility type, severity level, or both. To override the priority of the triggering event, configure the **priority-override** statement at the [**edit event-options policy policy-name then**] hierarchy level. To override the facility type with which the triggering event is logged, include the **facility** statement and the new facility type. To override the severity level with which the triggering event is logged, include the **severity** statement and the new severity level.

[*Junos OS Configuration and Operations Automation Guide*]

- **Junos XML protocol support for <get-configuration> attributes commit-scripts="apply" and commit-scripts="apply-no-transients"**—Starting with Junos OS Release 12.1, the <get-configuration> operation supports two new values for the **commit-scripts** attribute. The **commit-scripts="apply"** attribute value displays the configuration with commit script changes applied, including both transient and non-transient changes. The **commit-scripts="apply-no-transients"** attribute value

displays the configuration with commit script changes applied, but excludes transient changes.

[[Junos XML Management Protocol Guide](#)]

- **jcs:open() extension function support for routing-instance**—The `jcs:open()` extension function returns a connection handle that is used to execute RPCs on a local or remote device. To redirect the SSH connection to originate from within a specific routing instance, include the name of the routing instance in the connection parameters. The routing instance must be configured at the **[edit routing-instances]** hierarchy level, and the remote device must be reachable either using the routing table for that routing instance or from one of the interfaces configured under that routing instance.

[[Junos OS Configuration and Operations Automation Guide](#)]

- **Support for commit script access to the pre-inheritance candidate configuration in configure private sessions**—Commit scripts can invoke the `<get-configuration>` RPC in a private configuration session to retrieve the private, pre-inheritance candidate configuration for that session. The `<get-configuration>` RPC includes the **database-path** attribute, which is used to specify the location of the pre-inheritance configuration database. In addition, the global variable `$junos-context` contains a new **commit-context/database-path** element, which stores the location of the session's pre-inheritance candidate configuration.

To construct a commit script that retrieves the pre-inheritance candidate configuration specific to that session, include the `<get-configuration>` RPC in the commit script, and set the `<database-path>` attribute to `$junos-context/commit-context/database-path`. This feature is available in Junos OS Release 12.1R3 and subsequent 12.1 releases.

Layer 2 Ethernet Services

- **Support for Layer 2 Ethernet interface service features on T4000-FPC5-3D**—Starting with Junos OS Release 12.1, all Layer 2 features are supported on T4000-FPC5-3D, with the following exceptions:

- Media access control (MAC) filtering



NOTE: Because destination MAC filtering is not supported, the hardware is configured to accept all the multicast packets. This enables the OSPF protocol to work.

- MAC learning
- MAC policing
- MAC accounting

[*Network Interfaces*]

MPLS Applications

- **Nonstop active routing support for RSVP point-to-multipoint ingress LSPs**—Starting with Junos OS Release 12.1, Junos OS extends nonstop active routing support to RSVP point-to-multipoint ingress LSPs. Nonstop routing support for RSVP point-to-multipoint egress and transit routers was added in Junos OS Release 11.4.

During the switchover, the LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover.

However, Junos OS nonstop active routing support for RSVP point-to-multipoint LSPs does not include support for dynamically created point-to-multipoint LSPs, such as VPLS and next-generation MVPNs.

[*MPLS*]

- **MPLS label removal on T Series routers**—T Series routers in passive-monitor-mode support removing up to five MPLS labels.

[*MPLS*]

- **MPLS Transport Profile (MPLS-TP)**—The MPLS Transport Profile (MPLS-TP) introduces new capabilities for Operations, Administration, and Management (OAM) when MPLS is used for transport services and transport network operations. This includes a generic mechanism to send OAM messages. This mechanism contains two main components:
 - Generic Alert Label (GAL)—A special label that enables an exception mechanism that informs the egress label-switching router (LSR) that a packet it receives on an LSP belongs to an associated control channel or the control plane.
 - Generic Associated Control Channel Header (G-Ach)—A special header field that identifies the type of payload contained in the MPLS label-switched paths (LSPs). G-Ach has the same format as a pseudowire associated control channel header.

For more information about MPLS-TP, see RFC 5654, *Requirements of an MPLS Transport Profile*. For more information about GAL and G-Ach, see RFC 5586, *MPLS Generic Associated Channel*.

The following capabilities are supported in the Junos OS implementation of MPLS-TP:

- MPLS-TP OAM can send and receive packets with GAL and G-Ach, without IP encapsulation.
- Two unidirectional RSVP LSPs between a pair of routers can be associated with each other to create an associated bidirectional LSP for binding a path for the GAL and G-Ach OAM messages. The associated bidirectional LSP model is only supported for associating the primary paths. A single Bidirectional Forwarding Detection (BFD) session is established for the associated bidirectional LSP.

The current Junos OS implementation of MPLS-TP does not support:

- P2MP RSVP LSPs and BGP LSPs
- Loss Measurement (LM) and Delay Measurement (DM)

To enable GAL/G-Ach OAM operation without IP encapsulation on all LSPs, include the **mpls-tp-mode** configuration statement at the **[edit protocols mpls oam]** hierarchy level.

```
[edit protocols mpls oam]
mpls-tp-mode;
```

To enable GAL/G-Ach OAM operation without IP encapsulation on a specific LSP, include the **mpls-tp-mode** statement at the **[edit protocols mpls label-switched-path *lsp-name* oam]** hierarchy level.

```
[edit protocols mpls label-switched-path lsp-name oam]
mpls-tp-mode;
```



NOTE: Include this statement at the **[edit protocols mpls oam]** hierarchy level only if all the LSPs are point-to-point LSPs.

To configure associated LSPs on the two ends of the LSP, include the **associate-lsp *lsp-name* oam from *from-ip-address*** statement at the **[edit protocols mpls lsp *lsp-name* oam]** hierarchy level.

```
[edit protocols mpls lsp lsp-name oam]
associate-lsp lsp-name {
  from from-ip-address;
}
```

The **from *from-ip-address*** configuration for the LSP is optional. If omitted, it is derived from the **to** address of the LSP configuration.

To associate two LSPs at a transit router, include the **transit-lsp-association** statement at the **[edit protocols mpls]** hierarchy level.

```
[edit protocols mpls]
transit-lsp-association transit-association-lsp-group-name {
  lsp-name-1 name-of-associated-lsp-1;
  from-1 address-of-associated-lsp-1;
  lsp-name-2 name-of-associated-lsp-2;
  from-2 address-of-associated-lsp-2;
}
```

The association in the transit nodes is useful for the return LSP path for TTL-expired LSP ping packets or traceroute.

To view details of associated bidirectional LSPs, issue the **show mpls lsp** command. To view detailed information, issue the command with the **detail** or **extensive** option. In addition, you can also use the **show mpls lsp bidirectional** command to view associated bidirectional LSP information.

[MPLS]

- **MPLS support on Type 5 FPC (T4000 routers)**—Starting with Junos OS Release 12.1, MPLS support is extended to the Type 5 FPC on T4000 Core Routers.

The existing MPLS labels **push**, **pop**, **swap**, **multiple push**, and **swap and push** are supported on T4000 routers.

The Type 5 FPC on T4000 routers supports the following features, which are also supported on aggregated Ethernet interfaces:

- Layer 2 VPNs, Layer 2 circuits, and Layer 2 switching cross-connects (with circuit cross-connect (CCC) and VLAN CCC encapsulation)
- Layer 3 VPNs applicable on IPv4 and IPv6 routes:
 - With tunnel services
 - With the **vrf-table-label** statement with LSI and no tunnel services
 - With per-prefix load balancing with no tunnel services
- Interprovider and carrier-of-carriers VPNs and BGP MPLS multicast VPNs
- MPLS LSP tunnel cross-connects and LSP stitching cross-connects
- Ethernet translational cross-connect (TCC) and VLAN TCC encapsulation
- IPv4, MPLS, and ISO packets for TCC
- Point-to-multipoint CCC support (ingress and egress) using RSVP point-to-multipoint LSPs
- Class-of-service (CoS)-based features such as:
 - MPLS EXP classification and rewrites
 - Fixed-CoS value for a label-switched path (LSP) and RSVP bypass LSPs
 - CoS-based forwarding support for MPLS
- LDP-signaled LSPs, LSP accounting, LSP policers
- LSP ping and traceroute, which includes LSP traceroute for LDP LSPs with equal-cost multipath (ECMP) support
- RSVP-signaled LSPs
- RSVP-signaled point-to-multipoint LSPs, which includes link protection for the LSPs, and maximum transmission unit (MTU) signaling in RSVP
- MPLS fast reroute, node protection, and link protection
- Time to live (TTL) propagation and explicit NULL label support (ultimate-hop popping) for IPv4 and IPv6
- MPLS load balancing based on IP header and MPLS labels
- Static and explicit-path LSP, including support for **push** label and **swap-push** label operations
- MPLS over GRE tunnels and IPv6 tunnels over MPLS
- MPLS firewall filters
- Generalized MPLS (GMPLS)

- Source class usage (SCU) on label-switched interfaces (LSIs)
- Diffserv-aware traffic engineering

Note that the point-to-multipoint LSP traceroute and bidirectional Protocol Independent Multicast (PIM) features are currently not supported on T4000 routers.

All the MPLS features that are currently supported on existing T Series Core Routers are also supported on the T4000 routers.

[MPLS]

- **New features supported on T4000 Core Routers**—Starting with Junos OS Release 12.1R2, the following features are supported on T4000 routers:
 - **BFD for IPv6 BGP sessions**—You can configure Bidirectional Forwarding Detection (BFD) liveness detection on IPv6 BGP connections. For more information, see [Configuring Bidirectional Forwarding Detection for BGP](#).
 - **Advertise multiple paths to a destination for IBGP**—For IPv4 unicast (family inet unicast) routes, you can enable an internal BGP (IBGP) peer to advertise multiple exit points to reach a destination. For more information, see [Advertising Multiple Paths in BGP](#).
 - **Nonstop active routing support to Layer 2 circuits on RSVP-TE based LSPs**—Junos OS extends nonstop active routing support to label-switching routers (LSRs) and Layer 2 circuits that are part of an RSVP-TE based label-switched path (LSP). For more information, see [Nonstop Active Routing Support for RSVP-TE LSPs](#).

Multicast

- **Bidirectional PIM (RFC 5015) (M120, M320, MX Series, and T Series routers)**—Provides an alternative to other PIM modes, such as PIM sparse mode (PIM-SM), PIM dense mode (PIM-DM), and PIM source-specific multicast (SSM). In bidirectional PIM, multicast groups are carried across the network over bidirectional shared trees. This type of tree minimizes the amount of PIM routing state information that must be maintained, which is especially important in networks with numerous and dispersed senders and receivers.

To configure designated forwarder election parameters, include the **bidirectional** statement and child statements at the **[edit protocols pim interface *interface-name* | all]** hierarchy level. To configure the PIM mode, include the **bidirectional-sparse** or **bidirectional-sparse-dense** statement at the **[edit protocols pim interface (*interface-name*) | all] mode]** hierarchy level. To configure tracing operations, include the **bidirectional-df-election** statement at the **[edit protocols pim traceoptions flag]** hierarchy level. To configure an RP address, include an IP address at the **[edit protocols pim rp bidirectional address]** hierarchy level.

Useful monitoring and troubleshooting commands include all of the existing **show pim...**, **show route...**, and **show multicast...** commands, plus the **show pim bidirectional df-election** command.

[Multicast Protocols]

- **Load-balancing PIM join messages on multicast VPNs**—Starting with Release 12.1, Junos OS supports customer PIM (C-PIM) join messages to be load-balanced across unequal EBGP and IBGP paths in a Draft Rosen MVPN and a next-generation MVPN in the following ways:

- In the case of a Draft-Rosen MVPN, unequal EBGP and IBGP paths are utilized.
- In the case of next-generation MVPN:
 - Available IBGP paths are utilized when no EBGP path is present.
 - Available EBGP paths are utilized when both EBGP and IBGP paths are present.

This feature is applicable to IPv4 C-PIM join messages over the Layer 3 MVPN service.

[*Multicast, VPNs*]

Network Management

- **Updated enterprise-specific MIB and support for existing system log messages and operational commands for T4000 routers**—Starting with Junos OS Release 12.1, the following features are supported on T4000 Core Routers:

- **Updated MIB**—The Juniper Networks enterprise-specific Chassis MIB provides information about the router and its components. The enterprise-specific Chassis Definitions for Router Model MIB—**jnx-chas-defines.mib**—is updated for T4000 routers with object identifiers (OIDs) that are used by the Chassis MIB to identify platform and chassis components.
- **Support for existing system log messages**—The following system log messages apply to T4000 routers:
 - CHASSISD_PEM_IMPROPER
 - CHASSISD_UNSUPPORTED_FPC
 - CHASSISD_UNSUPPORTED_SIB

The other alarms related to T4000 routers can be viewed by executing the **show chassis alarms** operational command.

On T4000 Type 5 FPCs, there are no **top temperature sensor** or **bottom temperature sensor** parameters. Instead, **fan intake temperature sensor** and **fan exhaust temperature sensors** parameters are displayed.

Starting with Junos OS Release 12.1, you can configure five input feeds of the six-input DC power supply for the T4000 router. By default, the power supply is configured to have all the six input feeds connected. Note that all the power supplies in the router must use the same number of input feeds. Before configuring the input feeds, see the *T4000 Core Router Hardware Guide* for special considerations and for the number of input feeds supported by the router.

[*System Log Messages Reference, SNMP MIBs and Traps Reference, Interface Command References*]

- **SNMP MIB support for OSPFv3**—Starting with Release 12.1, Junos OS supports RFC 5643, Management Information Base for OSPFv3, and thus extends the SNMP support to OSPFv3. Junos OS support for RFC 5643 is read-only, and does not include `ospfv3HostTable` and `ospfv3CfgrTable`.

[*SNMP MIBs and Traps Reference*]

- **Junos OS support for proxy SNMP agent**—Junos OS enables you to assign one of the devices in the network as a proxy SNMP agent through which the network management system (NMS) can query other devices in the network. When you configure a proxy, you can specify the names of devices to be managed through the proxy SNMP agent.

When the NMS queries the proxy SNMP agent, the NMS specifies the community name (for SNMPv1 and SNMPv2) or the context and security name (for SNMPv3) associated with the device from which it requires the information.



NOTE: If you have configured authentication and privacy methods and passwords for SNMPv3, those parameters are also specified in the query for SNMPv3 information.

To configure a proxy SNMP agent and specify devices to be managed by the proxy SNMP agent, you can include the following configuration statements at the [**edit snmp**] hierarchy level:

```
proxy proxy-name {
  device-name device-name;
  <version-v1 | version-v2c> {
    snmp-community community-name;
    no-default-comm-to-v3-config;
  }
  version-v3 {
    security-name security-name;
    context context-name;
  }
  logical-system logical-system {
    routing-instance routing-instance;
  }
  routing-instance routing-instance;
}
```

- The **proxy** statement enables you to specify a unique name for the proxy configuration.
- The **version-v1**, **version-v2**, and **version-v3** statements enable you to specify the SNMP version.
- The **no-default-comm-to-v3-config** statement is an optional statement at the [**edit snmp proxy proxy-name <version-v1 | version-v2>**] hierarchy level that when included in the configuration requires you to manually configure the statements at the [**edit snmp v3 snmp-community community-name**] and [**edit snmp v3 vacm**] hierarchy levels.

If the **no-default-comm-to-v3-config** statement is not included at the [**edit snmp proxy proxy-name <version-v1 | version-v2>**] hierarchy level, the [**edit snmp v3**

`snmp-community community-name`] and `[edit snmp v3 vacm]` hierarchy level configurations are automatically initialized.

- The **logical-system** and **routing-instance** statements are optional statements that enable you to specify logical system and routing instance names if you want to create proxies for logical systems or routing instances on the device.



NOTE: The community and security configuration for the proxy should match the corresponding configuration on the device that is to be managed.



NOTE: Because the proxy SNMP agent does not have trap forwarding capabilities, the devices that are managed by the proxy SNMP agent send the traps directly to the network management system.

You can use the **show snmp proxy** operational mode command to view proxy details on a device. The **show snmp proxy** command returns the proxy names, device names, SNMP version, community/security, and context information.

Routing Protocols

- **New option propagate-ttl for IP traceroute**—The **traceroute** command has a new option **propagate-ttl** that can be used on a PE router to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

[*Routing Protocols Command Reference*]

- **Support for RFC 4861**—In Junos OS Release 12.1 and later, Junos OS supports Neighbor Discovery features as described in RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*, along with RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*, and RFC 4862, *IPv6 Stateless Address Auto Configuration for IPv6*. A new configuration statement, **do-not-disable-ip6-op**, introduced at the `[edit system]` hierarchy level, prevents IPv6 operation on an interface from being disabled when the duplicate address detection process fails on link-local addresses that are based on hardware addresses.

[*System Basics*]

- **Support for Layer 3 features on T4000 routers**—Support for Layer 3 protocols and Layer 3 forwarding is extended to T4000 routers. [Table 17 on page 256](#) lists the protocols, features, and services supported on T4000 routers.

Table 17: Protocols, Features, and Services Supported on T4000 Routers

Protocols	Features	Services
BGP	Equal-cost multipath (ECMP)	Layer 3 virtual private network (Layer 3 VPN)
OSPF	Loop-free Alternate	Layer 2 virtual private network (Layer 2 VPN)

Table 17: Protocols, Features, and Services Supported on T4000 Routers (*continued*)

Protocols	Features	Services
IS-IS	Unicast reverse path forwarding (unicast RPF)	Layer 2 circuit
RIP		
Bidirectional Forwarding Detection (BFD)		
SNMP		
Address Resolution Protocol (ARP)		
Neighbor Discovery Protocol (NDP)		
RSVP		

[Routing Protocols]

- **IEEE 802.3ah OAM functionality extended to 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T640, T1600, and TX Matrix routers with T640-FPC4-ES, T1600-FPC4-ES, and T640-FPC4-1P-ES)**—Enables you to perform Operation, Administration, and Maintenance (OAM)-related operations such as link fault management and link discovery.

Support for the following OAM operations on the T Series routers is extended to the 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- Link fault management
- Link discovery
- Graceful Routing Engine switchover (GRES)
- Layer 2 and Layer 3 control protocol packets (OSPF, OSPF3, VRRP, IGMP, RSVP, PIM, BGP, BFD, LDP, IS-IS, RIP, RIPV6, LACP, ARP, IPv6 NDP, CFM, and LFM) are mapped to the control queue. In the control queue, these packets are not dropped even if there is oversubscription or congestion on a port group.



NOTE: The following OAM features are not supported on the 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- Remote Loopback
- Unified in-service software upgrade (unified ISSU)

[Channelized Interfaces, Routing Protocols]

- **IEEE 802.3ag OAM functionality extended to 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T640, T1600, and TX Matrix routers with T640-FPC4-ES,**

T1600-FPC4-ES, and T640-FPC4-1P-ES)—Enables you to perform Operation, Administration, and Maintenance (OAM)-related operations. Support for the following OAM operations on the T Series routers is extended to the 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- Connectivity fault management (CFM)
- Linktrace
- Loopback
- Graceful Routing Engine switchover (GRES)
- Layer 2 and Layer 3 control protocol packets (OSPF, OSPF3, VRRP, IGMP, RSVP, PIM, BGP, BFD, LDP, IS-IS, RIP, RIPV6, LACP, ARP, IPv6 NDP, CFM, and LFM) are mapped to the control queue. In the control queue, these packets are not dropped even if there is oversubscription or congestion on a port group.



NOTE: OAM unified in-service software upgrade (unified ISSU) is not supported on the 10-Gigabit Ethernet LAN/WAN PIC with SFP+.

[Channelized Interfaces, Routing Protocols]

- **Accumulated IGP (AIGP) attribute for BGP**—Enables deployment in which a single administration can run several contiguous BGP autonomous systems. Such deployments allow BGP to make routing decisions based on the IGP metric. In such networks, it becomes possible for BGP to select paths based on metrics as is done by IGP. In this case, BGP chooses the shortest path between two nodes, even though the nodes might be in two different autonomous systems. To enable accumulated IGP processing, include the **aigp** statement in the BGP configuration on a protocol family basis. Junos OS supports accumulated IGP for **family inet labeled-unicast** and **family inet6 labeled-unicast**. The **aigp** statement can be configured for a given family at the global BGP, group, or neighbor level. By default, the value of the AIGP attribute for a local prefix is zero. An AIGP-enabled neighbor can originate an AIGP attribute for a given prefix by export policy, using the **aigp-originate** policy action. The value of the AIGP attribute reflects the IGP distance to the prefix. Alternatively, you can specify a value, by using the **aigp-originate distance distance** policy action.

[Routing Protocols]

- **Point-to-multipoint support for RIP**—The demand circuit (DC) feature implementation in RIP required the use of a single RIP peer. The point-to-multipoint feature enables a RIP device to have multiple peers on an interface irrespective of whether it uses demand circuits. To enable this feature, include:
 - The **interface-type p2mp** statement at the **[edit protocols rip group group-name neighbor interface-name]** hierarchy level to enable the neighbor to function as a point-to-multipoint endpoint.
 - The **dynamic-peers** statement at the **[edit protocols rip group group-name neighbor interface-name]** hierarchy level to enable or disable dynamic peer discovery at the neighbor level.

- The **peer address** statement at the **[edit protocols rip group group-name neighbor interface-name]** hierarchy level to manually configure peers.

As a result of this feature, the following **show** commands have been introduced to view the statistics of all peers or a given peer:

- **show rip statistics peer all**
- **show rip statistics peer address**

As a result of this feature, the following **clear** commands have been introduced to clear the statistics of all peers or a given peer:

- **clear rip statistics peer all**
- **clear rip statistics peer address**

[*Routing Protocols*]

- **PE routers cache data MDT advertisements of new data MDT group addresses**—In a draft-rosen Layer 3 multicast virtual private network (MVPN) configured with service provider tunnels, provider edge (PE) routers cache the information received in any “MDT join” type-length-value (TLV) packets sent over the default multicast distribution tree (MDT) for a multicast-enabled VPN routing and forwarding (VRF) instance. You can configure PE routers so that when the multicast source within a site exceeds a traffic rate threshold, the PE router to which the source site is attached creates a new data MDT. For more information, see [Example: Configuring Data MDTs and Provider Tunnels Operating in Source-Specific Multicast Mode](#).
- **Interdomain point-to-multipoint LSPs**—Interdomain point-to-multipoint LSPs can be used to transport traffic in the following applications in a multi-area or multi-AS network:
 - Layer 2 broadcast and multicast over MPLS
 - Layer 3 BGP/MPLS VPN
 - Virtual private LAN service (VPLS)

For more information, see [Configuring Inter-domain P2MP LSPs](#).

- **Support for the pimNeighborLoss trap**—The pimNeighborLoss trap is defined in RFC 2934, *Protocol Independent Multicast MIB for IPv4*. This trap is generated when the device loses the adjacency with its only neighbor that has an IP address lower than that of the interface to which the neighbor is connected. For more information, see [Standard SNMP MIBs Supported by Junos OS](#).
- **Internet multicast using ingress replication provider tunnels**—Enables a faster path for multicast traffic between sender and receiver routers in large-scale implementations. The **mpls-internet-multicast** routing instance type uses existing Junos OS technology and ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud. This configuration is available under PIM and multicast virtual private network (MVPN) infrastructure. For more information, see [Example: Configuring Ingress Replication for IP Multicast Using MBGP MVPNs](#).

- **LDP signaling for point-to-multipoint LSPs in next-generation MBGP multicast VPNs**—A point-to-multipoint label-switched path (LSP) is an LDP-signaled (or RSVP-signaled) LSP with a single source and multiple destinations. This feature enables point-to-multipoint LSPs for multicast Border Gateway Protocol (MBGP) VPNs in an intra-autonomous system (AS) environment (within an AS). For more information, see [Example: Configuring Point-to-Multipoint LDP LSPs as the Data Plane for Intra-AS MBGP MVPNs](#).
- **Nonstop active routing support for RSVP point-to-multipoint transit, ingress, and egress LSPs**—During a Routing Engine switchover, an LSP comes up on the backup Routing Engine that shares and synchronizes the state information with the master Routing Engine before and after the switchover. Nonstop active routing support for point-to-multipoint LSPs ensures that the switchover remains transparent to the network neighbors, and preserves the LSP information across the switchover. For more information, see [Nonstop Active Routing Support for RSVP-TE LSPs](#).
- **Enhancements to nonstop active routing Protocol Independent Multicast (PIM) support**—Junos OS Release 11.1 and 11.3 extended the nonstop active routing support for PIM to include the following features:

Support for the following features was extended in Junos OS Release 11.1:

- Local rendezvous point (RP)-set information synchronization is supported for local RP and bootstrap router (BSR) (on IPv4 and IPv6), auto-RP (on IPv4), and embedded RP (on IPv6).
- Anycast RP-set information synchronization and anycast RP register state synchronization on IPv4 and IPv6 configurations.
- Flow maps
- Unified ISSU
- Nonstop active routing for PIM configuration on devices that have both IPv4 and IPv6 configured on them.

Support for the following features was extended in Junos OS Release 11.3:

- Policy features such as neighbor policy, bootstrap router export and import policies, scope policy, flow maps, and reverse path forwarding (RPF) check policies.
- Upstream assert synchronization
- PIM join load balancing

These features are supported on T4000 routers as well. For more information about nonstop active routing PIM, see [Nonstop Active Routing PIM Support](#) and [Example: Configuring Nonstop Active Routing for PIM](#).

- **Support for RSVP-signaled point-to-multipoint LSPs extended to logical systems**—Junos OS Release 11.3 extended the support for RSVP-signaled point-to-multipoint LSPs to logical systems. The following topologies were supported:
 - A single logical system in a physical router. The logical system is one node in an RSVP-signaled point-to-multipoint label-switched path (LSP).

- Multiple logical systems in a physical router, with each logical system acting as a label-switching router (LSR). The multiple logical systems can be unconnected, connected to each other internally with logical tunnel (lt) interfaces, or connected to each other externally with back-to-back connections.
- One RSVP-signaled point-to-multipoint LSP, with some nodes being logical systems and other nodes being physical routers.

These topologies are supported on T4000 routers as well. For more information, see [Configuring RSVP-Signaled Point-to-Multipoint LSPs on Logical Systems](#).

- **Source-specific multicast (SSM) map definition for different groups to different sources**—You can configure multiple source-specific multicast (SSM) maps so that different groups map to different sources, which enables a single multicast group to map to different sources for different interfaces. For more information, see [Example: Configuring SSM Maps for Different Groups to Different Sources](#).
- **Support for testing the operability of BGP-signaled MPLS LSP connections**—You can use the `ping mpls bgp` command to test the operability of BGP-signaled MPLS LSP connections. Press Ctrl+c to interrupt a `ping mpls` command. If the LSP changes, the label and interface information displayed when you issued the ping command continues to be used.



NOTE: You must configure MPLS at the `[edit protocols mpls]` hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only BGP forwarding equivalence classes (FECs).

When you enter the `ping mpls` command, you are provided feedback on the status of your request. An exclamation point (!) indicates that an echo reply was received. A period (.) indicates that an echo reply was not received within the timeout period. An x indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately. If you issue the `ping mpls bgp` command with the `detail` option, the error codes are displayed in the command output.

- **MIB support for VRF route entries**—Junos OS extends the SNMP support to Layer 3 virtual private network (VPN) routing and forwarding table (VRF) entries as defined in RFC 4382, *MPLS/BGP Layer 3 Virtual Private Network (VPN) MIB*. For more information, see [Standard SNMP MIBs Supported by Junos OS](#).
- **NTP support for IPv4 VPN routing and forwarding (VRF) requests**—Enables a Network Time Protocol (NTP) server running on a provider edge (PE) router to respond to NTP requests from a customer edge (CE) router. As a result, a PE router can process any NTP request packet coming from different routing instances. For more information, see [NTP Overview](#).

Subscriber Access Management



NOTE: Although present in the code, the subscriber management features are not supported in Junos OS Release 12.1R7. Documentation for subscriber management features is included in the Junos OS Release 12.1 documentation set.

- **Junos OS subscriber management scaling values (M120, M320, and MX Series routers)**—A spreadsheet is available online that lists scaling values supported for Junos OS subscriber management beginning with Junos OS Release 10.1. Access the *Subscriber Management Scaling Values (XLS)* spreadsheet from the Downloads box at http://www.juniper.net/techpubs/en_US/junos12.1/information-products/pathway-pages/subscriber-access/index.html. You can also substitute the number of the latest Junos OS release for the 12.1 release number. For example, ...en_us/junos12.2/...

[*Subscriber Management Scaling*]

- **Delay in removing subscriber routes after graceful Routing Engine switchover (M120, M320, and MX Series routers)**—For a subscriber network in which either nonstop active routing (NSR) or graceful restart has been configured, you can configure the router to wait for 180 seconds (3 minutes) before removing access routes and access-internal routes after a graceful Routing Engine switchover (GRES) takes place.

This 3-minute delay provides sufficient time for the appropriate client process (jpppd or jdhcpd) or routing protocol process to reinstall the access routes and access-internal routes before the router removes the stale routes from the forwarding table. As a result, the risk of traffic loss is minimized because the router always has available subscriber routes.

To configure the router to wait for 180 seconds before removing (flushing) access routes and access-internal routes after a graceful Routing Engine switchover, include the **gres-route-flush-delay** statement at the [edit system services subscriber-management] hierarchy level.

Using the **gres-route-flush-delay** statement in your subscriber management configuration offers the following benefits:

- Provides sufficient time to reinstall subscriber routes from the previously active Routing Engine

In subscriber networks with graceful restart and routing protocols such as BGP and OSPF configured, the router purges any remaining stale routes as soon as the graceful restart operation completes, which can occur very soon after completion of the graceful Routing Engine switchover. Using the **gres-route-flush-delay** statement causes the router to retain the stale routes for a full 180 seconds, which provides sufficient time for the jpppd or jdhcpd client process to reinstall all of the subscriber routes.

- Prevents loss of subscriber traffic due to unavailable routes

In subscriber networks with nonstop active routing and routing protocols such as BGP and OSPF configured, the routing protocol process immediately purges the stale routes that correspond to subscriber routes. This removal results in a loss of subscriber traffic. Using the **gres-route-flush-delay** statement causes the router to retain the stale routes for a full 180 seconds, which prevents potential traffic loss due to unavailable routes.

[*Junos OS Subscriber Access Configuration Guide*]

- **Configuring retransmission of L2TP control messages (MX Series Routers)**—L2TP peers maintain a queue of control messages that must be sent to the peer device. After a message is sent, the local peer waits for a response from the remote peer. If a response is not received, the local peer retransmits the message. You can configure how many times an unacknowledged message is retransmitted by the LAC or the LNS. For tunnels that have been established, include the **retransmission-count-established** statement at the **[edit services l2tp tunnel]** hierarchy level. For tunnels that are not yet established, include the **retransmission-count-not-established** statement.

You can specify a maximum count in the range 0 through 30. The default count for established tunnels is 7; for not-established tunnels, the default is 5. The local peer waits 1 second for the first response to a control message. The retransmit timer then doubles the interval between each successive retransmission, up to a maximum interval of 16 seconds. This behavior allows the remote peer more time to respond. If the maximum retransmission count is reached and no response has been received, the tunnel and all its sessions are cleared.



BEST PRACTICE: Before you downgrade to a Junos OS release that does not support these statements, we recommend that you explicitly unconfigure the feature by including the **no retransmission-count-established** statement and the **no retransmission-count-non-established** statement at the **[edit services l2tp tunnel]** hierarchy level.



BEST PRACTICE: During a unified in-service software upgrade (ISSU) on an MX Series router configured as the LAC, the LAC might not respond to control messages from the LNS. This can result in dropping LAC L2TP sessions. You can avoid this situation by ensuring that the maximum retransmission count on the LNS is set to 16 or higher.

[*Subscriber Access*]

- **Configuring the idle timeout for L2TP tunnels without sessions (MX Series routers)**—You can configure the LAC or the LNS to specify how long a tunnel without any sessions remains active. The idle timer starts when the last session on the tunnel is terminated. When the timer expires, the tunnel is disconnected. This idle timeout frees up resources otherwise consumed by inactive tunnels.

To configure how long the tunnel remains active, include the **idle-timeout** statement at the **[edit services l2tp tunnel]** hierarchy level. You can set the timer in the range 0 through 86,400 seconds; the default value is 80 seconds. If you set the idle-timeout value to 0, the tunnel is forced to remain active indefinitely after the last session is terminated until one of the following occurs:

- You issue the **clear services l2tp tunnel** command.
- The remote peer disconnects the tunnel.



BEST PRACTICE: Before you downgrade to a Junos OS release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no idle-timeout** statement at the **[edit services l2tp tunnel]** hierarchy level.

[Subscriber Access]

- **Configuring how long L2TP maintains dynamic information (MX Series routers)**—You can configure the LAC or the LNS to specify how long the router attempts to maintain dynamic destinations, tunnels, and sessions after they have been torn down. This *destruct* timeout aids debugging and other analysis by saving underlying memory structures of those destinations, tunnels, or sessions. Any specific dynamic destination, tunnel, or session might not be maintained for this entire time period if the resources must be reclaimed early to allow new tunnels to be established. To set the destruct timeout, include the **destruct-timeout** statement at the **[edit services l2tp]** hierarchy level. You can set the timer in the range 10 through 3600 seconds; the default value is 300 seconds.



BEST PRACTICE: Before you downgrade to a Junos OS release that does not support this statement, we recommend that you explicitly unconfigure the feature by including the **no destruct-timeout** statement at the **[edit services l2tp]** hierarchy level.

[Subscriber Access]

- **Support for shaping rate and overhead accounting on dynamic subscriber interfaces based on access line parameters in PPPoE discovery packets (MX Series routers)**—Enables you to configure access line parameters in PPPoE discovery packets to set the shaping rate and overhead accounting attributes on dynamic subscriber interfaces in a broadband access network. This feature is supported on MPC/MIC interfaces on MX Series routers.

When you enable this feature, the values supplied by the PPPoE vendor-specific tags override the parameters that you have configured for **shaping-rate** and **overhead-accounting** statements at the **[edit dynamic-profiles profile-name class-of-service traffic-control-profile]** hierarchy level.

The shaping rate is based on the actual data rate downstream attribute, and is only overridden if the vendor specific-tag value is less than the configured value. The

overhead accounting value is based on the access loop encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode).

You can mix ANCP and PPPoE vendor-specific tags for dynamically instantiated static and interface sets so that the shaping rate is first using PPPoE vendor-specific tags and is later adjusted by ANCP. In this case, the shaping rate value overrides the PPPoE value.

To enable this feature, include the **actual-data-rate-downstream** or **access-loop-encapsulation** option with the **vendor-specific-tags** statement at the **[edit dynamic-profiles *profile-name* class-of-service dynamic-class-of-service-options]** hierarchy level.

[Subscriber Access, Class of Service]

- **Support for processing subscriber-initiated PPP fast keepalive requests (MX Series routers with MPCs/MICs)**—Enables the Packet Forwarding Engine on an MPC/MIC in an MX Series router to process and respond to Link Control Protocol (LCP) Echo-Request packets that the PPP subscriber (client) initiates and sends to the router. LCP Echo-Request packets and LCP Echo-Reply packets are part of the PPP keepalive mechanism that helps determine whether a link is functioning properly.

In earlier Junos OS releases, processing of LCP Echo-Request packets and LCP Echo-Reply packets was handled by the Routing Engine. In this release, the Packet Forwarding Engine on the MPC/MIC receives LCP Echo-Request packets from the PPP subscriber and transmits LCP Echo-Reply packets in response, without having to send the LCP packets to the Routing Engine for processing. The mechanism by which LCP Echo-Request packets are processed by the Packet Forwarding Engine instead of by the Routing Engine is referred to as *PPP fast keepalive*. Transmission of keepalive requests from the Packet Forwarding Engine on the router is not enabled in the current release.

Relieving the Routing Engine of having to process LCP Echo-Request packets provides increased bandwidth on the router to support a larger number of subscribers with improved performance.

No special configuration is required on an MX Series router with MPCs/MICs to enable processing of PPP fast keepalive requests on the Packet Forwarding Engine. The feature is enabled by default, and cannot be disabled.

When you issue the **show interfaces pp0.logical statistics** operational command to display interface statistics, the display does not include the number of keepalive packets processed or the amount of time since the router processed the last keepalive packet.

[Junos OS Subscriber Access Configuration Guide]

- **Support for DHCP Subscriber IP Session BFD Liveness Detection**—Liveness detection for DHCP subscriber IP sessions utilizes an active liveness detection protocol to institute liveness detection checks for relevant clients. Clients must respond to liveness detection requests within a specified amount of time. If the responses are not received within that time for a given number of consecutive attempts, then the liveness detection check fails and a failure action is implemented.

You can configure DHCP liveness detection either globally or per DHCPv4 or DHCPv6 group. To configure liveness detection, include the **liveness-detection** statement at the following hierarchy levels:

```
[edit system services dhcp-local-server],
[edit system services dhcp-local-server dhcpv6],
[edit forwarding-options dhcp-relay],
[edit forwarding-options dhcp-relay dhcpv6],
[edit system services dhcp-local-server group group-name],
[edit system services dhcp-local-server dhcpv6 group group-name],
[edit forwarding-options dhcp-relay group group-name],
[edit forwarding-options dhcp-relay dhcpv6 group group-name]
```

In Junos OS Release 12.1, the only method supported for liveness detection is Bidirectional Forwarding Detection (BFD). To configure BFD as the liveness detection method, include the **bfd** statement at the following hierarchy levels:

```
[edit system services dhcp-local-server liveness-detection method],
[edit system services dhcp-local-server dhcpv6 liveness-detection method],
[edit forwarding-options dhcp-relay liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 liveness-detection method],
[edit system services dhcp-local-server group group-name liveness-detection method],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection method],
[edit forwarding-options dhcp-relay group group-name liveness-detection method],
[edit forwarding-options dhcp-relay dhcpv6 group group-name liveness-detection method]
```

At the **bfd** hierarchy level, you can configure any of the following BFD-related statements:

• detection-time	• no-adaptation
• holddown-interval	• session-mode
• minimum-interval	• transmit-interval
• minimum-receive-interval	• version
- **multiplier**

After configuring the liveness detection method, configure the action the router takes when a liveness detection failure occurs by including the **failure-action** statement at the following hierarchy levels:

```
[edit system services dhcp-local-server liveness-detection],
[edit system services dhcp-local-server dhcpv6 liveness-detection],
[edit system services dhcp-relay liveness-detection],
[edit system services dhcp-relay dhcpv6 liveness-detection],
[edit system services dhcp-local-server group group-name liveness-detection],
[edit system services dhcp-local-server dhcpv6 group group-name liveness-detection],
[edit system services dhcp-relay group group-name liveness-detection],
[edit system services dhcp-relay dhcpv6 group group-name liveness-detection]
```

You can choose from the following three options when configuring a liveness detection failure option:

- **clear-binding**—The client session is cleared when a liveness detection failure occurs.

- **clear-binding-if-interface-up**—The client session is cleared only when a liveness detection failure occurs and the local interface is detected as being up.
- **log-only**—A message is logged to indicate the event; no action is taken and DHCP is left to manage the failure.

[Subscriber Access]

- **DHCPv6 Rapid Commit (M120, M320, and MX Series routers)**—The extended DHCPv6 local server supports the DHCPv6 Rapid Commit option (DHCPv6 option 14). When rapid commit is enabled on the extended DHCPv6 local server, the server supports a two-message exchange (Solicit and Reply) to configure clients, rather than the traditional four-method exchange (Solicit, Advertise, Request, and Reply). The two-message exchange provides faster client configuration, which is useful in environments in which client attachment points frequently change, such as mobile networks.

You can configure the DHCPv6 local server to support the Rapid Commit option globally, for a specific group, or for a specific interface. The DHCP client must also be configured to include the DHCPv6 Rapid Commit option in the Solicit messages sent to the DHCP local server.

To configure DHCP local server to support rapid commit, use the **rapid-commit** statement at the [edit dhcp-local-server dhcpv6 overrides] hierarchy level for global configuration, at the [edit dhcp-local-server dhcpv6 group group-name overrides] hierarchy level for group configurations, or at the [edit dhcp-local-server dhcpv6 group group-name interface interface-name overrides] hierarchy level for interface-specific configurations.

[Subscriber Access]

- **Support for interface-style services for PPPoE subscribers**—Starting with Junos OS Release 12.1, all interface-style services are supported for dynamic Point-to-Point Protocol over Ethernet (PPPoE) subscribers on all MX Series routers with Modular Port Concentrators (MPCs).

[Services Interfaces]

- **Support for rewrite rules and classifiers on Ethernet pseudowires (MX Series routers)**—Enables you to configure rewrite rules and classifiers on Ethernet pseudowires that are configured on logical tunnel interfaces. This feature is supported on MPC/MIC modules on MX Series routers.

Logical interfaces such as **lt-fpc/pic/port**, which are required to configure this feature, are created while configuring tunnel services for the router (using the **set chassis** command). For example, the **set chassis fpc 4 pic 0 tunnel-services** command creates a tunnel interface lt-4/0/0. You must then specify the Ethernet encapsulation type and **inet** as the family on the lt interface to configure this feature.

To configure the CoS parameters, include the **rewrite-rules** and **classifier** statements at the [edit class-of-service] hierarchy level. You can specify **inet-precedence** or **dscp** as the rewrite rule or the classification type.

You can use logical tunnel interfaces to create pseudowires by connecting two virtual routing forwarding (VRF) instances. A pseudowire can be used to represent a single subscriber (for example, a business subscriber).

[Class of Service, Subscriber Access]

- **Support for RADIUS attribute Local-Loopback-Interface for L2TP**—With Junos OS Release 12.1, the RADIUS attribute Local-Loopback-Interface [26-3] is supported for an L2TP LAC configured on an MX Series router.

You can configure the Local-Loopback-Interface attribute on a RADIUS server to manage multiple LAC devices. This attribute is used as the LAC source address on an LNS tunnel for PPPoE subscribers tunneled over L2TP.

When you use the Tunnel-Client-Endpoint attribute as the LAC source address, you must configure the Tunnel-Client-Endpoint attribute for each MX Series router that uses the same RADIUS server. Starting with Junos OS Release 12.1, you can use the Local-Loopback-Interface attribute, which needs to be configured only once.

When the LAC initiates an Access-Request message to RADIUS for authentication, RADIUS returns the Local-Loopback-Interface attribute in the Access-Accept message. This attribute contains the name of the loopback interface, either as a generic interface name such as “lo0” or as a specific name like “lo0.0.” The MX Series router then uses the configured loopback interface IP address as the source address during tunnel negotiation with the LNS.

[Broadband Subscriber Management Solutions]

- **Support for removing inactive dynamic subscriber VLANs (MX Series routers)**—Junos OS Release 12.1 supports the removal of dynamic subscriber VLANs when an inactivity threshold has been reached. To define the threshold at which a subscriber VLAN is removed, include the **client-idle-timeout** statement, along with the timeout value (from 10 through 1440 minutes), at the **[edit access profile *profile-name* session-options]** hierarchy level. The **client-idle-timeout** value specifies a maximum period of time that the subscriber can be idle. By default, no timeout is present.

In addition to removing inactive dynamic subscriber VLANs, the **client-idle-timeout** statement removes dynamic VLANs if no client sessions are ever created (for example, due to inactivity or error during creation or authentication).

When configuring dynamic VLAN removal upon inactivity timeout, keep the following in mind:

- The idle timeout period begins after a dynamic subscriber VLAN interface is created or traffic activity stops on a dynamic subscriber VLAN interface.
- If a new connection is created or a client session is reactivated successfully, the client idle timeout resets.
- The removal of inactive subscriber VLANs functions only with VLANs that have been authenticated.

[Subscriber Access, Network Interfaces]

- **Support for PTSP application on aggregated and redundant services PICs**—You can configure the packet-triggered subscribers and policy control (PTSP) feature on

aggregated Multiservices (ams) and redundant Multiservices (rms) interfaces. For the PTSP functionality, you can use a 1:1 redundancy model to pair two services PICs in high availability mode using a virtual redundant Multiservices PIC (rms) interface. You can also employ an N:1 redundancy mechanism to cluster together N number of ms-interfaces in an ams group that supports load sharing for PTSP subscribers.

With service PICs configured as rms interfaces, if a failover of a service PIC occurs, the subscribers on the failed service PIC are logged out, and the traffic is switched over to the redundant service PIC that initiates the subscribers to log in again.

With Multiservices interfaces configured in an ams group and traffic distributed over the service PICs in the ams group, the traffic of subscribers that are logged in to PTSP partitions on a particular service PIC is redistributed to other service PICs in the ams group. Such a replication of subscriber traffic for PTSP partitions enables the same subscriber detail to be present on different service PICs. In such a scenario, when one of the service PIC fails, the subscribers that are logged in over that PIC remain connected; the sessions are terminated only when the subscriber logs out or the idle timeout period is exceeded.

[Subscriber Access]

- **Enhanced show binding and clear binding commands for DHCP local server and DHCP relay agent (MX Series routers)**—The **show binding** and **clear binding** commands for extended DHCP local server and extended DHCP relay agent (including DHCPv6) have been enhanced to include additional options, which enable you to display or clear DHCP binding information by FPC, PIC, port, VLAN, and S-VLAN. The new options are supported on all underlying interfaces that support DHCP bindings.

The enhancement includes the following two new options:

- **<interfaces-vlan>**—The interface VLAN ID and S-VLAN ID on which to show or clear bindings.
- **<interfaces-wildcard>**—The set of interfaces on which to show or clear bindings. This option supports the use of the wildcard character (*), which enables you to identify interfaces based on FPC, PIC, and port.

The new options are supported by the operational commands shown in [Table 18 on page 269](#).

Table 18: Supported show and clear Commands

show dhcp server binding	show dhcp relay binding
show dhcpv6 server binding	show dhcpv6 relay binding
clear dhcp server binding	clear dhcp relay binding
clear dhcpv6 server binding	clear dhcpv6 relay binding



NOTE: IP demux interfaces are not supported by the show and clear DHCP bindings commands for DHCP local server and DHCP relay agent.

[Subscriber Access]

- **DHCPv6 relay agent support for DHCP snooping (M120, M320, and MX Series routers)**—Extends support for DHCP snooping to the DHCPv6 relay agent configured on the router. In multi-relay topologies where more than one DHCPv6 relay agent is between the IPv6 client and the DHCPv6 server, snooping enables intervening DHCPv6 relay agents between the client and the server to correctly receive and process the unicast traffic from the client and forward it to the DHCPv6 server.

The DHCPv6 relay agent snoops incoming unicast DHCPv6 packets by setting up a filter with UDP port 547 (the DHCPv6 UDP server port) on a per-forwarding table basis. The DHCPv6 relay agent then processes the packets intercepted by the filter and forwards the packets to the DHCPv6 server.

DHCPv6 relay agent snooping is disabled on the router by default. You can override the default DHCPv6 snooping configuration to explicitly enable or disable DHCPv6 snooping.

To enable snooping for the DHCPv6 relay agent, do one of the following:

- To enable DHCPv6 snooping support globally, include the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay dhcpv6 overrides]** hierarchy level.
- To enable DHCPv6 snooping support for a named group of interfaces, include the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides]** hierarchy level.
- To enable DHCPv6 snooping support for a specific interface within a named group of interfaces, include the **allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides]** hierarchy level.

To disable snooping for the DHCPv6 relay agent after you have enabled it, do one of the following:

- To disable DHCPv6 snooping support globally, include the **no-allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay dhcpv6 overrides]** hierarchy level.
- To disable DHCPv6 snooping support for a named group of interfaces, include the **no-allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides]** hierarchy level.
- To disable DHCPv6 snooping support for a specific interface within a named group of interfaces, include the **no-allow-snooped-clients** statement at the **[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides]** hierarchy level.

[Junos OS Subscriber Access Configuration Guide]

- **Offload status for PTSP and local L-PDF flows**—Starting with Junos OS Release 12.1, the **show services application-aware-access-list flows subscriber *subscriber-name*** and **show services subscriber flows client-id *client-id*** commands display flow offload status for each flow. Offloading is supported only on MX Series routers with Modular Port

Concentrators (MPCs) by the packet-triggered subscribers and policy control (PTSP) and local L-PDF plug-ins. No configuration is necessary to enable offloading.

[*Services Interfaces*]

- **Extended DHCPv4 local server and extended DHCPv4 relay agent and relay proxy support on M120 and M320 routers**—The extended DHCPv4 local server and extended DHCPv4 relay agent and relay proxy features used for subscriber management are supported and qualified on M120 and M320 routers. In addition to the configuration and command support for extended DHCPv4, M120 and M320 routers support the subscriber management features listed in [Table 19 on page 271](#).

Table 19: Additional Supported Features on M120 and M320 Routers

Ethernet interfaces on the following PICs: • 8-port Gigabit Ethernet IQ2 PIC (Type 3) • 2-port Gigabit Ethernet IQ PIC • 10-Gigabit Ethernet PIC with XENPAK • 10-Gigabit Ethernet IQ2 PIC with XFP	Aggregated Ethernet interfaces
MAC address validation	Service accounting (on FPCs that support bulk statistics)
Change of Authorization (CoA)	Filters (one input and one output filter per VLAN)
Class of Service (CoS), single-shaper per VLAN on IQ2 PICs	Non-default routing instances

[Table 20 on page 271](#) lists subscriber management features that are not supported on M120 and M320 routers.

Table 20: Unsupported Features on M120 and M320 Routers

IP and VLAN demux logical interfaces	Extended DHCPv6 local server and DHCPv6 relay agent
Layer 2 and Layer 3 wholesale	Subscriber secure policy traffic mirroring
Bidirectional Forwarding Detection (BFD)	Autosensed VLANs
NWay active link aggregation (LAG)	—

[*Subscriber Access*]

- **Support for link redundancy for CoS configured on an L2TP**—You can configure multiple ports on the same IQ2 and IQ2E PICs to support link redundancy for CoS on L2TP tunnels configured on an Ethernet interface.

Link redundancy is useful when the active port is unavailable due to events such as:

- Disconnection of a cable
- Rebooting of a remote end system

- Re-routing the traffic through a different port due to network conditions

When link redundancy is enabled, the traffic to the LAC devices is re-routed through another Ethernet interface configured on the same IQ2 or IQ2E PIC, and the L2TP sessions are maintained.

[*Subscriber Access*]

- **Support to display CoS statistics for L2TP session**—In Junos OS Release 12.1, the **show services l2tp cos-policer** command has been introduced to display CoS statistics for a policed or shaped L2TP session configured on an IQ2 or IQ2E PIC.

[*Subscriber Access*]

- **Support to display disconnect cause summary for L2TP sessions on M Series routers**—In Junos OS releases earlier than Release 12.1, there was no mechanism to view the statistics for the disconnect cause summary of L2TP sessions. The following new commands are now supported on M Series routers to view the disconnect cause summary statistics for an L2TP session:

- **show services l2tp disconnect-cause-summary**
- **clear services l2tp disconnect-cause-summary**

[*Subscriber Access*]

- **Support to display LCP statistics for an L2TP session**—You can view LCP statistics among other statistics for an L2TP session. The output of the following commands has been modified to include LCP statistics information:

- **show services l2tp session**
- **show services l2tp summary**

- **Support to display MLPPP statistics for an L2TP session**—You can view MLPPP (Multilink Point-to-Point Protocol) statistics along with other L2TP statistics for an L2TP session. The **show services l2tp multilink** command has been modified to display the MLPPP statistics for an L2TP session.

- **Default subscriber service (MX Series routers)**—Subscriber management enables you to specify a default subscriber service for DHCP subscribers. The default service (dynamic profile) is applied to subscribers who are not assigned a service by a remote server, such as RADIUS or a provisioning server (for example, a JSRC server or Gx-Plus server) when the subscriber logs in.

When a subscriber logs in, subscriber management examines the subscriber's access profile and uses a predetermined sequence to activate the subscriber's service.

To provide support for default subscriber services, use the following statements at the [**edit system services dhcp-local-server...**] and [**edit forwarding-options dhcp-relay...**] hierarchy levels:

- **service-profile *dynamic-profile-name***—Specifies that the service (dynamic profile) is used as the default subscriber service.
- **dynamic-profile *dynamic-profile-name***—Specifies that the default service is used globally, for a group of interfaces or for a specific interface.

[Subscriber Access]

- **CLI-based activation and deactivation of subscriber services (MX Series routers)**—Subscriber management enables you to use the Junos OS CLI to locally activate and deactivate dynamic subscriber services. CLI-based activation and deactivation provides local control for dynamic subscriber services that is similar to subscriber management's change of authorization (CoA) feature, which activates services remotely. Both the CoA and CLI-based methods enable you to manage services for subscribers who are currently logged in to the network—you can activate a new service for the subscriber or deactivate the subscriber's current service.

The CLI-based feature can activate a new service, but you cannot use it to modify a subscriber's dynamic profile instantiation or to modify a dynamic profile's user-defined variables, such as CoS variables.

[Subscriber Access]

- **Using the CLI to modify active CoS traffic-control profiles (MX Series routers)**—Subscriber management provides both global and per-subscriber methods for modifying traffic-control profiles that are in use. The global method modifies the traffic control profile for all subscribers currently using the profile. The per-subscriber method modifies the traffic control profile for a particular subscriber—all other subscribers currently using the traffic control profile remain unaffected.

To make a global modification for all current subscribers, use the **traffic-control-profiles** statement at the **[edit dynamic-profiles business-profile class-of-service]** hierarchy level to change one or more parameters for the existing traffic-control profile.

To make a per-subscriber modification for a specific subscriber, use the **request network-access aaa subscriber modify session-id** command and specify the name of the new traffic-control profile to use for the subscriber session.

[Subscriber Access]

System Logging

- **New and deprecated system log tags**—The following set of system log messages are new in this release:
 - **JTASK**—This chapter describes messages with the **JTASK** prefix.
 - **JTRACE**—This chapter describes messages with the **JTRACE** prefix.
 - **LACP**—This chapter describes messages with the **LACP** prefix. The Link Aggregation Control Protocol (LACP) provides a method to control the bundling of several physical ports together to form a single logical channel. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer (directly connected device that also implements LACP).
 - **SSL**—This chapter describes messages with the **SSL** prefix. They are generated by the Secure Sockets Layer (SSL) services. These logs contain information about logical system names, SSL proxy whitelists, policy information, and SSL proxy information.

SSL is an application-level protocol that provides encryption technology for the Internet. SSL, also called Transport Layer Security (TLS), ensures the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity. SSL relies on certificates and private-public key exchange pairs for this level of security.

- **VM**—This chapter describes messages with the **VM** prefix.

The following system log messages are new in this release:

- CHASSISD_FM_FABRIC_DEGRADED
- CHASSISD_FRU_FIRE_TEMP_CONDITION
- CHASSISD_FRU_HIGH_TEMP_CONDITION
- CHASSISD_FRU_OVER_TEMP_CONDITION
- CHASSISD_MASTER_CG_REMOVED
- CHASSISD_NO_CGS
- CHASSISD_PDU_BREAKER_TRIP
- CHASSISD_PDU_NOT_OK
- CHASSISD_PIC_PORT_ERROR
- CHASSISD_POWER_ON_CHECK_FAILURE
- CHASSISD_PSM_NOT_OK
- CHASSISD_VSERIES_LICENSE_ERROR
- CHASSISD_ZONE_BLOWERS_SPEED
- CHASSISD_ZONE_BLOWERS_SPEED_FULL
- EVENTD_CONFIG_CHANGE_FAILED
- EVENTD_CONFIG_CHANGE_SUCCESS
- JTASK_ACTIVE_TERMINATE
- JTASK_ASSERT
- JTASK_ASSERT_SOFT
- JTASK_CFG_CALLBACK_LONGRUNTIME
- JTASK_CFG_SCHED_CUMU_LONGRUNTIME
- JTASK_EXIT
- JTASK_LOCK_FLOCKED
- JTASK_LOCK_LOCKED
- JTASK_MGMT_TIMEOUT
- JTASK_OS_MEMHIGH
- JTASK_PARSE_BAD_LR_NAME

- JTASK_PARSE_BAD_OPTION
- JTASK_PARSE_CMD_ARG
- JTASK_PARSE_CMD_DUPLICATE
- JTASK_PARSE_CMD_EXTRA
- JTASK_PTHREAD_CREATE
- JTASK_SCHED_CUMU_LONGRUNTIME
- JTASK_SCHED_MODULE_LONGRUNTIME
- JTASK_SCHED_TASK_LONGRUNTIME
- JTASK_SIGNAL_TERMINATE
- JTASK_SNMP_CONN_EINPROGRESS
- JTASK_SNMP_CONN_QUIT
- JTASK_SNMP_CONN_RETRY
- JTASK_SNMP_INVALID_SOCKET
- JTASK_SNMP_SOCKOPT_BLOCK
- JTASK_SNMP_SOCKOPT_RECVBUF
- JTASK_SNMP_SOCKOPT_SENDBUF
- JTASK_START
- JTASK_SYSTEM
- JTASK_TASK_CHILDKILLED
- JTASK_TASK_CHILDDSTOPPED
- JTASK_TASK_DYN_REINIT
- JTASK_TASK_FORK
- JTASK_TASK_GETWD
- JTASK_TASK_MASTERSHIP
- JTASK_TASK_NOREINIT
- JTASK_TASK_PIDCLOSED
- JTASK_TASK_PIDFLOCK
- JTASK_TASK_PIDWRITE
- JTASK_TASK_REINIT
- JTASK_TASK_SIGNALIGNORE
- JTRACE_FAILED
- L2TPD_SESSION_IP_DUPLICATE
- LACP_INTF_DOWN

- RPD_PIM_FOUND_NON_BIDIR_NBR
- RPD_PIM_NON_BIDIR_RPF
- RT_SCREEN_WHITE_LIST
- RT_SCREEN_WHITE_LIST_LS
- SNMPD_ENGINE_ID_CHANGED
- SSL_PROXY_ERROR
- SSL_PROXY_INFO
- SSL_PROXY_SESSION_IGNORE
- SSL_PROXY_SESSION_WHITELIST
- SSL_PROXY_SSL_SESSION_ALLOW
- SSL_PROXY_SSL_SESSION_DROP
- SSL_PROXY_WARNING
- UFDD_GROUP_ACTION_COMPLETE
- UI_RESTART_FAILED_EVENT
- VM_DCF_PB_COMMUNICATION_FAILED
- VM_DCF_PB_INVALID_IMAGE
- VM_DCF_PB_INVALID_UUID
- VM_DCF_PB_RESOURCE_FAILURE

The following system log messages are no longer documented, either because they indicate internal software errors that are not caused by configuration problems or because they are no longer generated. If these messages appear in your log, contact your technical support representative for assistance:

- CFMD_PARSE_BAD_SWITCH
- CFMD_PARSE_CMD_ARG
- ESWD_SYSTEM_CALL_FAILED
- KMD_PM_DUPLICATE_LIFE_DURATION
- KMD_PM_IKE_SERVER_LOOKUP_FAILED
- KMD_PM_IKE_SERVER_NOT_FOUND
- KMD_PM_ILLEGAL_REMOTE_GW_ID
- KMD_PM_INCONSISTENT_P2_IDS
- KMD_PM_INVALID_LIFE_TYPE
- KMD_PM_NO_LIFETIME
- KMD_PM_NO_LIFE_TYPE
- KMD_PM_NO_PROPOSAL_FOR_PHASE1

- KMD_PM_NO_SPD_PHASE1_FUNC_PTR
- KMD_PM_P1_POLICY_LOOKUP_FAILURE
- KMD_PM_P2_POLICY_LOOKUP_FAILURE
- KMD_PM_SA_PEER_ABSENT
- KMD_PM_SPI_DELETE_REJECT
- KMD_PM_UNEQUAL_PAYLOAD_LENGTH
- KMD_PM_UNINITIALISE_ERROR
- KMD_PM_UNKNOWN_P1_IDENTITIES
- KMD_PM_UNKNOWN_PHASE2_ENTITIES
- KMD_PM_UNKNOWN_QM_NOTIFICATION
- KMD_PM_UNSUPPORTED_KEY
- KMD_PM_UNSUPPORTED_MODE
- LIBJNX_AUDIT_ERROR
- LIBJNX_COMPRESS_EXEC_FAILED
- LIBJNX_DEFAULT_IP_ADDR_NOT_SET
- LIBJNX_EVLIB_FAILURE
- LIBJNX_EXEC_EXITED
- LIBJNX_EXEC_FAILED
- LIBJNX_EXEC_PIPE
- LIBJNX_EXEC_SINGALED
- LIBJNX_EXEC_WEXIT
- LIBJNX_FILE_COPY_FAILED
- LIBJNX_FILE_SYSTEM_FAIL
- LIBJNX_FILE_SYSTEM_SPACE
- LIBJNX_INVALID_CHASSIS_ID
- LIBJNX_INVALID_RE_SLOT_ID
- LIBJNX_INVALID_XML_DATA
- LIBJNX_LOGIN_ACCOUNTS_NOT_LOCKED
- LIBJNX_LOGIN_ACCOUNT_LOCKED
- LIBJNX_LOGIN_ACCOUNT_NOT_LOCKED
- LIBJNX_LOGIN_ACCOUNT_UNLOCKED
- LIBJNX_LOGIN_ACCT_NOT_UNLOCKED
- LIBJNX_PRIV_LOWER_FAILED

- LIBJNX_PRIV_RAISE_FAILED
- LIBJNX_REPLICATE_RCP_ERROR
- LIBJNX_REPLICATE_RCP_EXEC_FAILED
- LIBJNX_SNMP_ENGINE_SCAN_FAILURE
- LIBJNX_SOCKET_FAILURE
- LIBJNX_XML_DECODE_FAILED
- RPD_RT_CFG_EIBGP_VTL_CONFLICT
- SNMP_GET_ERROR1
- SNMP_GET_ERROR2
- SNMP_GET_ERROR3
- SNMP_GET_ERROR4
- SNMP_TRAP_LINK_DOWN
- SNMP_TRAP_LINK_UP
- SSH_RELAY_SERVER_ERROR
- UFDD_SYSTEM_CALL_FAILED
- UI_AUTH_BAD_LOCATION
- UI_AUTH_BAD_TIME
- UI_BOOTTIME_FAILED
- UI_CFG_AUDIT_NEW
- UI_CFG_AUDIT_OTHER
- UI_CFG_AUDIT_SET
- UI_CFG_AUDIT_SET_SECRET
- UI_CHILD_ARGS_EXCEEDED
- UI_CHILD_CHANGE_USER
- UI_CHILD_EXEC
- UI_CHILD_EXITED
- UI_CHILD_FOPEN
- UI_CHILD_OUTPUT
- UI_CHILD_PIPE_FAILED
- UI_CHILD_STOPPED
- UI_CHILD_WAITPID
- UI_CLASS_MODIFIED_USERS
- UI_CMD_AUTH_REGEX_INVALID

- UI_COMMIT
- UI_COMMIT_AT_COMPLETED
- UI_COMMIT_AT_FAILED
- UI_COMMIT_COMPRESS_FAILED
- UI_COMMIT_CONFIRMED_REMINDER
- UI_COMMIT_EMPTY_CONTAINER
- UI_COMMIT_PROGRESS
- UI_COMMIT_SYNC_FORCE
- UI_COND_GROUPS
- UI_CONFIGURATION_ERROR
- UI_CONFIGURATION_WARNING
- UI_DBASE_ACCESS_FAILED
- UI_DBASE_CHECKOUT_FAILED
- UI_DBASE_EXTEND_FAILED
- UI_DBASE_LOGIN_EVENT
- UI_DBASE_LOGOUT_EVENT
- UI_DBASE_MISMATCH_EXTENT
- UI_DBASE_MISMATCH_MAJOR
- UI_DBASE_MISMATCH_MINOR
- UI_DBASE_MISMATCH_SEQUENCE
- UI_DBASE_MISMATCH_SIZE
- UI_DBASE_OPEN_FAILED
- UI_DBASE_REBUILD_FAILED
- UI_DBASE_REBUILD_SCHEMA_FAILED
- UI_DBASE_REBUILD_STARTED
- UI_DBASE_RECREATE
- UI_DBASE_REOPEN_FAILED
- UI_DUPLICATE_UID
- UI_INVALID_REMOTE_PERMISSION
- UI_JUNOSCRIPT_ERROR
- UI_LOAD_EVENT
- UI_LOAD_JUNOS_DEFAULT_FILE_EVENT
- UI_LOGIN_EVENT

- UI_LOGOUT_EVENT
- UI_LOST_CONN
- UI_MOTD_PROPAGATE_ERROR
- UI_NETCONF_ERROR
- UI_READ_FAILED
- UI_READ_TIMEOUT
- UI_SCHEMA_CHECKOUT_FAILED
- UI_SCHEMA_MISMATCH_MAJOR
- UI_SCHEMA_MISMATCH_SEQUENCE
- UI_SCHEMA_SEQUENCE_ERROR
- UI_VERSION_FAILED

[*System Log*]

User Interface and Configuration

- **Support for CLI edit mode wildcard range command**—The wildcard range command enables you to specify ranges in activate, deactivate, delete, protect, set, show, and unprotect commands. You can use ranges to specify a range of interfaces, logical units, VLANs, and other numbered elements. The wildcard range command expands the command you entered into multiple commands, each of which corresponds to one item in the range. For example, wildcard range interfaces deactivate ge-0/0/[1-3] expands to deactivate interfaces ge-0/0/1, deactivate interfaces ge-0/0/2, and deactivate interfaces ge-0/0/3.

[*CLI User Guide*]

- **Support for batch commits**—The batch commit feature aggregates or merges multiple configuration edits from different CLI sessions or users and adds them to a batch commit queue. A commit server running on the device takes one or more jobs from the batch commit queue, applies the configuration changes to the shared configuration database, and then commits the configuration changes in a single commit operation. When compared to the regular commit operation where all commits are independently committed sequentially, batch commits save time and system resources by committing multiple small configuration edits in a single commit operation.

Batch commits are performed from the **[edit batch]** hierarchy level in the configuration mode. The commit server properties can be configured at the **[edit system commit server]** hierarchy level.

[*CLI User Guide*]

- **Support for 16G Routing Engine (RE-DUO-C1800-16G) on standalone T4000 routers**—The T4000 router supports 64-bit Junos OS only. Before installing RE-DUO-C1800-16G in a T4000 router, it must be running 64-bit Junos OS. Except during upgrade, both Routing Engines must have the same model number.

The output for the following show commands has been updated to reflect the new Routing Engine:

- **show chassis routing engine**—Shows available memory (16G) in output.
- **show chassis hardware**—Shows FRU identification and displays the correct Routing Engine type and model number in the output.
- **show system alarms**—Shows alarms currently active in output.

[*T4000 Hardware Guide*]

[*System Basics and Services Command Reference*]

- **Support for Chassis Management (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**—The Application Services Modular Line Card (AS MLC) is a Modular Port Concentrator (MPC) that is designed to run services and applications on MX240, MX480, and MX960 routers.

The following CLI operational mode commands support the chassis management operations of the AS Modular Carrier Card on the AS MLC:

- show chassis environment fpc
- show chassis firmware
- show chassis fpc
- show chassis hardware
- show chassis pic
- show chassis temperature-thresholds
- request chassis fpc
- request chassis mic
- request chassis mic fpc-slot mic-slot

[*System Basics, Junos OS System Basics and Services Command Reference*]

- **Support for fabric management features (MX240, MX480, MX960 Routers with Application Services Modular Carrier Card)**—The Application Services Module Line Card (AS MLC) is supported on MX240, MX480, and MX960 routers. The AS MLC consists of the following components:
 - Application Services Modular Carrier Card (AS MCC)
 - Application Services Modular Processing Card (AS MXC)
 - Application Services Modular Storage Card (AS MSC)

The AS MCC plugs into the chassis and provides the fabric interface. On the fabric management side, the AS MLC provides redirection functionality using a demultiplexer. The following CLI operational mode commands display fabric-related information for the AS MCC:

- show chassis fabric fpcs
- show chassis fabric map

- show chassis fabric plane
- show chassis fabric plane-location
- show chassis fabric reachability
- show chassis fabric summary

[*Junos OS System Basics Configuration Guide, Junos OS System Basics and Services Command Reference*]

VPNs

- **Support for Layer 3 VPN composite next hop (T4000 routers)**—Layer 3 VPN composite next hop is supported on T4000 Type 5 FPCs. Next-hop chaining (also known as composite next hop) is a composition function that concatenates the partial rewrite strings associated with individual next hops to form a larger rewrite string that is added to a packet.

A chained next hop contains the inner label and the downstream indirect next-hop destination information. The outer labels and Layer 2 rewrite bytes are associated with the unicast next hop. The T4000 Type 5 FPC Packet Forwarding Engine supports full chaining—that is, the ingress Packet Forwarding Engine adds the inner VPN label to the packet and sends the modified packet, along with a token that corresponds to the unicast next hop, to the egress Packet Forwarding Engine. The egress Packet Forwarding Engine adds the transport label and Layer 2 encapsulation information based on the token to the packet.

For Layer 3 VPNs configured on Juniper Networks routers, Junos OS normally allocates one inner VPN label for each VPN network on the customer edge (CE)-facing interfaces of a provider edge (PE) router. However, other vendors allocate one VPN label for each BGP route on the CE-facing interfaces of a PE router. This practice increases the number of VPN labels exponentially, which leads to slow system processing and slow convergence time. To account for this difference, configure the **l3vpn-composite-nexthop** statement at the **[edit routing-options]** hierarchy level on the Juniper Networks routers participating in a mixed vendor network. The **l3vpn-composite-nexthop** statement is disabled by default.

[VPNs]

- **EBGP and IBGP load-balancing support for MVPN and PIM**—The multipath PIM join load-balancing feature enables customer PIM (C-PIM) join messages to be load-balanced across unequal EBGP and IBGP paths in a draft-rosen MVPN, and across all available IBGP paths in a next-generation MVPN that does not have any EBGP upstream path toward the source or RP.
- **Load balancing and IP header filtering for Layer 3 VPNs**—You can simultaneously enable load balancing and IP header filtering for traffic in a network with both internal and external BGP paths. To enable these features, include the **equal-external-internal** statement at the **[edit routing-instances routing-instance-name routing-options multipath vpn-unequal-cost]** hierarchy level and the **vrf-table-label** statement at the **[edit routing-instances routing-instance-name]** hierarchy level.

[VPNs]

- **Inter-AS multicast Layer 3 VPNs**—You can configure multicast Layer 3 VPNs (also known as multiprotocol BGP (MBGP)-based multicast VPNs) across autonomous systems (ASs). Previously, you could only configure unicast Layer 3 VPNs across ASs. Although there are a number of different network configurations that can be used to enable inter-AS support for multicast VPN traffic, Junos OS Release 12.1 supports only next generation VPN option A and option C. For both option A and option C, the customer service provider depends on the VPN service provider to deliver a VPN transport service between the customer service provider's points of presence (POPs) or regional networks. This functionality might be used by a VPN customer who has connections to several different Internet service providers (ISPs), or different connections to the same ISP in different geographic regions, each of which has a different AS number.
 - Option A—In this implementation, the VPN routing and forwarding (VRF) table in the ASBR of one AS is linked to the VRF table in the ASBR in the other AS. Each ASBR must contain a VRF instance for every VPN configured in both service provider networks. In addition, PIM needs to be configured between the VRF instances, and IGP or BGP must be configured between the ASBRs. Option A is a relatively simple interprovider VPN solution. However, it is less scalable relative to option C.
 - Option C—In this implementation, only routes internal to the service provider networks are announced between ASBRs. This is achieved by using the **family inet labeled-unicast** statements in the IBGP and EBGP configuration on the PE routers. Labeled IPv4 (not VPN-IPv4) routes are exchanged by the ASBRs to support MPLS. An MP-EBGP session between the end PE routers is used for the announcement of VPN-IPv4 routes. In this manner, VPN connectivity is provided while keeping VPN-IPv4 routes out of the core network.

The existing configuration statements and documented procedures for the implementation of interprovider Layer 3 VPNs option A and option C can be applied to interprovider multicast Layer 3 VPNs. Of course, interprovider multicast Layer 3 VPNs requires the configuration of the multicast features in addition to the interprovider features.

[VPNs]

Related Documentation

- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 284](#)
- [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 298](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 404](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 432](#)

Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers

- [Changes in Default Behavior and Syntax on page 284](#)
- [Changes Planned for Future Releases on page 297](#)

Changes in Default Behavior and Syntax

The following are changes made to Junos OS default behavior and syntax.

- [High Availability on page 284](#)
- [Interfaces and Chassis on page 285](#)
- [J-Web on page 288](#)
- [Junos XML API and Scripting on page 288](#)
- [MPLS Applications on page 289](#)
- [Multicast on page 289](#)
- [Network Address Translation \(NAT\) on page 289](#)
- [Routing Protocols on page 290](#)
- [Services Applications on page 290](#)
- [Subscriber Access Management on page 292](#)
- [User Interface and Configuration on page 296](#)
- [VPNs on page 297](#)

High Availability

- **Updates to command forwarding support for MX Series Virtual Chassis (MX240, MX480, and MX960 routers with MPC/MIC interfaces)**—The following operational commands support command forwarding in an MX Series Virtual Chassis configuration:
 - **show chassis craft-interface**
 - **show chassis fan**
 - **show chassis pic**
 - **show chassis power**

Command forwarding enables you to monitor and manage an MX Series Virtual Chassis as a single network element by running operational commands on a specific member router in the Virtual Chassis, on both member routers, or on the local member router from which you issue the command. With command forwarding, the router sends the command to the specified member router or routers and displays the results as if the command were processed on the local router.

To forward operational commands to one or both member routers in an MX Series Virtual Chassis, use one of the following options when you issue these commands:

- To forward the command to a specified member router, use the **member member-id** option, where *member-id* can be 0 or 1.

- To run the command on the local member router on which the command was issued, use the **local** option.
- To forward the command to both member routers, use the **all-members** option. This is the default command forwarding option for the **show chassis craft-interface**, **show chassis fan**, **show chassis pic**, and **show chassis power** commands.

[*Junos OS System Basics and Services Command Reference, Junos OS High Availability Configuration Guide*]

Interfaces and Chassis

- **MIC for non-queuing MPC**—MIC comes online for non-queuing MPC even though it is not supported for non-queuing MPC (MX Series routers).
- The **allow-sram-parity-errors** statement is made visible at the **[edit chassis fpc slot-number]** hierarchy level (for T Series routers only).

[*System Basics*]

- **New command to monitor PPP recovery after a GRES or restart (MX Series 3D Universal Edge Routers)**—The **show ppp statistics recovery** command monitors the progress of PPP recovery after a GRES or restart. When the PPP subscriber sessions have been recovered, the command output displays **Recovery state: recovery done** to indicate that it is safe to force another GRES or restart. When you issue this command during the recovery process, the command might time out or fail silently rather than display output. Recovery is not complete until the command displays recovery done.

[*Interfaces Command Reference*]

- **Options for Multichassis Link Aggregation (MC-LAG)**—For MC-LAG, you can specify one of two actions to take if the Inter-Chassis Communication Protocol (ICCP) peer of the switch or router goes down. To bring down the interchassis link logical interface if the peer goes down, include the **force-icl-down** statement at the **[edit interfaces aeX aggregated-ether-options events iccp-peer-down]** hierarchy level. To have the router or switch become the active node when a peer goes down, include the **prefer-status-control-active** statement at the **[edit interfaces aeX aggregated-ether-options mc-ae events iccp-peer-down]** hierarchy level. When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement at the **[edit interfaces aeX aggregated-ether-options-mc-ae]** hierarchy level. If you do not configure the **status-control** as **active** with the **prefer-status-control-active** statement, the router or switch does not become the active node if a peer goes down.

[*Junos OS Ethernet Interfaces Configuration Guide*]

- **Enhancement to show interfaces queue command**—The output for the **show interfaces queue** command displays rate-limit statistics for class-of-service schedulers for all IQ and Enhanced IQ (IQ2E) PICs when rate-limiting is configured, even when no traffic is dropped. When rate limiting is configured but no traffic is dropped, the output for the **RL-dropped packets** and **RL-dropped-bytes** fields display the value zero (0). Previously, these fields were not displayed when no traffic was dropped and rate-limiting was configured. To configure rate-limiting for queues before packets are queued for output,

you include the **rate-limit** statement at the **[edit class-of-service schedulers transmit-rate rate]** hierarchy level.

[Interfaces Command Reference]

- **Multichassis Link Aggregation (MC-LAG)**—When you configure the **prefer-status-control-active** statement at the **[edit interfaces aex aggregated-ether-options mc-ae events iccp-peer-down]** hierarchy level, you must also configure the **status-control active** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning.

[Junos OS Ethernet Interfaces Configuration Guide]

- **Enhancement to output for show chassis fabric plane extensive command**—(TX Matrix Plus Routers) The output for the **show chassis fabric plane extensive** command includes more detailed information about link errors caused by CRC saturation. The output displays whether a CRC-related link error occurred under the following two conditions:
 - CRC exceeded the rate threshold and reached saturation without optical errors, that is, no cable has been cut, removed, or otherwise experienced an error. In this situation, the output displays the following: **Link error crc saturated**.
 - CRC exceeded the rate threshold and reached saturation with optical errors, that is, a cable has been cut, removed, or otherwise experienced an error. In this situation, the output displays the following: **Link error crc saturated with optical errors**.

If a link error is caused when the CRC exceeds the rate threshold but does not reach saturation, the output for the command continues to display the following: **Link error**.

[Interfaces Command Reference]

- **Enhancement to set date ntp command**—You can specify an authentication-key number for the NTP server used to synchronize the date and time on the router or switch. Include the new **key number** option with the **set date ntp** command. The key number you include must match the number you configure for the NTP server at the **[edit system ntp authentication-key number]** hierarchy level.
- **New Link Aggregation Control Protocol (LACP) commands and SNMP MIB**—You can view and clear LACP timeout entries. To display information about LACP timeout entries, use the **show lacp timeouts** command. Include the **interfaces interface-name** option to view timeout information about a specific interface only. To clear LACP timeout entries, use the **clear lacp timeouts** command. Include the **interfaces interface-name** option to clear timeout information for a specific interface only. A new SNMP MIB is now also available. The **jnxLacpAggTimeout** MIB lists all interfaces where the **jnxLacpTimeOut** trap is sent.

[Interfaces Command Reference]

- **Enhancement to Link Layer Discovery Protocol (LLDP)**—(MX Series and T Series routers) You can configure LLDP to generate the interface name as the port ID Type TLV. To generate the interface name as the port ID Type, Length, and Value, include the **interface-name** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level. The default behavior is to generate the SNMP Index of the interface as the port

ID TLV. If you have changed the default behavior, include the **locally-assigned** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level to reenabling the default behavior of generating the SNMP Index of the interface as the port ID TLV. When you configure LLDP to generate the interface name as the port ID TLV on the remote neighbor, the **show lldp neighbors** command displays the interface name in the **Port ID** field. The default behavior is for the command to display the SNMP index of the interface of the remote neighbor in the **Port ID** field.

[Ethernet Interfaces Configuration Guide, Interfaces Command Reference]

- **Configuring the flow-tap service for IPv6 traffic:** The **family inet | inet6** statement at the **[edit services flow-tap]** hierarchy enables you to specify the type of traffic for which you want to apply the flow-tap service. If the family statement is not included, the flow-tap service is, by default, applied to the IPv4 traffic. To apply flow-tap service to IPv6 traffic, you must include the **family inet6** statement in the configuration. To enable the flow-tap service for IPv4 and IPv6 traffic, you must explicitly configure the family statement for both **inet** and **inet6** families.

However, you cannot configure the flow-tap service for IPv6 along with port mirroring or sampling of IPv6 traffic on routers that support LMNR-based FPCs. This restriction is true even if the router does not have any LMNR-based FPC installed on it. There is no restriction on configuring the flow-tap service on routers that are configured for port mirroring or sampling of IPv4 traffic.

- On MX80 routers, the FPC Slot output field has been changed to TFEB Slot for the **show services accounting flow inline-jflow**, **show services accounting errors inline-jflow**, and **show services accounting status inline-jflow** commands.
- **Changes to DDoS protocol groups (MX Series routers)**—The **ipv4-unclassified** and **ipv6-unclassified** DDoS protocol groups have been deprecated in the **protocols** statement at the **[edit system ddos-protection ddos]** hierarchy level. These two protocol groups have also been deprecated from the **show ddos-protection protocols** commands. These groups formerly were used to police all unclassified IPv4 and IPv6 host-bound traffic.

In their place, 10 new protocol groups have been added to the **protocols** statement and the **show ddos-protection protocols** commands:

- **control-layer2**—Unclassified layer 2 control packets.
- **control-v4**—Unclassified IPv4 control packets.
- **control-v6**—Unclassified IPv6 control packets.
- **filter-v4**—Unclassified IPv4 filter action packets; sent to the host because of reject terms in firewall filters.
- **filter-v6**—Unclassified IPv6 filter action packets; sent to the host because of reject terms in firewall filters.
- **host-route-v4**—Unclassified IPv4 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
- **host-route-v6**—Unclassified IPv6 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.

- **other**—All unclassified packets that do not belong to another type.
- **resolve-v4**—Unclassified IPv4 resolve packets sent to the host because of a traffic request resolve action.
- **resolve-v6**—Unclassified IPv6 resolve packets sent to the host because of a traffic request resolve action.

[DDoS Configuration]

- **Configuration support to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure**—You can now configure the **prefer-status-control-active** statement with the **status-control standby** configuration at the **[edit interfaces aeX aggregated-ether-options mc-ae]** hierarchy level to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP does not go down unless the router is down. You must also configure the **hold-time down** value (at the **[edit interfaces interface-name]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents traffic loss by ensuring that when the router with the **status-control active** configuration goes down, the router with the **status-control standby** configuration does not go into standby mode.

J-Web

- On all M Series, MX Series, and T Series platforms, the username field does not accept HTML tags or the < and > characters. The following error message appears: **A username cannot include certain characters, including < and >.**

Junos XML API and Scripting

- **<get-configuration> RPC with inherit="inherit" attribute returns correct time attributes for committed configuration**—In Junos OS Release 12.1R5 and earlier releases, when you configured some interfaces using the interface-range configuration statement, if you later requested the committed configuration using the <get-configuration> RPC with the **inherit="inherit"** and **database="committed"** attributes, the device returned **junos:changed-localtime** and **junos:changed-seconds** in the RPC reply instead of **junos:commit-localtime** and **junos:commit-seconds**. This issue is fixed in Junos OS Release 12.1R6 and later releases so that the device returns the expected attributes in the RPC reply.
- **IPv6 address text representation is stored internally and displayed in command output using lowercase**—Starting with Junos OS Release 11.1R1, IPv6 addresses are stored internally and displayed in the command output using lowercase. Scripts that match on an uppercase text representation of IPv6 addresses should be adjusted to either match on lowercase or perform case-insensitive matches.

MPLS Applications

- Starting in Junos OS Release 12.1, for point-to-multipoint LSPs, the output of the **show route** command does not show multiple sub-LSPs sharing the forwarding branch. In Junos OS Release 11.4 and earlier, the **show route** command showed the LSP name on the penultimate-hop and egress routers.

Multicast

- In a bootstrap router (BSR)-enabled bidirectional PIM domain, mixing Junos OS Release pre-12.1R7 releases and later releases can cause unexpected outages. If you have a deployment with routers running Junos OS Release pre-12.1R7 and if you upgrade a subset of the routers to Junos OS release 12.1R7 or later, the group-to-RP mapping across the domain breaks and an outage occurs.

Network Address Translation (NAT)

- **Protection of MX Series, M Series, and T Series routers from denial of service (DOS) attacks**—New CLI options provide improved protection against DOS attacks.
 - NAT mapping refresh behavior—Prior to Release 12.1R7, a conversation was kept alive when either inbound or outbound flows were active. This remains the default behavior. As of this release, you can also specify mapping refresh for only inbound flows or only outbound flows. To configure mapping refresh behavior, include the **mapping-refresh (inbound | outbound | inbound-outbound)** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.
 - EIF inbound flow limit—Previously the number of inbound connections on an EIF mapping was limited only by the maximum flows allowed on the system. You can now configure the number of inbound flows allowed for an EIF. To limit the number of inbound connections on an EIF mapping, include the **EIF-flow-limit number-of-flows** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.

[Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions]

- The method for computing the block size for deterministic port block allocation for network port translation (NAPT) when the configured block size is zero has changed, and is computed as follows:

$$\text{block-size} = \text{int}(\text{ceil}[(\text{Nr\_Addr\_PR\_Prefix}/\text{Nr\_Addr\_PU\_Prefix})])$$

where:

64512 is the maximum available port range per public IP address.

Nr_Addr_PR_Prefix is the number of usable pre-NAT IPv4 subscriber addresses in a **from** clause match condition

Nr_Addr_PU_Prefix is the number of usable post-NAT IPv4 addresses configured in the NAT pool

Routing Protocols

- When a route with an improved metric is added to the IS-IS internal routing table, IS-IS flushes all next-hop information (including LSP next-hop information) for a route. To override this default behavior for load balancing, the **lsp-equal-cost** statement is added at the **[edit protocols isis traffic-engineering multipath]** hierarchy level to retain the equal cost path information in the routing table.

[Routing Protocols]

- **Enhancements to DDoS Protection (MX Series routers)**—The configuration of DDoS Protection has changed slightly. You can now include the **disable-fpc** statement at the **[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]** hierarchy level to disable policers on all line cards for a particular packet type or aggregate within a protocol group. The ability to configure this statement globally or for a particular line card remains unchanged.

You can also now include the **disable-logging** statement at the **[edit system ddos-protection protocols protocol-group (aggregate | packet-type)]** hierarchy level to disable router-wide logging of DDoS violation events for a particular packet type or aggregate within a protocol group. The **disable-logging** statement has been removed from the **[edit system ddos-protection violation]** hierarchy level, and that hierarchy level has been removed from the CLI.

The **show ddos-protection protocols** command displays **Partial** in the **Enabled** field to indicate when some of the instances of the policer are disabled. The **Routing Engine Information** section of the output includes fields for bandwidth, burst, and state.

The **show ddos-protection protocols parameters** command and the **show ddos-protection protocols statistics** command include a **terse** option to display information only for active protocol groups—that is, groups that show traffic in the **Received (packets)** column. The **show ddos-protection protocols parameters** command also displays **part** for policers that have some disabled instances.

[DDoS Protection]

- If you configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**. To avoid a conflict in the two route distinguisher values, we recommend ensuring that the first half of the route distinguisher obtained by configuring the **route-distinguisher** statement is different from the first half of the route distinguisher obtained by configuring the **route-distinguisher-id** statement.

[Routing Protocols Guide]

Services Applications

- **Reusability of NAT source pools between service sets**—Prior to Junos OS Release 11.4R3, you could only use a source NAT pool in a single service set. As of Junos OS Release 11.4R3 and subsequent releases, you can reuse a source NAT pool in multiple service sets, provided that the service interfaces associated with the service sets are in different virtual routing and forwarding (VRF) instances. To enable sharing, include the **allow-overlapping-nat-pools** statement at the **[edit services nat]** hierarchy level.

The service sets that are associated with the reused NAT pool cannot use service interfaces in the same VRF (virtual routing and forwarding) instance.

Use of the **allow-overlapping-nat-pools** option prevents the following commit errors:

- For translation types **basic-nat44**, **napt-44**, or **dynamic-nat44** when the same NAT pool is configured across multiple service sets, the commit error:

NAT pool *pool name* is already used with service set *service set name*

- For translation type **napt-44**, when the NAT pool 1 address (or address-range) overlaps with pool 2 configured across multiple service sets (for example: **rule1, pool1** under **service-set1**, and **rule2, pool2** under **service-set2**), the commit error:

NAT pool *<pool name 1>* overlaps with NAT pool *<pool name 2>* used by service set *<service set name>*

- For translation type **basic-nat44**, when the same NAT pool is configured across multiple service sets (one rule in one service set and another rule in another service set), the commit error:

With translation-type *basic_nat_44*, same pool cannot be shared by multiple rules or terms

- For translation type **basic-nat44**, when the NAT pool 1 address (or address-range) overlaps with pool 2 configured across multiple service sets (for example: **rule1, pool1** under **service-set1**, and **rule2, pool2** under **service-set2**), the commit error:

static nat pools *<pool name 1>* and, *<pool name 2>* overlap

- For translation type **basic-nat44**, when the NAT pool 1 address (or address-range) overlaps with pool 2 configured in a single service set (for example: **pool1** under term 1 and **pool2** under term 2 of **rule1**), the commit error:

static nat pools *<pool name 1>* and *<pool name 2>* overlap

- For translation type **dynamic-nat44**, when the same NAT pool is configured across multiple service sets, the commit error:

NAT pool *<pool name>* is already used by service set *<service set name>*

- For translation type **dynamic-nat44**, when the NAT pool 1 address (or address-range) overlaps with **pool2** configured across multiple service sets (for example: **rule1, pool1** under **service-set1**, and **rule2, pool2** under **service-set2**), the commit error: .

Dynamic nat pools *<pool name 1>* and *<pool name 2>* overlap

- For translation type **basic-nat44**, when the NAT pool 1 address (or address-range) overlaps with pool 2 configured in a single service-set (for example: **pool1** under term 1 and **pool2** under term 2 of **rule1**), the commit error:

Dynamic nat pools *<pool name 1>* and *<pool name 2>* overlap

[*Next-Generation Network Addressing Solutions*]

- Option available to clear the DF (do not fragment) flag in packet headers under NAT 64**—You can specify clearing of the DF flag in IPv4 headers when performing NAT64 translation. The flag is cleared when packets are less than 1280 bytes but might need to be fragmented when an IPv4-path-mtu is less than 1280. When the flag is not cleared, a **packet-too-big** error message is generated and sent to IPv6 host, which

replies with a fragment-header added. RFC 6145, *IP/ICMP Translation Algorithm*, provides a full discussion of the use of the DF flag to control fragmentation.

To clear the DF flag, include the **stateful-nat64 clear-dont-fragment-bit** statement at the **[edit services service-set service-set-name nat-options]** hierarchy level.

Subscriber Access Management



NOTE: Although present in the code, the subscriber management features are not supported in Junos OS Release 12.1R7. Documentation for subscriber management features is included in the Junos OS Release 12.1 documentation set.

- **Additional option for RADIUS NAS-Port attribute (MX Series routers)**—You can configure the width of the aggregated Ethernet identifier field used in the RADIUS NAS-Port attribute (attribute 5). To configure the width, include the **ae-width** option in the **nas-port-extended-format** statement at the **[edit access profile profile-name radius options]** hierarchy level. The **ae-width** field can be from 0 through 32 bits. The total width of the NAS-Port attribute can be a maximum of 32 bits.

[Subscriber Access]

- **Enhanced Subscriber Service Accounting Information**—Subscriber access accounting has been enhanced in Junos OS Release 11.4R2 and later, and Release 12.1R1 and later. RADIUS VSA 26-83 (Service-Session), which is included in RADIUS service accounting start and stop packets, now includes the parameter values used to activate a subscriber service, in addition to the service name. In earlier releases, only the service name was included. When VSA 26-83 is not available from the RADIUS server, subscriber management sends the service name in the accounting message (as was the case in earlier releases).

[Subscriber Access]

- **Keepalive statistics display for PPP fast keepalive (MX Series routers with MPCs/MICs)**—PPP fast keepalive, which is enabled by default on MX Series routers with MPCs/MICs, is the mechanism by which LCP Echo-Request packets are processed by the Packet Forwarding Engine instead of by the Routing Engine. When the router is using PPP fast keepalive for a PPP link, the output of the **show interfaces pp0.logical** operational command does not include the number of keepalive packets received or sent, or the amount of time since the router received or sent the last keepalive packet.

With PPP fast keepalive and **no-keepalives** configured, the **show interfaces pp0.logical** command does not display output statistics for keepalive packets sent, as expected, or input statistics for keepalive packets received from PPP subscribers (clients). Even if **no-keepalives** is configured to prevent the router from sending PPP keepalive messages, the router is still able to respond to PPP keepalive messages that it receives from clients.

When PPP fast keepalive is *not* in use on MX Series routers without MPCs/MICs, or on routers other than MX Series routers, you can view input statistics for received PPP keepalive packets with the **show interfaces pp0.logical** command, even if you have

issued the **no-keepalives** statement to disable sending keepalive messages on the PPP interface.

[*Junos OS Subscriber Access Configuration Guide*]

- **Configuring an AAA local access profile for L2TP clients on the LNS (MX240, MX480, MX960 routers)**—You can configure a local access profile that specifies a particular RADIUS server configuration to override the global access profile and the tunnel group AAA access profile for a LAC client. The AAA access profile for the client takes precedence over the tunnel group profile, which in turn takes precedence over the global access profile for the routing instance.

Include the **aaa-access-profile** statement at the [edit access profile access-profile-name client *client-name*] hierarchy level. In earlier releases, you included this statement only at the [edit services l2tp tunnel-group *name*] hierarchy level to configure a profile for an L2TP tunnel group. The global access profile is configured at the routing instance that hosts the L2TP tunnel with the **profile access-profile-name** statement at the [edit access] hierarchy level.

[*Subscriber Access*]

- **Enhanced monitoring of RADIUS server status and information (MX Series routers)**—RADIUS server monitoring is updated to include a new operational command for displaying RADIUS server information, and an enhanced system log message for RADIUS server events. The **show network-access aaa radius-servers** command enables you to display status and information related to RADIUS servers. The command also supports the **detail** option, which displays additional information. The AUTHD_AUTH_SERVER_STATUS_CHANGE system log (syslog) message is reported at the **warning** Severity level when a RADIUS server state changes from **up** to **down**, or vice versa.

[*Subscriber Access*]

- **Session-ID added to output of show network-access aaa subscribers username command (MX Series routers)**—The output of the **show network-access aaa subscribers username *username*** command includes a column showing the Session-ID for the specified username.

[*Subscriber Access*]

- **Managing CoA requests that include unapplied changes to client profile dynamic variables**—You can manage the way that subscriber management processes CoA requests that include changes to a client profile dynamic variable that cannot be applied. For example: a CoA request might include several changes to client profile dynamic variables, one of which contains updates to a firewall filter that does not exist and so cannot be applied.

In the default behavior, subscriber management does not apply the incorrect firewall filter update but makes the other changes to the client profile dynamic variables, and then responds with an ACK message. Subscriber management supports an optional configuration, which replaces the default behavior. The optional configuration specifies that when a CoA operation is unable to apply a requested change to a client profile dynamic variable, subscriber management does not apply any changes to client profile dynamic variables in the request and then responds with a NACK message.

To configure subscriber management to override the default behavior, use the **coa-dynamic-variable-validation** statement at the **[edit access profile *profile-name* radius options]** hierarchy level.

[Subscriber Access]

- **Enhanced output for show network-access aaa statistics authentication command (MX Series routers)**—The **show network-access aaa statistics authentication** command supports the **detail** option, which displays additional authentication information. The following example shows the output for the enhanced command. The new fields are described in the table after the example.

```
user@host> show network-access aaa statistics authentication detail
Authentication module statistics
  Requests received: 2118
  Multistack requests: 0
  Accepts: 261
  Rejects: 975
    RADIUS authentication failures: 975
      Queue request deleted: 0
      Malformed reply: 0
      No server configured: 0
      Access Profile configuration not found: 0
      Unable to create client record: 0
      Unable to create client request: 0
      Unable to build authentication request: 0
      No server found: 975
      Unable to create handle: 0
      Unable to queue request: 0
      Invalid credentials: 0
      Malformed request: 0
      License unavailable: 0
      Redirect requested: 0
      Internal failure: 0
    Local authentication failures: 0
    LDAP lookup failures: 0
  Challenges: 0
  Requests timed out: 882
```

Table 21: New show network-access aaa statistics authentication Output Fields

Field Name	Field Description
Multistack requests	Number of authentication requests for dual stack subscribers
RADIUS authentication failures	Number of RADIUS authentication requests that have failed
Queue request deleted	Number of queue requests that have been deleted
Malformed reply	Number of malformed replies received from the RADIUS authentication server
No server configured	Number of authentication requests that failed because no authentication server is configured
Access Profile configuration not found	Number of authentication requests that failed because no access profile is configured

Table 21: New show network-access aaa statistics authentication Output Fields (*continued*)

Field Name	Field Description
Unable to create client record	Number of times that the router is unable to create the client record for the authentication request
Unable to create client request	Number of times that the router is unable to create the client request for the authentication request
Unable to build authentication request	Number of times that the router is unable to build the authentication request
No server found	Number of requests to the authentication server that have timed out; the server is then considered to be down
Unable to create handle	Number of authentication requests that have failed because of an internal allocation failure
Unable to queue request	Number of times the router was unable to queue the request to the authentication server
Invalid credentials	Number of times the router did not have proper authorization to access the authentication server
Malformed request	Number of times the router request to the authentication server is malformed.
License unavailable	Number of times the router did not have a license to access the authentication server
Redirect requested	Number of authentication requests that have been redirected based on routing instance
Internal failure	Number of internal failures
Local authentication failures	Number of times local authentication failed
LDAP lookup failures	Number of times the LDAP lookup operation failed

[*Subscriber Access*]

- The **user *username*** option for the **clear services l2tp session** command is no longer available in the CLI for LNS on MX Series routers. Added to the option's previous unavailability for LAC on MX Series routers, this means that L2TP on MX Series routers does not support clearing L2TP sessions based on subscriber username. As an alternative, you can determine the session ID for the username by issuing the **show subscribers detail** command, and then remove the session with the **clear services l2tp session local-session-id *session-id*** command.

[*Subscriber Access*]

- The **user *username*** option for the **show services l2tp session** command is no longer available in the CLI for L2TP LAC or L2TP LNS on MX Series routers. To view L2TP session information organized by subscriber username, you can issue the **show subscribers detail** command or the **show network-access aaa subscribers username** command.

[Subscriber Access]

- L2TP support for SNMP statistics (MX Series routers)**—By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in the table have a default value of zero.

Table 22: SNMP Counters for L2TP Statistics

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	Tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 30,000 sessions, none of these statistics is more than 15 minutes old.



BEST PRACTICE: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

Subscriber Access

User Interface and Configuration

- Enhancement to test configuration operational statement**— The option **syntax-only** allows the user to check a partial configuration without checking for commit errors.

[System Basics and Services Command Reference]

- **Support for displaying logical system and routing instance for L2TP tunnels (MX Series 3D Universal Edge Routers)**—When you issue the `show services l2tp tunnel` command with the `detail` or `extensive` option on either the LAC or LNS, the output displays both the logical system and the routing instance in which the L2TP tunnel is brought up.

[*System Basics and Services Command Reference*]

- **Removal of the sampling action modifier for IPv4 firewall filters**—In the J-Web interface, the **Sample** check box is not available for configuration from the Other Actions section under the Actions tab of the Configure > Security > Filters > IPv4 Firewall Filters page. This configuration of the sample action modifier is not enabled because the `set firewall filter foo term bar then sample` configuration command has been deprecated in the Junos OS CLI in Release 12.1 and later.

[*J-Web Online Help*]

VPNs

- **Route resolution for Layer 2 VPN route in the `bgp.l2vpn.0` route table**—Starting in Junos OS Release 12.1, route resolution for Layer 2 VPN routes in the `bgp.l2vpn.0` route table is based on the routes in the `inet.3` only instead of both the `inet.3` and `inet.0` tables. This behavior makes the route resolution scheme for Layer 2 VPNs and Layer 3 VPNs consistent. No CLI change is introduced as result of this change. However, if a route reflector (RR) is connected to its clients through a route in `inet.0` only and if the RR reflects Layer 2 VPN routes, you must add the following configuration on the RR. This is already required for Layer 3 VPNs in the similar situation.

```
user@RR# show routing-options
resolution {
  rib bgp.l2vpn.0 {
    resolution-ribs [ inet.3 inet.0 ];
  }
}
```

[*VPNs*]

- **Vrf-import policies must reference a target community in the "from" clause**—Starting in Junos OS Release 11.4, vrf-import policies must reference a target community in the "from" clause. If the import policy does not reference a specific community target or if the referenced community is a wildcard, the commit operation fails. As an exception, the policy does not need to reference a community target in the "from" clause when the policy action in the then clause is "reject." Prior to Junos OS Release 11.4, when the vrf-import policy did not reference a specific community target in the "from" clause, the commit operation was successful, but the import policy had a nondeterministic effect.

Changes Planned for Future Releases

The following are changes planned for future releases.

Network Management

- **Change in the Junos OS support for IS-IS MIB**—In Junos OS Release 11.3 and later, the currently supported version of IS-IS MIB, as defined in Internet draft draft-ietf-isis-wg-mib-07.txt, is replaced with the latest version of IS-IS MIB, as defined in RFC 4444. Junos OS can support only one of these versions of IS-IS MIB in a release. Therefore, Junos OS Release 11.2 and earlier continue to support IS-IS MIB as defined in Internet draft draft-ietf-isis-wg-mib-07.txt, whereas Junos OS Release 11.3 and later support only the updated RFC 4444–based version of IS-IS MIB.

[*Network Management, SNMP MIBs and Traps Reference*]

Related Documentation

- [New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 223](#)
- [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 298](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 404](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 432](#)

Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers

The current software release is Release 12.1. For information about obtaining the software packages, see “[Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers](#)” on page 432.

- [Current Software Release on page 298](#)
- [Previous Releases on page 333](#)

Current Software Release

Outstanding Issues

Class of Service (CoS)

- When GRES is configured, a core file might be seen on the backup Routing Engine. This happens if VRRP priority changes continuously and the following configurations are changed on the distribution routers: 1. IRB interfaces 2. Routing and spanning-tree protocols 3. Multicast 4. VPLS routing-instance 5. Bidirectional Forwarding Detection (BFD). The core file does not affect traffic or functionality. [PR668695](#)
- M7i and M10i routers with non-enhanced CFEB support only four forwarding classes. [PR786081](#)

Forwarding and Sampling

- When "burst-size-limit" is configured to a value more than 100,000,000 bytes on M7i or M10i routers with CFEB-E, the correct burst size is not reflected by CFEB-E on the router. [PR706272](#)
- On MX Series routers, in the output of 'show interfaces statistics detail' command, the number of input bytes under the field Transit Statistics randomly modifies and does not increment after that. [PR721445](#)
- Firewall logs do not record rejected or accepted packets on the loopback filter though packets are rejected or accepted. [PR724059](#)
- In Junos OS Release 10.2 and later, FBF does not disable uRPF when setting the uRPF(rpf-check) and FBF on an interface, which is in a routing instance. [PR735697](#)
- When master Routing Engine switchover is performed because the backup Routing Engine is out of synchronization with the master and if GRES(graceful Routing Engine switchover)/NSR(nonstop active routing)/NSB(nonstop bridging) is enabled, L2ald (Layer 2 address learning daemon) might generate a core file. [PR735913](#)
- PFED crashes while processing a corrupted message sent from Packet Forwarding Engine. To fix this issue, changes are being made to PFED. However, the scope of this issue is limited due to the unknown nature of the corruption. A separate PR (771108) is created to make some changes to Packet Forwarding Engine. [PR750143](#)
- For each interface accounting-profile configured on the interface, accounting node is created. On a timer expiration, statistics are fetched for this accounting node. Accounting node pointer is used as the user data in the request. When the response arrives, the pointer is used to write the statistics to output file. However, if accounting node gets deleted before the response arrives, Packet Forwarding Engine daemon (pfed) may crash and generate a core file. [PR782281](#)
- When the set client-idle-timeout is over, VLAN is removed even though there is ingress/egress traffic on PPPoE clients. [PR784529](#)
- In DHCP subscriber management environment, the session database may get into deadlock or recovery state (maybe caused by pfed process restarts), when jdhcpd process tries to access the session database, then it may get into an infinite loop and cause high cpu usage, causing all DHCP packets to be dropped. When issue happens, the following behaviors could be observed: user@router> show dhcp relay statistics error: the dhcp-service subsystem is not responding to management requests And below log messages displays continuously: mgd[32758]: %DAEMON-5-UI_READ_TIMEOUT: Timeout on read of peer 'dhcp-service' mgd[32758]: %DAEMON-3-UI_RECONN_READ_FAILED: Invalid response from peer 'dhcp-service': 421 Connection unexpectedly terminated [PR791816](#)
- In T4000 platforms with ES-FPC, for IPv6 firewall filters with match conditions on address prefixes longer than 64 bits, in some corner cases, the filter may not be correctly evaluated and packet loss may occur. [PR879829](#)
- Sampled deleting from Packet Forwarding Engine whenever LSP path switchover happens in scaled scenario.

- 10 million routes and 800K active routes
- 16 LSPs on active paths
- Chassis level knobs mentioned in the RSI
- 100k IPv6 routes

[PR891899](#)

General Routing

- A feature to support OSI/CLNS type frames on Trinity platform was added in 11.2R1, and this feature is not supported on ichip based DPC. Despite DPC not supporting CLNS type next-hops, the following log message is seen: > \$ zcat messages.log.gz | grep 50556 > Sep 20 02:27:52 dub2-vc-edg-r2 fpc0 Table nexthop for unsupported proto. NH(50556) CLNP:2617 > Sep 20 02:27:52 dub2-vc-edg-r2 fpc1 Table nexthop for unsupported proto. NH(50556) CLNP:2617 > Sep 20 02:27:52 dub2-vc-edg-r2 fpc3 Table nexthop for unsupported proto. NH(50556) CLNP:2617 > Sep 20 02:53:20 dub2-vc-edg-r2 fpc0 NH: Failed to find nh (50556) for deletion > Sep 20 02:53:20 dub2-vc-edg-r2 fpc1 NH: Failed to find nh (50556) for deletion > Sep 20 02:53:20 dub2-vc-edg-r2 fpc3 NH: Failed to find nh (50556) for deletion This fix masks this log message for DPC. [PR680782](#)
- The knob route-memory-enhanced(hierarchy: set chassis) is hidden in platforms M320 and MX Series. There is no functionality break but this knob shouldn't be hidden. It is fixed now. [PR690100](#)
- For an IPv4 pool, only the all-0 host and the all-1 host addresses are precluded from allocation, both for gateway-assigned and external address assignment. [PR729144](#)
- On an MX Series router with GRES enabled, when the system reboots after "dynamic-profile-options versioning" is committed, the DHCP/PPPoE subscriber fails to log in. [PR770140](#)
- When dynamic-profile versioning is enabled, CoA requests are not processed if authd restarts. [PR796416](#)
- When a large number of subscribers log in and log out simultaneously in a scaled configuration, errors might occur when logical interfaces are created. The errors occur because the rpd process fails to read ifstate notifications related to logical interface deletions. As a workaround, restart the rpd. [PR775033](#)
- When XE or GE interface in promiscuous mode moves to an ae bundle, the AE member link is rejecting packets due to bad unicast MAC. [PR783332](#)
- In subscriber management environment, with GRES enabled, subscriber management infrastructure daemon (smid) maintains a directory /mfs/var/sdb which contains several logs/stats and databases that must be cleaned up and compacted every couple of hours. If smid process compacts databases during berkeley database replication daemon (bdbrepd) hasn't finished the initial replication between Routing Engines, smid may get stuck in a loop cause it does not checkpoint or archive the log files anymore, so the /mfs directory eventually fills up. Then the router locks up with no telnet/console access, and no subscribers logging into the router. When issue happens, the following errors could be seen: /kernel: Process (1977,bdbrepd) has exceeded 85%

of RLIMIT_DATA: used 114692 KB Max 131072 KB /kernel: Process (1977,bdbrepd) attempted to exceed RLIMIT_DATA: attempted 131076 KB Max 131072 KB /kernel: pid 1977 (bdbrepd), uid 0 inumber 636230 on /mfs: filesystem full The storage usage of bdbrepd process could be observed by following command (the 'SIZE' field means Total size of the process (text, data, and stack), in kilobytes): user@router> show system processes extensive | match bdbrepd PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND 1568 root 1 96 0 55068K 37940K select 28:32 0.00% bdbrepd [PR796430](#)

- In subscriber management environment with knob versioning enabled for dynamic-profile, during subscriber flapping or login/logout, the authd (authentication daemon) process may leak memory. [PR800028](#)
- Dynamic VLAN might remain stuck in "terminating" state after line card reboot. [PR800533](#)
- In a BRAS environment with PPPoE subscribers and parameterized firewalls and policers, the FPC memory usage will go high after a CoA change for the services' parameters. [PR805922](#)
- BFD packets sent from FPC (distributed mode) over normal physical interfaces are set with ttl 0 so that it gets decremented by 1 and becomes 255 once it is sent out on the wire. This behavior is not the case when the BFD packets are sent over IPsec routed tunnels where the packets are sent from the corresponding service PIC. In this case, the ttl should be set to 255 as no such decrement action takes place when it is sent from a service PIC. But in the current scenario, the ttl is set to 0 as a result of which the service pic drops the outgoing packet. This was an untested scenario till date. [PR808545](#)
- In subscriber management scenario, during DHCP or PPPoE subscribers login, cosd process may leak memory when cosd is initializing cos related parameters for subscribers. It happens only if the subscriber has one COS service session. [PR815777](#)
- If dynamic profile versioning is configured and In-service software upgrade (ISSU) is performed from 11.4x27.35 (GA build) to 11.4x27.38 (Nov-2012), exiting subscribers may either loose traffic or may get terminated. [PR817018](#)
- In subscriber management environment, after scale DHCP/PPPoE subscribers (about 16k) login/logout multiple times, some of them fail to login again due to system can not allocate memory for them. When issue happens, following errors could be seen: [LOG: Err] IFRT: 'IFL demux enable' (opcode 175) failed [LOG: Err] rt_table_index_alloc: failed to allocate index with error: memory allocation failure !!! [LOG: Err] ifdemux_create_table(): ifl 17716:IPv4 => failed to allocate index for demux rt table with error: memory allocation failure !!! [LOG: Err] IFRT: 'IFL demux enable' (opcode 175) failed [LOG: Err] ifl 17716:IPv4; demux create table failed with error=8 [PR819614](#)
- ICMP redirects are not disabled even after configuring no-redirects on irb interface. [PR819722](#)
- In subscriber management environment, with dynamic-profile configured for subscribers, with high churn rate of subscribers, memory leak is observed in authd process. This was observed from a login/logout or flapping of 1000 subscribers every 3mins. [PR835204](#)

- FPC restart can lead to dfwd being terminated. This will require a restart of the openflow daemon for the OpenFlow functionality to work properly again. [PR842923](#)
- Output of "show subscribers physical-interface aex" displays multiple AE links. [PR864555](#)
- In subscriber management environment, with scaling subscribers login (110K DHCP and 20K PPPoE), after restarting one of the line cards which has subscribers, autoconf process might crash and generate a core fledue to memory corruption or memory double free. Only 11.4X27.45 is affected by this issue. [PR870661](#)
- In a scenario with scale Routing Instances (RIs) configured, after deactivating/activating two RIs, the routing protocol processdaemon (rpd) might try to free a specific pointer pointing to an incorrect structure that is actively in use. Then the rpd process crashes and generates core files. [PR870683](#)
- authd reports syntax error, although the syntax is correct, when trying to activate service profile for subscriber and fails to activate the service. [PR883065](#)
- The Routing Engine might become non-responsible due to the exhaustion of kernel mbufs with following messages. /kernel: Mbuf: High Utilization Level: (Low) Throttling low priority requests (10 ms) /kernel: Mbuf: High Utilization Level: (Medium) Throttle low priority requests (150 ms) /kernel: Mbuf: High Utilization Level: (High) Block low priority requests. [PR886083](#)
- When multiple framed-route(type-22) AVPs are present in Radius access accept message, the router will install only the first route into the routing table. [PR891036](#)

High Availability (HA) and Resiliency

- On TX or TXP Line Card Chassis (LCC) with graceful Routing Engine switchover (GRES) enabled, if a mastership switch is being requested on an LCC whose backup Routing Engine's em0 interface has physically failed (due to hardware failure or driver stops working), this will cause all FPCs on the LCC to disconnect from the old master Routing Engine, but cannot reconnect to the new master one either. [PR799628](#)
- RPD on the backup Routing Engine may crash when it receives a malformed message from the master. This can occur at high scale with non-stop routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare and backup RPD will restart and the system will recover without intervention. [PR830057](#)
- With minimal flow configuration, RPD cores during the shutdown process due to missing safety checks. [PR852766](#)

Infrastructure

- IPv6 access and access-internal routes for IP demux interfaces are not installed successfully when the dynamic profile that creates the IP demux interface is configured with the router's unnumbered loopback address. This results in the failure of IPv6 egress traffic forwarding. This problem does not affect IPv4 egress traffic on IP demux interfaces. As a workaround, configure the dynamic profile that creates the IP demux interface to use a numbered address rather than the router's unnumbered loopback address. [PR747029](#)
- Aggregate Bundle interface with IPV6 Interface stuck in Tentative state. Trigger was deactivation/activation of ae-interface. [PR844177](#)
- Kernel may crash when delete routing instance under the donor and unnumbered address borrower scenario. When the deleting for the donor is before the deleting of the corresponding unnumbered borrower, in this window, the donor interface does not have an address, arp processing over the borrower interface during this window may trigger the crash. The core files could be seen by executing CLI command "show system core-dumps". [PR880179](#)
- Every 10 minutes kernel reports "%KERN-6: MTU for 2001:4c0:1:1301:0:1:0:250 reduced to 1500" after reducing MTU once. There is no impact to the system due to this additional log message. [PR888842](#)

Interfaces and Chassis

- For Automatic Protection Switching (APS) on SONET/SDH interfaces, there are no operational mode commands that display the presence of APS mode mismatches. An APS mode mismatch occurs when one side is configured to use bidirectional mode, and the other side is configured to use unidirectional mode. [PR65800](#)
- In a SAToP pseudowire on a 4-port COC3/CSTM1 or 12-port T1/E1/J1 CE PIC, when there is data loss from the pseudowire or due to an alarm condition (LOS/LOF/AIS) at the peer end of the SAToP pseudowire, the local PIC does not transmit AIS. This PR addresses that issue. [PR602563](#)
- Interfaces on 3D 2x10GE XFP MIC or 3D 4x10GE XFP MIC in WAN-PHY mode, when changed from LAN-PHY mode may not detect RDI-P error from remote. [PR700097](#)
- Deactivating the 'default-chap-secret' knob, committing, then activating it again can cause memory corruption. This may trigger an Routing Engine to panic and go to the "db>" prompt. A restart of the Routing Engine is required to recover. [PR718634](#)
- When you have the following configuration on a logical interface, unit 2000 { encapsulation vlan-bridge; vlan-tags outer 40 inner-list [20 3000]; family bridge; } And you execute "show interface intf-name extensive", you will see the following: Under "Flags: SNMP-Traps Redundancy-Device 0x20004000 VLAN-Tag [0x8100.40 0x8100.2000 20,3000] ", you will see the unit number 2000 between outer and inner tags configured. This is just a display issue and no functionality is affected. [PR723188](#)
- The following message is displayed in syslog for every PADI received: pfe_ppp_distributed_keepalive: fail (15) index=453 this PR fixed this. [PR734110](#)

- With a large number of PPPoE subscribers, a queue overflow occurs when a burst of PPP control packets are received. As a consequence, the PPPoE session cannot be established while bringing up subscribers because the system drops control packets during PPP negotiation. Also, when the subscriber sessions are up, the system loses LCP echo replies from the client and over time this might cause keepalive failures. This is caused by the improper PPP packet queue size set in the kernel. [PR735769](#)
- IP header compression (VJ compression) is not rejected by PPP LCP. This is a bug in Junos OS. [PR737981](#)
- Occasionally, an MX Series router that is connected to a customer premises equipment (CPE) using PPP fails to send an LCP configure-ack/IPCP configure-ack for the first LCP configure-request/IPCP configure-request received from the CPE. For LCP, this occurs when the MX Series router receives the LCP configure-request from the CPE soon after sending PADS to the CPE. In the case of IPCP, this occurs when the router receives IPCP configure-request from the CPE while PPP connection is transitioning to the next step for configuring Layer 3 protocols using NCP. [PR773950](#)
- With multiple nondefault routing instance configuration, jpppd seems to bring up the subscribers fine to ACTIVE state but the access-internal routes are not added to all subscribers. It appears as if only the first nondefault RI subscribers might have the routes installed and the rest of them have no routes. [PR775715](#)
- In a DHCPv6oPPPoE environment, DHCPv6 bindings are terminated if the IPv6 address is removed or added on a static IPv6 interface, not related to the PPPoE configuration. [PR780607](#)
- In MLPPP scenario, in rare conditions (such as FPC crash), kernel may try to delete a MLPPP bundle with an invalid (although within the max bundle limit) bundle ID. This will casue vmcore and Routing Engine switchover. [PR780784](#)
- jpppoed memory utilization spikes after GRES or jpppoed restart event. [PR800650](#)
- In PPPoE subscriber management environment, during subscribers logging in, PPP daemon (jpppd) may leak memory due to two software defect issues: 1 - During PPPoE subscribers logging in who use CHAP authentication method, if the CHAP protocol taking longer than 30 seconds to negotiate, PPPoE daemon (pppoed) will try to delete the underlying logical interface (IFL) through which the subscribers log in. Then pppd process try to remove the subscriber entry (which involves ensuring the IFL is correctly deleted). But sometimes jpppd will not remove some of the memory associated with the subscriber and cause jpppd memory leak. 2 - Because software defect in PPP protocol NCP negotiation stage, jpppd process may leak pending buffers during subscribers logging in. The memory leak of jpppd process could be observed by following command (The RES field means "Current amount of resident memory, in kilobytes"):
user@router> show system processes extensive | match "PID | jppp" PID USERNAME
THR PRI NICE SIZE RES STATE TIME WCPU COMMAND 1460 root 3 20 0 66912K
37420K sigwai 20:17 0.00% jpppd [PR802915](#)
- With LSQ interface, the MLPPP fragments cannot use the egress queue 4 to 7 on the MLPPP member links. There is no workaround. [PR805307](#)
- Junos OS doesn't generate vrrpv4 mastership change syslogs while it generates vrrpv6 logs. [PR807217](#)

- DCD reports error when configuring hierarchical-scheduler on MX80 with QX chipset. This is cosmetic error and it should not have functional impact. [PR807345](#)
- Incorrect Detection timestamp in "show chassis fabric reachability". [PR811846](#)
- VC powerdown of VCMm leaves Ip demux interfaces in a "hardware-down" state. [PR813902](#)
- Stale addresses in local address pool. [PR815331](#)
- Planes might go into faulty state during the SCB initialization when the SERDES on the SF chip failed to come up. [PR839509](#)
- ERA events are not credited back by jpppoed. ERA has a purge timer of 10 minutes which reclaims stale events so new connections are allowed after the purge timer fires. In a high scaled scenario this can lead to slow PPPoE connections. [PR842935](#)
- In a scenario of PPP sessions over L2TP tunnels, on L2TP network server (LNS), if authentication is none or if authentication is enabled but radius does not return any Framed-IP-Address/Framed-Pool, jpppd process is not setting the IP address key of subscriber to "255.255.255.254" thereby resulting in address allocation failure in authd process. Then the L2TP tunnels can not be established, hence subscribers can not login. When issue happens, the following logs of authd process could be seen: client type jpppd client type REQUESTING: OldStyle 0 OldStyleFilled 0 hint null network null client pool name. [PR849191](#)
- The device configuration daemon (dcd) may crash when a partial demux subinterface configuration is attempted to be committed. There is no impact to traffic forwarding but before the configuration can be committed, it must provide a valid 'underlying-interface' for the demux subinterface. [PR852162](#)
- M7i Routing Engine Crashed with last reboot reason panic:page fault and kernel core, after commit. [PR868212](#)
- snmpwalk of "jnxPPPoEIfLockoutTable" didnt capture pppoe locked out clients. [PR869024](#)
- Chassisd core generated on initializing process on MX-VC. [PR870457](#)
- MC-LAG will no longer change just the LACP System Identifiers directly, but will also remove the "Synchronization, Collecting, Distributing" bits from the Actor State bits advertised in the PDU. [PR871933](#)
- PPP interface MTU changes after making any configuration changes to the system. [PR897940](#)

J-Web

- On M10i, MX240, and MX480 routers; EX Series switches; and on SRX3400, SRX3600, SRX5600, and SRX5800 devices; when you use the Microsoft Internet Explorer browser to open reports from the following pages in the J-Web interface, the reports open in the same browser session: -Files page (Maintain > Files) -History page (Maintain > Config Management > History) -Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port) -Static Routing page (Monitor > Routing > Route Information) -Support Information page (Maintain > Customer Support > Support Information) -View Events page (Monitor > Events and Alarms > View Events).

[PR433883](#)

- On MX Series routers, the options on the J-Web interface such as Access Concentrator, Idle Timeout, and Service Name for PPPoE logical interfaces are not supported.

[PR493451](#)

- Add IPv4/IPv6 Filters does not work as expected with IE. As a workaround, use IE after configuring the filters and click APPLY on the summary page. As this action does not take you back to the configuration page, navigate back to the IPV4/V6 firewall page by clicking the Security > IPv4/v6 firewall pages. [PR543607](#)

- On the J-Web interface, HTTPS access might work with an invalid certificate. As a workaround, after you change the certificate, issue the "restart web-management" command to restart the J-Web interface. [PR700135](#)

- On the J-Web interface, Configure > Routing > OSPF > Add > Interface Tab is showing only the following three interfaces by default: - pfh-0/0/0.16383 - lo0.0 - lo0.16385 To overcome this issue and to configure the desired interfaces to associated ospf area-range, perform the following operation on the CLI: - set protocols ospf area 10.1.2.5 area-range 12.25.0.0/16 - set protocols ospf area 10.1.2.5 interface fe-0/3/1 [PR814171](#)

- In the J-Web interface on EX Series switches, you cannot configure OSPFv3 by using the point and click function (Configure > Point and Click > Protocols > Configure > Ospf3). As a workaround, configure OSPFv3 options by using the CLI. You can then view and edit the OSPFv3 parameters by using the point and click function in the J-Web interface. [PR857540](#)

Layer 2 Features

- L2VPN/BGP-VPLS site-preference value has changed its encoding format. Value has been traditionally encoded with an erroneous endianness format, which prevented Junos OS platforms from interoperating with other vendors. While the current fix corrects this behavior, it should be taken into account, that updated Junos OS platforms (containing this fix) might be unable to interoperate with those that are running older releases. [PR725197](#)
- In L2vpn scenario, if there is no mpls route to neighbor and there is a static route with discard nexthop in inet.3 table as follows: user@router# show routing-options rib inet.3 { static { route 0.0.0.0/0 discard; } } Then the L2vpn connection will use the above static route in inet.3 table to connect its neighbor as follows: user@router# run show route table mpls mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, * = Both 0 *[MPLS/0] 00:01:42, metric 1 Receive 1 *[MPLS/0] 00:01:42, metric 1 Receive 2 *[MPLS/0] 00:01:42, metric 1 Receive 13

*[MPLS/0] 00:01:42, metric 1 Receive ge-0/0/1.601 *[L2VPN/7] 00:01:38, metric2 0 Discard In this situation, routing protocol daemon (rpd) core dump while walking snmp mib 'jnxVpnPwEntry'. [PR816821](#)

- On TRIO systems after the changes performed within PR/686399 Junos OS 10.4R9 or higher, traffic destined towards mac addresses learned from the core interfaces are aged out every aging interval and added again. During this very short event, vpls traffic will get flooded. [PR820726](#)

Layer 2 Ethernet Services

- DHCP server in a vrf responds with incorrect server identification address option 54 in DHCPOFFER. [PR776222](#)
- In DHCP subscriber management environment, when system deletes a DHCP client while processing login request for another client via the same logical interface, this login request will be dropped and could cause jdhcpd process crash and core dumped. There will have a window where DHCP is recovering the existing subscribers and during this time no new logins will be allowed. However existing subscribers shouldn't be affected. When issue happens, the following logs could be seen: jdhcpd: %USER-4-DH_SVC_ROUTE_ADDITION_FAILURE: DHCP Persistent route addition failure for DHCPv4 client SDB session 51 on incoming interface ae1.0 init: %AUTH-3: dhcp-service (PID 3115) terminated by signal number 11. Core dumped! jdhcpd: %USER-4-DH_SVC_ROUTE_ADDITION_FAILURE: DHCP Persistent route addition failure for DHCPv4 client SDB session 113 on incoming interface ae1.0 [PR782535](#)
- In an MX Series Virtual Chassis environment, a traffic outage longer than 2 minutes occurs when a member of the VC-M LAG fails while LACP is in active mode with link protection on the VC-M router. The outage occurs while the LACP process restarts on the new VC-M router. To avoid this situation, make sure that LACP is running in active link protection mode on the device in front of the VC-M router. This device cannot be an EX Series Ethernet Switch, because the switch does not support LACP in link protection mode. [PR784965](#)
- To free the socket used by jdhcpd for bootp helper deactivate dhcp-service traceoptions. [PR817515](#)
- jdhcpd interface traceoptions are not saved to the default log file jdhcpd and require an explicit file name. [PR823129](#)
- MX wasnt using its link local address to send DHCPv6 replies to IPv6 host. [PR826107](#)
- In a scenario where a DHCP stray request is received on MX Series router acting as the DHCP relay and authentication for this request fails, the MX Series router is generating DHCP NACK back to the DHCP client. [PR835794](#)
- MXVC-DHCP bindings stuck in a "RELEASE(RELAY_STATE_WAIT_AUTH_REQ_RELEASE" state. [PR850187](#)
- In certain cases when an MX Series router is configured as an DHCPv6 server and servicing DHCPv6 clients through LDRA relay, it might send advertisements with UDP port 546 instead of 547. [PR851642](#)

- In DHCP subscriber management environment, while DHCP subscribers login, in rare conditions, system calls of these subscribers fail, due to only on success does system free the memory, resulting in a memory leak for the jdhcpd process. If memory usage of jdhcpd process goes to its limit, no new DHCP subscribers can login. When issue happens, high weighted CPU usage of jdhcpd process and following logs could be observed. /kernel: %KERN-5: Process (31403,jdhcpd) has exceeded 85% of RLIMIT_DATA: used 2825132 KB Max 3145728 KB jdhcpd:
%USER-3-DH_SVC_RTsock_FAILURE: Error with rtsock: rtslib: ERROR Failed to allocate new block of size 16384 jdhcpd: %USER-3-DH_SVC_RTsock_FAILURE: Error with rtsock: rtslib: ERROR Failed to allocate new block of size 16384 jdhcpd:
%USER-3-DH_SVC_RTsock_FAILURE: Error with rtsock: rtslib: ERROR Allocation Failure for (16384) bytes authd[1822]: %DAEMON-3:
../../../../src/junos/usr.sbin/authd/plugin/radius/authd_plugin_radius_module.cc:1090 Failed to get SDB snapshot for session-id:3549005. [PR856024](#)
- In DHCP subscriber management environment, with scaled DHCP subscribers login, after executing "clear dhcp relay binding all/interface" or "clear dhcp server binding all/interface", new subscribers login are delayed, and it shows high CPU usage for a while. [PR857006](#)
- New knob is provided to set the prefix to compare requested ip and server address. Knob is configured as - [edit system services dhcp-local-server] #set requested-ip-network-match <0-31> For V6 [edit system services dhcp-local-server] #set dhcpv6 requested-ip-network-match <0-127> Default will be 8 for v4 and 16 for v6 (first terms). [PR872145](#)
- When IPv6 is configured on integrated routing and bridging (IRB) interfaces that have AE interfaces as child links, after GRES was enabled and after one child link failure or removal, the kernel crashes. [PR878470](#)

Multiprotocol Label Switching (MPLS)

- For point-to-multipoint LSPs configured for VPLS, the "ping mpls" command reports 100 percent packet loss even though the VPLS connection is active. [PR287990](#)
- If you configure the TTL propagation behavior for individual VRFs and VRFs have LDP instances configured, LDP transit routes are not updated when TTL propagation statements are changed. This is also the current behavior for TTL propagation. In order for LDP transit routes to get updated with TTL propagation-related statements, you must run the clear ldp session instance <vrf-instance-name> command. The clear ldp session command disables and then reenables sessions and might cause temporary traffic loss. Clearing the sessions is required only if the propagation configuration changes. You do not need to run the clear command if the no-propagate-ttl, vrf-propagate-ttl, or no-vrf-propagate-ttl statement is in the initial configuration. [PR540175](#)
- Unsupported feature warning missing for mLDP+NSR while doing ISSU. [PR849178](#)
- The LDP protocol might use the lowest IP address configured on an interface even if there is another (higher) address that is explicitly configured as primary. This can lead to unexpected LDP session flap if the lowest but non-primary address is being removed from the configuration. [PR858838](#)

- In an RSVP environment with AutoBw, the Bandwidth Adjustment timer for new LSPs added simultaneously is not smeared along with the rest of the existent LSPs when the smearing algorithm is triggered. [PR874272](#)
- In a scenario where scaled MPLS tag labels exist, while MPLS flapping (which could be triggered by routing protocol flapping), the routing protocol daemon (rpd) might crash and create a core file due to the system trying to delete an already freed MPLS tag label Element. [PR878443](#)
- When BGP labeled-unicast route has BGP label as null and its indirect next-hop requires adding 2 or more labels, traffic using the BGP label may not be forwarded properly. [PR881571](#)
- LSP metric will not be correctly changed as the new configured one after committed when cspf finds an Explicit Route Object (ERO) different from the current ERO and the Path State Block (PSB) re-signaling fails. This is because a change in metric is a local PSB change, but after a configuration change (for example, the bandwidth requirement was changed), PSB and associated routes used to get this change only after a cspf computation followed by a session refresh or re-signaling. If the re-signaling fails, the configured metric value is not updated in the existing PSB and the route metric. [PR894035](#)

Network Management and Monitoring

- SNMP selects the first valid configured IP address on the loopback interfaces for the source address of SNMP Trap. It should select the lowest valid IP address configured on the loopback address. [PR729699](#)
- Any 'show snmp mib walk' cli command, in case of getting in a loop, will introduce SNMPD memory leak, then eventually crash. Note: The same snmp walk from remote NMS won't trigger this issue. [PR732852](#)
- When there are mix of OIDs in a PDU Get request, the subtree information of certain OIDs can be NULL if the respective subagent hasn't registered completely the corresponding MIB objects or if there is some other underlying issue which would cause the subtree information to be NULL. Accessing NULL information caused SNMP to crash. Fix: Put on a safety check for NULL before accessing it. [PR779346](#)
- Junos OS version later than 10.0 reserves separate index pool for private and public interfaces. After an upgrade from version prior to 10.0 to version later than 10.0, some of the public interfaces may be included within the private index pool. There is no operational risk caused by this issue. [PR815028](#)
- When we do snmp polling via the CLI on a big MIB node which has lots of OIDs and huge data, like "show snmp mib walk 1.3.6.1.4.1", that CLI might not be able to consume data at the rate it was being generated by snmpd, so the snmpd buffer is occupied more and more. Eventually this would cause snmpd to reach its rlimit then crash. [PR864704](#)

- SNMP query from valid client on routing-instance-1 with community string that belongs to routing-instance-2 gets the details of routing-instance-2 instead of blocking such queries based on community. [PR865023](#)
- When you perform the following MIB Walk on interfaces, for some interfaces the ifLastChange value will show a value of zero. show snmp mib get ifLastChange.<SNMP ifIndex> will show a value of zero. ifLastChange.<SNMP ifIndex> = 0 [PR886624](#)

Platform and Infrastructure

- RMOPD crash is due to sort of buffer overflow crash and library function being used improperly. It is not caused by RPM scaling, This issue happens randomly and hard to point out the specific trigger. [PR277900](#)
- Following a large output with jcs:get-input() will truncate the output after the ---(more)--- prompt and cause all following user input to be hidden. As a workaround, add the "| no-more" option when executing op scripts. [PR473816](#)
- Commit time warning is changed to trace message. [PR480082](#)
- On the process details page (Monitor > System View > Process Details) of the J-Web interface, there are multiple entries listed for a few processes that do not impact any functionality. [PR661704](#)
- On certain M-Series routers (M20, M40, M40e, and M7i/M10i without Enhanced CFEB), the following error message is displayed on the Packet Forwarding Engine console periodically: "pfe_get_ifl_stats failed" This error is seen only if aggregate interfaces (like AE or AS) is configured on the router. There is no functional impact because of this error message. [PR692081](#)
- The output of the "file list detail" command can display a capital 'S' instead of a lower case 's' for the file or directory permission. The system shell output, however, displays the correct values. [PR736474](#)
- The "request system zeroize" command deletes the /var/db/scripts directory and all subdirectories but does not re-create them. The directories and subdirectories need to be manually re-created via the root shell and the correct permission set. [PR736478](#)
- Junos OS 10.4R8 or higher on MX platforms, L3VPN application using l3vpn-composite-nexthop when the indirect-next-hop configuration statement is added or removed it might cause traffic drops affecting L3VPN flows. To recover from this condition all the l3vpn prefixes needs to get removed and installed new into the forwarding-table, like clearing the bgp peers where the routes are learned from. [PR741646](#)
- With l3vpn composite next-hops configured and 3 or more odd number of core uplinks every l3vpn route deletion will syslog the following error messages. [LOG: Err] JTREE: (jt_mem_free) size 0 for addr 1595452, seg 1, inst 0 [LOG: Emergency] Multiple Free :jt_mem_free There is no operational impact. An even number of core-uplinks will not trigger such error logs. [PR786993](#)
- On TX Series system platforms, under very special corner case condition, non-enhanced FPCs might drop most of the traffic send into the fabric. TXP Platforms are not exposed to this symptom. You will see lots of fabric drops reported via the "show pfe statistics

traffic' or "show class-of-service fabric statistics" command. FPC affected needs to be restarted to recover from this condition. Sometimes you might also see the following error log reported in the syslog. LOG: Err] NFAB(1/1): RODR offset overflow count incremented (65) LOG: Err] CMALARM: Error (code: 542, type:Minor) encountered, cmalarm_passive_alarm_signal LOG: Err] NFAB(1/1): PKTR ICELL signature error counter incremented [PR805682](#)

- When changing configuration repeatedly, in rare conditions, some internal errors may cause CLI process hogs memory and the utilization keeps on increasing due to memory leak. When the memory usage of CLI process increases to around 85% of system limit, the following logs could be seen: /kernel: Process (1383,cli) has exceeded 85% of RLIMIT_DATA: used 62048 KB Max 65536 KB The memory will be released once user logout from the router. [PR813673](#)
- Commit may fail, when a config object is deleted and re-added as transient change from a commit script. [PR814796](#)
- Packet Forwarding Engine may crash when receiving TCP packets with an incorrect format. [PR817318](#)
- In rare case, after no graceful FPC rebooting (i.e. temporary power failure on egress FPC), fabric ASIC on ingress FPC Type with "ES" string can run into temporary problematic status. This will cause a temporary large delay on fabric traffic from Stoli-FPC to the egress FPC. [PR831743](#)
- Cli command 'show route forwarding-table' would only display <= 16 ecmp paths when CBF is used. [PR832999](#)
- An FPC may reboot when a live-core is requested and the /var partition does not have sufficient space to store the live-core. [PR835047](#)
- IPv6 traceroute is not setting traffic class with TOS command. Traceroute packets may not get to the correct queue and the COS bits may not be reflected properly in the traceroute outputs. [PR835359](#)
- If there is a mix of high and low priority fabric traffic as can be seen by checking "show class-of-service fabric statistics", the following error messages can be seen when there are bursts of high priority fabric traffic, while low priority fabric traffic is present :- May 6 14:58:41 routername-re0 fpc1 MQCHIP(0) FI Reorder cell timeout May 6 14:58:41 routername-re0 fpc1 MQCHIP(0) FI Cell underflow at the state stage A second symptom with this mix of low and high priority fabric traffic present; if an FPC that is the recipient of this high and low priority fabric traffic restarts, it is possible for the ingress FPC forwarding ASIC to lockup. In this case the following log message might be simultaneously logged :- Jun 5 13:46:50 router fpc4 MQCHIP(0) CPQ Queue underrun error, Qsys1 Queue 42 Jun 5 13:46:50 router fpc4 MQCHIP(0) CPQ Freecnt nearing empty error, Qsys mask 0x2 [PR877123](#)
- When we are deleting a configuration hierarchy which has no groups applied, the corresponding group object hierarchy is also marked as changed in commit script view. [PR878940](#)
- PHP to PE link with MPLS MTU 1300 allows transit traffic more than 1268, that is up 1272. Note : PE to PHP has default MTU in this case, that is there is MTU mismatch

between the PHP and PE link. Max packet size allowed is 1300 - 20 (ip) - 8 (icmp) -4 (1 label due to PHP) = 1268 [PR879427](#)

- It is observed that in the setup route nexthop for destination of collector's IP address was of type indexed nexthop. [PR889884](#)
- Because of the hardware limit, the feature "maximum-labels" on FPC can't exceed 3. Whenever maximum mpls label is configured as 4 or 5 on unsupported FPC, the LDP/RSVP session will go down and cause MPLS traffic black hole for couple of minutes. This dark window will remain till the unicast next hops are installed and attached to the egress interface where the label has been configure. After that MPLS traffic will resume. [PR890992](#)
- Firewall filter counter doesn't count packets when firewall is configured on discard interface on MX80-T. [PR900203](#)

Routing Policy and Firewall Filters

- The auto-complete feature is not working for the "show policy" command. [PR471332](#)
- If RPF and/or SCU is enabled, then any change to an ingress firewall table filter will trigger RPF/SCU reconfiguration for every prefix in the routing table. This might cause transient high CPU utilization on the fpc which can result in SNMP stats request being timed out. [PR777082](#)

Routing Protocols

- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- When "passive" and "disable" knobs are both configured under [edit protocols isis interface <intf> level <N>] hierarchy the interface is treated as "passive" instead of being disabled. [PR697553](#)
- In some scenarios MVPN-routes with same RD:Prefix may get generated from multiple-VRFs on a PE-router. When such a PE-router is not a MVPN-RR and has no MVPN-EBGP peers, it is possible that the core-network may lose the MVPN-route because of an erroneous MVPN-withdrawal sent by the PE because of the MVPN-route getting deleted from one of the PE VRFs; even if there are other-VRFs on the PE still advertising the route. [PR698493](#)
- BGP family configuration expects that if any more specific configuration--for example neighbor vs. group, group vs. global--is present that no family configuration from less specific configuration is used. When family route-target advertise-default is configured at BGP peer group scope and BGP neighbor scope configuration contained family statements, advertise-default should not be propagated. [PR706925](#)
- Under scaled situations, BGP might core in bgp_rtarget_tsi_update when running family route-target. This may happen when all viable RT-Constrain routes to a destination have been deleted but the route has not been withdrawn from BGP peers. [PR725196](#)

- After upgrade to 10.4R9 following messages are seen "Cancelling deferral pp0 index 131" These messages are not indicative of any problem and only cosmetic. [PR742534](#)
- When a subscriber dynamic profile is configured to add access routes, some route additions might fail intermittently, preventing those subscribers from forwarding traffic. This issue occurs more frequently when the CPU utilization is high. As a workaround, do not specify access-internal routes in the dynamic profile. Omitting this configuration enables the route additions to occur correctly by means of a different internal mechanism. [PR747631](#)
- BFD flaps with AE bundle and 10 ms timer. [PR773101](#)
- With this fix, "jnxBgpM2PrefixesInPrefixesRejected" counter will return the number of prefixes from a BGP peer, that are not eligible to become active. This change makes the variable conform to definition in the specification <http://tools.ietf.org/html/draft-ietf-idr-bgp4-mibv2-03>. There is a new variable "jnxBgpM2PrefixInPrefixesActive" introduced, to return the number of active prefixes from a BGP peer. So the new sequence of variables for the table is as follows:

```
root@wfpro-mx4-c> show snmp mib walk jnxBgpM2PrefixCountersTable
jnxBgpM2PrefixCountersAfi.0.1.1 = 1 jnxBgpM2PrefixCountersSafi.0.1.1 = 1
jnxBgpM2PrefixInPrefixes.0.1.1 = 0 jnxBgpM2PrefixInPrefixesAccepted.0.1.1 = 0
jnxBgpM2PrefixInPrefixesRejected.0.1.1 = 0 jnxBgpM2PrefixOutPrefixes.0.1.1 = 3
jnxBgpM2PrefixInPrefixesActive.0.1.1 = 0
```

[PR778189](#)
- This issue reported captures a change in behavior observed from previous releases. The adjacency hold down is taking longer than expected on passive interfaces and subsequently the issue disappears. This will not cause any functionality break since the functionality is restored eventually and seen only on passive interfaces immediately after ISSU. [PR780684](#)
- In L3VPN scenario, if PE-CE's link is multi-access LAN, the direct subnet route on LAN PE-CE interface will be advertised with a matching nexthop label. In case there are multiple matching nexthops, one of the nexthop labels is selected randomly for the direct subnet route. If the chosen nexthop is unreachable, L3VPN customer's traffic destined to the direct subnet will be dropped. This is a day one issue. [PR781685](#)
- Problem: Routes not deleted from routing table when interface is deleted. Analysis: SPF calculation was not triggered in one particular code flow after the LSAs are deleted from database. Solution: SPF calculation is triggered when the LSA is being deleted due to zero links. [PR782029](#)
- In some specific cases spf calculation may be incomplete because of the specific order the IS-IS LSP fragments are received. [PR797278](#)
- Setting OSPF overload via the configuration sets both the metric field in router LSAs as well as te-metric field in opaque LSAs to 65535 or $2^{16}-1$. Since te-metric is a 32-bit field, it should be set to $2^{32}-1$. [PR797293](#)
- On a device that is running Protocol Independent Multicast (PIM) and with nonstop active routing (NSR) enabled on the device, if a PIM corresponding interface flaps continuously, a PIM thread might attempt to free a pointer that has already been freed, causing the routing protocol daemon (rpd) to crash and create a core file. [PR801104](#)

- With OSPFv3, PIMv6, or LDP configured, the periodic packet management daemon (ppmd) takes responsibility for these protocols' adjacencies. In a rare condition, kernel might send an invalid packet with a null destination in the message header to ppm process, causing the ppm process to crash and create a core file. [PR802231](#)
- AUTOEXPORTED secondary BGP routes that are advertised using "advertise-inactive" can miss the flash event; if so, this leaves the deleted secondary route in the stuck state. [PR818552](#)
- OSPF route will not be deleted from routing/forwarding table if configuration satisfies below simultaneously. 1. Router ID is not specified and it can be changed due to interface down. 2. There is an interface where OSPF is not running. Suppose OSPF is running on interface A and it is not running on interface B. IP address of interface A is selected as router ID. When interface A goes down and router ID is changed to the IP address of interface B, OSPF on interface A will lose adjacency to the remote OSPF router but router will keep routes learnt via OSPF. [PR820909](#)
- Multiple route nexthops will not be returned via SNMP for ipCidrRouteTable object. [PR831553](#)
- In BGP scenario, the initial peer flaps and goes down. Then a new peer is established which causes the rpd process to crash. [PR840652](#)
- IS-IS Prefix Export limit and NSR switchover might push routers into overload. [PR853328](#)
- When an import-policy change rejects a BGP-route previously contributing to BGP-Multipath formation, the Peer Active-route-counters in "show bgp neighbor" may not get updated correctly. [PR855857](#)
- Multicast packets coming with source address as 0.0.0.0 might cause the rpd to crash. [PR866800](#)
- If the SNMP MIB for BGP is walked, the AFI=1, SAFI=5 entries are missing. If an SNMP "get" is performed, the values can be retrieved. [PR868424](#)
- In VPLS multi-homing environment, with same route-distinguisher configured for the VPLS primary PE and backup PE, routing protocol daemon (rpd) may crash and dump a core file in each of following two scenarios: 1 - On VPLS backup PE, enable "advertise-external" knob, then rpd process crashes and dumps a core file on backup PE. 2 - On VPLS primary PE, enable "advertise-external" knob, after disabling the VPLS interface, rpd process crashes and dumps a core file on primary PE. When issue happens, the following behavior could be observed: user@router> show bgp neighbor error: the routing subsystem is not running user@router> show vpls connections error: the routing subsystem is not running [PR869013](#)
- If a static route is configured and exported into OSPF, and it has the same subnet as an OSPF interface address, then committing configuration change (even unrelated to OSPF, like device's hostname) may cause the static route related OSPF type-5 LSA to be removed from the OSPF database. [PR875481](#)
- The issue will happen under following conditions: * NSR is enabled; * The primary Routing Engine will send BGP route with two or more of aggregator attribute, cluster attribute, or an unknown attribute; * The backup Routing Engine may store the BGP route with a incorrectly data structure. Trigger: * Perform switchover. Because of the

incorrect format route stored in backup Routing Engine, the RPD will crash and generate a core. [PR885305](#)

Services Applications

- When you specify a standard application at the [edit security idp idp-policy <policy-name> rulebase-ips rule <rule-name> match application] hierarchy level, IDP does not detect the attack on the nonstandard port (for example, junos:ftp on port 85). Whether it is a custom or predefined application, the application name does not matter. IDP simply looks at the protocol and port from the application definition. Only when traffic matches the protocol and port does IDP try to match or detect against the associated attack. [PR477748](#)
- When the unit 0 for the multi-services PIC interface is not specified, the "monitor interface traffic" command doesn't display input packets number properly for that particular ms-I/F. e.g) user@lab# show interfaces ms-1/2/0 unit 1 { family inet; } [PR544318](#)
- Default Junos OS configuration contains an application definition for matching UDP based traceroute. It does not include port 33434 which is also used by unix traceroute implementation. [PR727825](#)
- Extensive CLI requests associated with l2tp (show services l2tp <switch>) may result in l2tpd process crash. [PR755948](#)
- The show services l2tp session 'interface' extension did not work prior to this PR at devices acting as L2TP LACs. With the implementation committed in this fix, 'interface' can be used as further discriminators to identify l2tp session details in a LAC. [PR780651](#)
- In an L2TP setup with MX Series router acting as an LAC, if the value used passed from RADIUS in VSA "Tunnel-Client-Endpoint" does not exist on the router, Junos OS will send SCCRQ message to LNS with random source addresses. [PR788081](#)
- When an MX Series router configured as an LNS sends an Access-Request message to RADIUS for an LNS subscriber, the LNS now includes the Called-Station-ID-Attribute when it receives AVP 21 in the ICRQ message from the LAC. [PR790035](#)
- If you set aggregated Ethernet interface ipfix sampling, IPv6 egress flow samples get snmp index of member link. Flow samples of IPv4 ingress, egress and IPv6 ingress do not experience this problem. They get snmp index of the aggregated Ethernet interface. [PR791619](#)
- MX CLI allows you to configure "*" for client name in the l2tp access profile leading to failure of establishing the l2tp connection. [PR799232](#)
- In scenario where MX is acting as LAC and radius server is returning tunnel-server-endpoint attribute but NOT returning tunnel-client-endpoint memory leak in jl2tpd process can occur. Additionally same memory leak can occur if unnumbered loopback attribute is returned from radius for tunneled subscribers. [PR800107](#)
- In L2TP subscriber management environment, on LAC, if Calling-Station-ID AVP (AVP Code 31) returned from Radius server is longer than 64 bytes, while LAC sending a call setup message which contains calling-number AVP (AVP Code 22) to LNS, since LAC

only sends upto 64 bytes in the calling-number AVP, the jl2tpd process will crash and dump core. During the restart of jl2tpd process, no new L2TP tunnels will be allowed to establish. If jl2tpd process crashes several times, it wouldn't recover. Existing L2TP sessions aren't affected. The core files could be seen by executing CLI command "show system core-dumps". [PR802044](#)

- Memory leak in key management daemon (kmd) causes some IPSec VPN tunnels to be dropped and don't get re-negotiated for over 10 minutes. Before issue happens, the following logs could be observed: /kernel: Process (1466,kmd) attempted to exceed RLIMIT_DATA: attempted 131080 KB Max 131072 KB /kernel: Process (1466,kmd) has exceeded 85% of RLIMIT_DATA: used 132008 KB Max 131072 KB [PR814156](#)
- Called-station-id always sent by LNS irrespective of configuration on LNS. The knob of "excluding" this attribute didnt work. The fix was tested working in 11.4X27.38 [PR818899](#)
- An LNS configuration on MX Series boxes can be reported as invalid even if it has been successfully committed previously. [PR823709](#)
- In L2TP subscriber management environment, on L2TP Access Concentrator (LAC), L2TP tunnel idle timer is started when the last session on the tunnel is deleted. If the tunnel idle timer expires, then L2TP keeps the tunnels/session/destinations in dying state for the duration of destruct timer (which by default is 5 minutes (300 secs)) before they get destructed. During this phase, jl2tpd process tries to resurrect the tunnel in dying state, causing jl2tpd process crash and dump core. When issue happens, the following logs could be observed: init: l2tp-universal-edge (PID 50230) terminated by signal number 6. Core dumped! /kernel: pid 50230 (jl2tpd), uid 0: exited on signal 6 (core dumped) The impact of l2tpd process crash is, for short period of time tunneled subscribers cannot connect while processes restarts, existing connections are not expected to drop. The unexpected result of continued crashes (which has been found in production and been replicated in the lab) is some subscribers are left in stale states, subscriber disconnects and reconnects but original session gets stuck on LAC in stale state. This will cause memory jump of many different processes (e.g. authd, jpppd, dcd, dfwd, rpd, cosd). [PR824760](#)
- When rollback from v9 to v5 is done, sampling logic was not rolling back, as sampling registers are not getting released from Packet Forwarding Engine and because in v5 the sampling is Routing Engine-based, it was not working. [PR824769](#)
- The lawful intercept message is cosmetic and does not impact functionality. [PR830457](#)
- In L2TP subscriber management environment, after issuing CLI command "commit full", jl2tpd process (l2tp daemon) deletes all tunnel profiles and brings down all L2TP subscribers. Even though there are no configuration changes. [PR834504](#)
- When DHCP subscribers login and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: GENCFG: op 24 (Lawful Intercept) failed; err 5 (Invalid)." [PR837877](#)
- If flow-tap or radius-flow-tap is configured and logging, dynamic flow control daemon (dfcd) may be leaking file descriptors. Over time these leaked file descriptors reach the limit and following error message will be seen. /kernel: kern.maxfiles limit exceeded

by uid 0, please see tuning(7). Then routing protocol daemon (rpd) may crash and generate a core file. [PR842124](#)

- When DHCP subscribers log in and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: rts_gencfg_dependency_ifstate(): dependency type (2) is not supported." [PR864444](#)
- If RSP1 and RSP10 interfaces are configured on the same box issuing the "request interface switchover rs1" or "request interface revert rsp1" causes both RSP1 and RSP10 to switchover or revert. [PR877569](#)
- The jl2tpd process generates a core file as follows:
"././src/bsd/lib/libc/stdlib/abort.c:69." [PR887662](#)
- SIP ALG - Service PIC might crash when SIP flows are cleared. [PR890193](#)
- Output interface' shown as 'Unknown' under show services accounting flow-detail.issue has been analyzed RCA;-At the time when a flow is created in PIC memory, if the route to the destination IP (in the flow) is not known, we set a flag indicating that there is no route to Destination IP in the flow structure. When the flows are queried using "show service accounting flow-detail", the picinfo daemon inspects this flag for each flow and prints the Output interface as "Unknown" if this flag is set. Now, after route record for that flow is downloaded to the Service PIC, the flow structure is updated to reflect the corresponding output interface, but, the above flag is NOT UNSET. So, picinfo daemon continues to print the output interface as "unknown" whenever "show services accounting flow-detail" is executed. [PR890324](#)

Subscriber Access Management

- A Change-of-Authorization request is being NAK-ed on MX80 when a PPP subscriber is terminated in a non-default routing-instance. It works as expected if the subscriber is terminated in the default routing-instance. [PR704560](#)
- "client-idle-timeout" under access profile stanza does not work when RADIUS accounting is not configured. [PR717870](#)
- After a large number of concurrent PPP session logouts and GRES operation some sessions may not complete logout (services activated from SRC). Sessions eventually will time out and clear. [PR742900](#)
- If Juniper Networks Session and Resource Control (JSRC) is used to manage subscribers policies, when subscribers log out, AAA messages may not be freed by the generic authentication service process (authd). So this may result in a memory leak in authd daemon. When the max memory is reached, new subscribers will not be able to log in. [PR753101](#)
- Subscriber management supports three methods for assigning addresses to DHCP clients. When multiple methods are configured, the router uses the following precedence to determine which address to assign to the client. 1. Address defined on the RADIUS server by Internet Assigned Numbers Authority (IANA) vendor ID 4874 attributes 26-4 (Primary-DNS) and 26-5 (Secondary-DNS). 2. Address defined on the RADIUS server by IANA vendor ID 2636 attributes 26-31 (Primary-DNS) and 26-33 (Secondary-DNS). 3. Address defined in the local address pool on the router. [PR772408](#)

- When adding new static DHCP binding where the newly added host address are already active in the dynamic DHCP pool, system will continue to use the dynamic mapping until it is available. If that subscriber release and reinitiate the DHCP request, the new static mapping would not take effect. [PR777257](#)
- In scenario where multiple address ranges are used in address assignment pool for pppoe subscribers. Problem is exhibited after 1st range address space is exhausted. The MX BNG accepts new subscribers at very low rate (few per minute) while "authd" process is running ~90% or RE CPU [PR778179](#)
- Authd attempts to remove VLAN when subscribers are idle but connected. [PR789009](#)
- The captive portal content delivery service applied on PPPoE subscriber to rewrite IPDA is not working. The subscriber traffic is altered but is dropped on MSDPC. [PR789368](#)
- In subscriber management scenario, if the subscribers (such as PPPoE, DHCP, IPOE) login through authenticated dynamic VLAN, after them logout, the authenticated dynamic VLAN will be removed, but the corresponding libstats iflstats entry will be left which should not be. Finally the memory leak will cause filesystem /mfs full and prevent subscribers from login. If issue happens, the following logs could be seen: /kernel: ifl(pp0.1073859148): ifl_config not found !!! smid: /mfs capacity 83% smid: /mfs capacity 85% authd[13198]: ===== Idle Timeout Exceeded Rid the subscriber ===== last message repeated 2409 times smid: /mfs capacity 108% /kernel: pid 13197 (jdhcpd), uid 0 inumber 24280 on /mfs: filesystem full jdhcpd: DH_SVC_LOGIN_FAILURE: DHCP pre-authentication failure for DHCPv4 client SDB session 8659554 on incoming interface demux0.1073841747. [PR796299](#)
- MX-VC:Authd keeps retrying the attempts to fetch final stats for already removed IFL. [PR806104](#)
- MX-VC:Authd sends duplicate requests to enable interim accounting in PFED for idle timeout configured on VLAN subscriber. [PR806112](#)
- MX-VC:Authd core:../../../../authd_aaa_dyn_req.cc:780 [PR806128](#)
- In a network with Session and Resource Control (SRC), after changing accounting-interim-interval on SRC, accounting start and stop messages for service sessions are generated properly, but interim accounting messages for these service sessions are not generated by MX Series BRAS. This may cause states for these service sessions expiring hence cause service accounting working incorrectly. When issue happens, the following messages could be observed: user@router> show network-access aaa subscribers session-id 39 detail Type: dhcp Stripped username: 7278010887 AAA Logical system/Routing instance: default:default Target Logical system/Routing instance: default:default Access-profile: ACCESS-PROFILE-GLOBAL Session ID: 39 Accounting Session ID: 39 Multi Accounting Session ID: 0 IP Address: 10.30.196.9 Authentication State: AuthStateActive Accounting State: Acc-Interim-Sent Provisioning Type: Jsrc Service name: 1410921251356016676 Service State: SvcActive Session ID: 40 Session uptime: 00:22:30 <-----22 minutes passed Accounting status: on/volume+time Service accounting session ID: 39:40-1345620109 Service accounting state: Acc-Start-Sent <-----Accounting interim is not sent Accounting interim interval: 900 Service accounting protocol: JSRC [PR806784](#)

- Berkeley DB is used by various daemons (pfed, dcd, cosd, smid). This database interaction has been found to leak file descriptors in various scenarios. This leak can lead to error messages and possibly a crash of one of these daemons or the kernel. Check the system limit: % `sysctl -a | grep kern.maxfiles ..` and current usage: % `fstat -vm | wc` [PR809189](#)
- In subscriber management environment, while receiving a Change of Authorization (COA) packet with session ID that is ID of service session and the COA packet doesn't contain username, the lookup for the session id finds the service session rather than the subscriber session which isn't handled gracefully, then causing authd process crash and core dumped. The core files could be seen by executing CLI command "show system core-dumps". When issue happens, the following behaviors could be observed: `user@router> show network-access aaa statistics authentication detail error: the general-authentication-service subsystem is not running` `user@router> show log messages %AUTH-3: general-authentication-service (PID 14502) terminated by signal number 11. Core dumped!` [PR811607](#)
- In a subscriber management environment, with Juniper Session and Resource Control (JSRC) configured for subscriber access, if the subscriber sessions with services activated from JSRC flap, authd process might crash and create a core file. [PR816036](#)
- In DHCP/PPPoE subscriber management environment, after terminating subscribers, authd process might crash and generate a core file because an invalid pointer is used. [PR821639](#)
- In situation when CoA message includes both LI attributes and CoA attributes authd process fails to respond to CoA. [PR821876](#)
- Subnet mask option is not returned to DHCP client when framed-ip-address is used with dhcp-local-server. [PR851589](#)
- Authd core experienced when multiple DHCP subscriber connection attempts require SRC for subscriber authentication. [PR862037](#)
- PPPoE subscribers do not always get disconnected after the client-session-timeout expires. [PR869559](#)
- DTCP - First 127 triggers are applied. [PR873013](#)
- The authdlib logout/terminate release notify request might experience a processing loop. [PR888281](#)

User Interface and Configuration

- The logical router administrator can modify and delete master administrator-only configurations by performing local operations such as issuing the load override, load replace, and load update commands. [PR238991](#)
- Selecting the Monitor port for any port in the Chassis Viewer page takes the user to the common Port Monitoring page instead of the corresponding Monitoring page of the selected port. [PR446890](#)
- User needs to wait until the page is completely loaded before navigating away from the current page. [PR567756](#)

- The J-Web interface allows the creation of duplicate term names in the Configure > Security > Filters > IPV4 Firewall Filters page. But the duplicate entry is not shown in the grid. There is no functionality impact on the J-Web interface. [PR574525](#)
- Using the IE7 browser, while deleting a user from the Configure > System Properties > User Management > Users page on the J-Web interface, the system is not showing warning message, whereas in the Firefox browser error messages are shown. [PR595932](#)
- If you access the J-Web interface using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR603669](#)
- The output of the "show system users no-resolve" command displays the resolved hostname. [PR672599](#)
- On the J-Web interface, next hop column in Monitor > Routing > Route Information displays only the interface address and the corresponding IP address is missing. The title of the first column displays "static route address" instead of "Destination Address." [PR684552](#)
- Protected sections of the group hierarchy do not have their protection status displayed correctly and are not prevented from adding new elements into existing groups. [PR717527](#)
- "annotate" was not valid under firewall filter then hierarchy level and displayed "No valid completions", and lead to the configuration could not be committed under "edit private" mode. [edit] liutao@mx480-a-re1# show | compare [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] + /* Don't process the packet here; it's IPv6, not IPv4. + * Accept it and have it be processed by the IPv6 ACL. */ accept; syntax error. liutao@mx480-a-re1# commit full [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] 'accept' outgoing comment does not match patch: [PR812111](#)
- In the cases of aggressive provisioning using scripts or automated tools, we recommend not to using rollback immediately after a successful commit. [PR874677](#)

VPNs

- When you modify the frame-relay-tcc statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level of a Layer 2 VPN, the connection for the second logical interface might not come up. As a workaround, restart the chassis process (chassisd) or reboot the router. [PR32763](#)
- BGP community 0xFF04 (65284) is a well known community (NOPEER), but it is incorrectly displayed as "mvpn-mcast-rpt" in the cli command "show route". This is a show command issue only. No operational mis-behavior will be observed on the router/network. [PR479156](#)
- VPLS traffic gets flooded back over the ingress interface on the local PE as the split-horizon gets disabled upon interface flap. [PR818926](#)

- RPD can crash when a cmcast leave is received after disabling the internet-multicast. [PR864304](#)
- If a logical interface is taken out of VPLS or L2VPN Pseudowire Routing Instance and placed in protocol l2circuit, after the above configuration changes are done in one commit, routing protocol process (rpd) crashes and generates a core file. [PR872631](#)

Resolved Issues

Class of Service (CoS)

- A few memory leaks have been fixed in the class of service daemon. [PR811613: This issue has been resolved.](#)
- When 'scheduler-map-chassis derived' configuration is used under class-of-service, interface related configuration changes can lead to cosd process crash. [PR863734: This issue has been resolved.](#)

Forwarding and Sampling

- Memory leak could happen to pfed, dcd, cosd, cfmd and dfcd processes if user frequently and repeatedly executes "show interface extensive" command from multiple telnet sessions under the conditions below. 1. Set screen-length value to small value. Screen length can be changed by the command "set cli screen-length <n>". 2. User enters "show interface extensive" command simultaneously from multiple telnet sessions. And cancel the output of the command with "q" as soon as "---(more)---" shows up at the end of the output. [PR843145: This issue has been resolved.](#)
- Possibility of duplicate packets when sampling and interface-style nat are configured [PR861984: This issue has been resolved.](#)
- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 or MX Series DPC, if there is distributed Bidirectional Forwarding Detection (BFD) running on Aggregated interface and a firewall filter is configured on loopback interface (lo0), the lo0 will bind an implicit filter, after FPC restarts or Routing Engine switchover, the next hop of the implicit filter is not updated with the corresponding link word to point to CLI filter, causing the CLI filter to be not executed. To resolve the issue, deactivate the firewall filter under loopback interface and then activate it again. Note: The default operational mode of bfd for all protocols is distributed mode (runs on Packet Forwarding Engine), one exception being ospf v3 which runs on Routing Engine by default (centralized mode).+ So ospf v3 is not affected by this issue. [PR864665: This issue has been resolved.](#)
- Outbound control traffic is not counted by accounting-profile which applied to logical interfaces of AE (Aggregated Ethernet). This is a variation of the PR-562964. [PR866181: This issue has been resolved.](#)
- lab@T1600-2_Critical_VZB_Manjit> show services accounting flow-detail
destination-prefix 20.1.1.2/32 Service Accounting interface: sp-2/0/0, Local interface
index: 147 Service name: (default sampling) Interface state: Accounting Protocol Input
Source Source Output Destination Destination Packet Byte Time since last Packet
count for Byte count for interface address port interface address port count count
active timeout last active timeout last active timeout udp(17) xe-0/0/3.0 10.1.1.2
whois++(63) xe-0/0/2.0 20.1.1.2 whois++(63) 1075917 49492182 00:17:55 1780922

81922412 tcp(6) xe-0/0/3.0 10.1.1.2 0 xe-0/0/2.0 20.1.1.2 0 106479 4898034 00:01:46
1835070 84413220 [PR881629: This issue has been resolved.](#)

General Routing

- It is possible that RPD's higher priority tasks (HPTs) are scheduled to run such that lower priority tasks (LPTs) may not be able to complete until HPTs are completed.
[PR836197: This issue has been resolved.](#)
- When MX Series router running with DPC is upgraded by ISSU, some of interface may show incorrect input packet/byte count. And the incorrect count is also seen to the related interface MIB. The value will be a large number. Physical interface: xe-3/1/0, Enabled, Physical link is Up Interface index: 138, SNMP ifIndex: 5449, Generation: 141 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback: Local, Source filtering: Disabled, Flow control: Enabled Device flags : Present Running Loop-Detected Interface flags: SNMP-Traps Internal: 0x4000 Link flags : None CoS queues : 8 supported, 8 maximum usable queues Hold-times : Up 0 ms, Down 0 ms Current address: 00:24:dc:9c:7c:30, Hardware address: 00:24:dc:9c:7c:30 Last flapped : 2013-01-13 14:36:25 JST (02:07:52 ago) Statistics last cleared: Never Traffic statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps ~snip~ Logical interface xe-3/1/0.0 (Index 196614) (SNMP ifIndex 5450) (Generation 140) Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Traffic statistics: Input bytes : 3867797326912475 Output bytes: 0 Input packets: 15108583308733 Output packets: 0 Local statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Transit statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps Protocol inet, MTU: 1500, Generation: 160, Route table: 0 Flags: Sendbcst-pkt-to-re Addresses, Flags: Is-Preferred Is-Primary Destination: 10.3.1/24, Local: 10.3.1.1, Broadcast: 10.3.1.255, Generation: 141 Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0 Policer: Input: __default_arp_policer__ gladiolus:Desktop\$ grep .5449 mib_value_after_issu.txt ifName.5449 = xe-3/1/0 ifInMulticastPkts.5449 = 0 ifInBroadcastPkts.5449 = 0 ifOutMulticastPkts.5449 = 0 ifOutBroadcastPkts.5449 = 0 ifHCInOctets.5449 = 3867797326912475 ifHCInUcastPkts.5449 = 0 ifHCInMulticastPkts.5449 = 0 ifHCInBroadcastPkts.5449 = 0 ifHCOctets.5449 = 0 ifHCOUcastPkts.5449 = 0 ifHCOUmulticastPkts.5449 = 0 ifHCOUbroadcastPkts.5449 = 0 gladiolus:Desktop\$ grep .5450 mib_value_after_issu.txt ifName.5450 = xe-3/1/0.0 ifInMulticastPkts.5450 = 0 ifInBroadcastPkts.5450 = 0 ifOutMulticastPkts.5450 = 0 ifOutBroadcastPkts.5450 = 0 ifHCInOctets.5450 = 3867797326912475 ifHCInUcastPkts.5450 = 15108583308733 ifHCInMulticastPkts.5450 = 0 ifHCInBroadcastPkts.5450 = 0 ifHCOctets.5450 = 0 ifHCOUcastPkts.5450 = 0 ifHCOUmulticastPkts.5450 = 0 ifHCOUbroadcastPkts.5450 = 0 [PR847106: This issue has been resolved.](#)
- It is possible for RPD core when the following conditions are met: - VRF with multipath knob configured - static routes with next-hops which are indirect type and needs further resolution - the numerically lowest (smallest IP) next-hop of indirect type becomes unreachable RPD core is NOT triggered in either of the following scenarios: - no multipath under VRF - if there is no static route entry - static route whose next-hops are indirect type requiring further resolution multipath under VRF is supported only for

BGP configurations. multipath in other conditions are not supported, and a bug in this detection phase is fixed in this PR. [PR847214: This issue has been resolved.](#)

- In certain Graceful Routing Engine Switchover (GRES) scenarios, with IPv6 address configured on at least two interfaces, Solicited node multicast addresses (SNMA) and link local addresses with same prefix might be created on the two interfaces. There is a possibility that there could be inconsistency in the Next Hop database between Master and Backup Routing Engines. When the Backup becomes Master in these scenarios, it'll try to program the Packet Forwarding Engines with the bad Next Hop data. This may cause undesired forwarding behavior on the Packet Forwarding Engines. [PR850625: This issue has been resolved.](#)
- Routing engine may kernel panic crash after SW upgrade to Junos OS including the fix of PR831233. [PR851086: This issue has been resolved.](#)
- Ptsd failed to append policy with multi-rules since 'msg over size limit' [PR852224: This issue has been resolved.](#)
- In a virtual chassis environment in the event power is loss on the Master virtual chassis the standby chassis has potential to experience slot resets during transition period. [PR859717: This issue has been resolved.](#)
- Multiple clksyncd core-dump may be seen after ISSU upgrade on the chassis. [PR861676: This issue has been resolved.](#)
- When a prefix next-hop address resolution requires a recursive lookup, the next-hop might not be updated correctly after an egress interface is disabled. [PR862989: This issue has been resolved.](#)
- When using BGP Flow Spec with rate-limit option, even though the value is in Bytes/second, the value being programmed is in bits/second. [PR864496: This issue has been resolved.](#)
- Configuration of Container Interfaces for APS on MX FPCs is not allowed since Junos OS 12.1. If this feature is needed on MX Series legacy FPCs use a release with this PR fixed [PR869192: This issue has been resolved.](#)
- When configuration stanza: [protocols router-advertisement] starts as: ## ## inactive: protocols router-advertisement ## interface ge-0/0/1.1 { virtual-router-only; } Then perform the following actions: Step 1 - activate protocols router-advertisement Step 2 - deactivate protocols router-advertisement interface ge-0/0/1.1 Step 3 - set protocols router-advertisement interface ge-0/0/1.2 After issuing "commit check", there are no problems. But after issuing "commit", routing protocol daemon (rpd) crashes and dumps core with following logs: rpd[1422]: RPD_RA_CFG_UNKNOWN_ACTION: Unknown configuration action 3 received. [PR871359: This issue has been resolved.](#)
- Under high scale, expiry of a kernel side reconnect timer would cause it to send a non-serviceable message to the Packet Forwarding Engine (asking the line cards to restart and resync since reconnect failed). Since there is no ack- to this kernel message, kernel thought it sent the message and untoggles the GRES flag. The Packet Forwarding Engine wasn't expecting anything so it continued along. The EFFECT: The system is permanently not ready for GRES... CLI GRES check will always report: [cmd] request chassis routing-engine master switch check Apr 14 19:03:13 [INFO] warning: Standby

Routing Engine is not ready for graceful switchover. [PR873679: This issue has been resolved.](#)

- The default setting for the sysctl "net.pfe.relayg_merge_enabled" is 0 (off), this results in a support limit of 16 line-cards within the VC. Even with the group merge disabled, line-cards may have been grouped at system start-up only presenting an issue after they restart. [PR874791: This issue has been resolved.](#)

Infrastructure

- The root cause of the problem was IFADDR change in VRRP context was not replicated to GRESS backup. [PR790485: This issue has been resolved.](#)
- In a L3VPN environment, when PE enabled with composite-nexthop receives a ICMP packet from local CE having TTL=1 and IP-Option set packets destined towards remote VPN end, then below messages are seen in syslog of PE router. /kernel: %KERN-3: tag_send_nh_chain(): no mbuf after tag_nh_chain_comp_label_output() This syslog message doesn't indicate real mbuf issue. Hence this can be treated as non-intrusive. [PR811406: This issue has been resolved.](#)
- Kernel fails to generate ICMP ttl expired when IP packet len is a multiple of 256 [PR829567: This issue has been resolved.](#)
- Delay in bringing ONLINE an FPC after it is inserted into the chassis [PR853304: This issue has been resolved.](#)
- In a scenario with scaling routes existing (e.g. 54k BGP routes), while these routes flapping, in a rare case, the TCP connection between Routing Engine and FPC is mistakenly enabling re-transmit timer for pure ACK's which is causing the FPC to reboot. [PR858489: This issue has been resolved.](#)
- With nonstop active routing (NSR) enabled, while performing graceful Routing Engine Switchover (GRES), Junos OS fails to restore BGP peers' TCP connections on the new master Routing Engine's replicated socket due to it is not able to find the BGP peer address's route, causing BGP peers to flap with following logs: /kernel: jsr_sdrL_merge: PSRM merge failed 65 rpd[xx]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer a.b.c.d (Internal AS X) changed state from Established to Idle (event TcpSocketReplicationError) [PR862796: This issue has been resolved.](#)
- When a sonet interface with PPP encapsulation is used as forwarding next hop for the IPv6 remote router loopback address on IPv6 BGP sessions, if the sonet link is down, the IPv6 BGP session might flap at same time although there is valid route via other interface. [PR863462: This issue has been resolved.](#)
- After enabling firewall filter of IPv6 on Aggregated Ethernet (AE) interface to block Micro BFD Packets (Dst Port 6784), kernel crashes continually on Master and Backup Routing Engine due to double free of memory. [PR864112: This issue has been resolved.](#)
- IPv6 Neighbor discovery(ND) failed after multiple GRES. Nexthop getting stuck in hold state forever. We also see that the neighbor state is in NO_STATE and it is on ND timer queue. In this condition, on ND timer expiry it never sends neighbor solicitation (NS) out and it never transitions to known ND states. Use "show route forwarding-table" CLI command to see the result of IPv6 route in hold state. root@ABC> show route forwarding-table Destination Type RtRef Next hop Type Index NhRef Netif 1234::56

/128 dest 0 1234::56 hold 1902 1 irb.5678 Use "show ipv6 neighbors" CLI command to see the result of IPv6 ND state in NO_STATE. root@ABC> show ipv6 neighbors IPv6 Address Linklayer Address State Exp Rtr Secure Interface 1234::56 none nostate 0 no no irb.5678 [PR864133: This issue has been resolved.](#)

Interfaces and Chassis

- In SFPC for M320 fabric hardening feature, self Ping Packets are sent periodically during each FPC periodic. Due to certain race condition introduced by this self ping feature, FPC may crash while sending self-ping packet or any host generated packet. Crash occurs due to the resource contention to use the shared resources among the different Packet Forwarding Engine modules. Similar issue was reported on MX Series platform and addressed with PR 701928. [PR732806: This issue has been resolved.](#)
- There can be a mismatch between the ifIndex value on IF-MIB-ifName and the ifIndex value on SONET-APS-MIB-apsMapGroupName and apsMapEntry. [PR771877: This issue has been resolved.](#)
- This issue is specific to the M120 hardware since there are two independent FRU's from where the PIC needs to be detached/attached. This IPC messages goes out-of-order due to the additional control-plane messages related to routing-change as a result of PIC restart which happens in this case due to the buffer configuration change. When PIC needs to be detached and at the same time there is still a lot of protocol information which should be processed as well, the detached messages will NOT be able to be delivered in time. After PIC restarts, it requests to be attached again but obviously this action failed because from other FRUs perspective the PIC has NOT been detached at all. [PR773081: This issue has been resolved.](#)
- Hash Key configuration not programmed in Packet Forwarding Engine correctly after system reboot. [PR818035: This issue has been resolved.](#)
- Faulty SCG causes continuous interrupts to HCFPC making its CPU Utilization 100% and unusable for any service. As a fix the monitoring mode for the SCG is changed to polling status of SCG device rather than interrupts based awake and monitoring system. [PR827489: This issue has been resolved.](#)
- Removing IP address on ATM interface after adding another IP address from the common subnet can lead to a race condition. New IP address configured on the interfaces still referring to shared broadcast-nexthop. Then when TCP/IP access, this broadcast-nexthop, kernel panic may happen. [PR833015: This issue has been resolved.](#)
- When packet has to be forwarded over NH topology unilist->indirect-indexed and when the packet size is greater than egress interface MTU w/ DF set, then we may log the following message and not send the message back to source indicating "frag needed and DF set". fpc0 NH: Can not find logical interface for nh 1048590 fpc0 NH(nh_get_mtu_iff) : get unilist mtu failed [PR844987: This issue has been resolved.](#)
- Whenever tunnel interface -pe/-pd got created using the MS-DPC instead of the MPC it will not be able to process register messages. Because MPC and MS-DPC have different multicast architectures and they are incompatible, if chassis is configured in "enhanced-ip" mode this issue will be seen. Necessary changes has been made to the code so that these interfaces will not be created on MS-DPC. [PR853995: This issue has been resolved.](#)

- SDG : After rebooting both Routing Engines together, the FPCs and MS-DPCs may come online, go offline (with "Chassis connection dropped" and "Chassis Manager terminated" error messages) and come back online again automatically. This issue is seen only when both Routing Engines are rebooting at once. There is exactly one additional reboot of the FPCs when this happens, and the FPCs come back up online, and system stabilized by itself within 2 to 3 additional minutes.[PR854519: This issue has been resolved.](#)
- After perform GRES, IPv6 Neighbor discovery on AE interfaces may fail, and explicit ping from Routing Engine will solve the issue. [PR854619: This issue has been resolved.](#)
- Routing Engine reset with Fatal trap 12: page fault while in kernel mode [PR855317: This issue has been resolved.](#)
- Interface hold-time-down is not working properly for PIC type 10x10GE(LAN/WAN) SFPP [PR859102: This issue has been resolved.](#)
- ISSU does no support VRRP. [PR862052: This issue has been resolved.](#)
- Injecting Enhanced RDI-P(G1 bit5-7:0x2 Payload defect) alarm to a MPC 10GbE WAN-PHY interface causes RDI_P and LCD-PAIS-V alarm on messages. This is due to string typo. RDI_P and LCD-P should be printed on messages.[PR872133: This issue has been resolved.](#)
- Both VRRP routers keep backup-backup state until "startup-silent-period" expires if both "startup-silent-period" and "delegate-processing" are configured.[PR873488: This issue has been resolved.](#)

Layer 2 Features

- When VPLS is configured with GRES, the backup Routing Engine responds to certain route replication requests by simulating address learning. If the route being replicated is associated with an LSI or VT interface, the address learning code references a special LSI or VT nexthop. Thus, there is a dependency between that route and that nexthop. This fix is to explicitly enforce this ifstate dependency , ensures that the special nexthop is seen by the peer before the route. [PR867929: This issue has been resolved.](#)

Layer 2 Ethernet Services

- DHCPv6 fails for clients using DUID type 2 (Vendor-assigned unique ID), the software was using the DUID to extract MAC address information. This behavior is fixed and tested [PR838404: This issue has been resolved.](#)
- When DHCPv4 relay is configured on an Integrated Routing and Bridging (IRB) interface with both IPv4 and IPv6 families configured, when remove "family inet6" configuration from the IRB, DHCPv4 relay function broken. This happens regardless of whether the "family inet6" is configured directly under the IRB or applied through an "apply-group" configuration. In versions that do not have the fix for this PR, the workarounds to get the dhcp relay functionality working again over the IRB are **either** of the following: 1) deactivate/activate the IRB configuration. 2) Restart dhcp daemon using the following command. `user@host> restart dhcp-service` [PR870543: This issue has been resolved.](#)

Multiprotocol Label Switching (MPLS)

- The cleanup procedures may leave transient inconsistent references when the interface address of an MPLS enabled GRE or IPIP tunnel is being deleted or the action taken implies an internal reconfiguration of the interface address (for example MTU change). During this period, if these references are being reused by a particular task, the kernel may report an invalid memory access and restart. [PR844790: This issue has been resolved.](#)
- The routing protocol daemon (rpd) might leak memory when there are MPLS LSP changes, the memory leak could eventually cause rpd process to crash. [PR847354: This issue has been resolved.](#)
- ASBR might not rewrite EXP correctly for egress MPLS packets on the Inter-AS link for the eBGP-LU LSP if the eBGP session is a multihop BGP session. [PR864914: This issue has been resolved.](#)

Network Management and Monitoring

- Flapping interfaces in combination with restarting chassisd/dcd/mib2d daemons and other abnormal scenarios coupled with SNMP polling create race conditions in mib2d. This results in a mib2d core. [PR812019: This issue has been resolved.](#)
- The 'timestamp year msec' command in syslog (not using structured data) is intended to include year and msec details in local syslog messages stored in rotating files, but not to be included in messages sent to a remote syslog collector. This fix corrects a wrong behavior introduced in 11.3R1 where such details were also included in syslog messages sent to a remote host. [PR820436: This issue has been resolved.](#)

Platform and Infrastructure

- XML tags for get-software-information output missing some elements of new Junos OS service release naming convention. [PR783653: This issue has been resolved.](#)
- Tunnel services (MT-x/x) using MPC on PE installs (S,G) route on receiving IGMP report. [PR821893: This issue has been resolved.](#)
- In a race condition where multiple interrupts are asserted, timer tick may not get well handled and remain asserted. This caused FEB panic and core. [PR828496: This issue has been resolved.](#)
- file-archive rpc throws out error - 'Operation allowed only from CLI' : [PR831865: This issue has been resolved.](#)
- Due to a bug in logical interface localization, a DPC restart/offline may cause a removal of legitimate CCC routes on other DPC's. This can also be triggered by removal of an unrelated family CCC logical unit. [PR835216: This issue has been resolved.](#)
- FPC core dump with the feature copy-plp-all enabled when add link to existing AE interface, which is part of downstream interface list of a multicast route. [PR842046: This issue has been resolved.](#)
- When a junoscript <get-configuration> RPC query, by default the query is done on candidate DB, a MGD process is spawned to handle this request. Now at the same time via another session if the configuration is deleted it is possible for the above

spawned MGD process performing the JunoScript query to crash. MGD process crashes while accessing a NULL parent which contained an object previously which was deleted. The fix addresses this by not exporting the object which has no parent. [PR844795: This issue has been resolved.](#)

- mlfrr/mlppp interfaces are not reachable after restart FPC (primary MSPIC) followed by deactivate and activate R.I or GRES followed by deactivate and activate R.I. This is because link FPC does not have the interfaces programmed towards the bundle [PR847278: This issue has been resolved.](#)
- If routing-instance is popping the mpls label through vt tunnel interface and the egress interface MTU of the vrf needs fragmentation and the dont-fragment bit is set in the ipv4 header, the egress vrf interface might stop forwarding traffic. The following syslog message will be reported fpc4 LCHIP(3): 1 new errors in LSIF To recover from this condition you can either bring the interface down via disable knob or deactivate/activate the interface from the configuration. The following platforms are exposed to this condition: M320 (excluding E3 FPCs), T/TX systems (excluding ES FPCs and FPC Type 5) [PR854806: This issue has been resolved.](#)
- In the T4000 Type 5 FPC platform, aperture management can lead to a collision between the sched tick timer and asic driver interrupt handlers, which will result in FPC crashes. [PR857167: This issue has been resolved.](#)
- On MX Series routers, with some logical interfaces of an aggregated Ethernet (AE) interface attached to a bridge-domain and LACP is enabled on the AE interface, after disabling/enabling or removing/adding one or more member links of the AE interface, because the receive channel of the AE interface is closed when LACP state is down, traffic loss might be observed for several seconds. [PR858124: This issue has been resolved.](#)
- In IPFIX context: 1. In an IPv6 single stack environment, when exporting Data and Template records for family IPv6, the Template records sequence number is not initialized and is always == 0 for all records. This is because the Template sequence numbers are blindly copied from family IPv4 and if this is not configured for IPFIX, then the Sequence Number is always 0. 2. In an IPv4 + IPv6 dual stack environment, since the Template records sequence numbers will be identical for both families, we will get Data and Template records sequence numbers being interleaved when exported. This could confuse the Flow Collector and mislead it into reporting random missing flows. [PR859169: This issue has been resolved.](#)
- BOOTP request packets might get dropped because of the DDOS protection feature in Trio based platforms. In this case, the bootp packet is coming with 1 byte option. So the length of bootp become 241 which is larger than 240. Then Packet Forwarding Engine will identify it not as BOOTP as per the current DDOS algorithm, and tries to parse it as DHCP. Since the packet lacks the options fields which need for DHCP, then pfe_nhdb_dhcpv4_msg_type() mark it as DHCPNOMSGTYPE. [PR862206: This issue has been resolved.](#)
- In some corner cases SPMB can stuck in READY state. Restarting the SPMB does not help to recover from the problem state. The issue is fixed by this PR. [PR866127: This issue has been resolved.](#)

- On MX Series routers with DPC (ICHIP based) type FPCs running a 11.4 (or newer) Junos OS release disabling uRPF on a logical interface might result in another logical interface on the router to drop all incoming packets. This problem happens only when the following conditions are met concurrently: a) 2 different logical interfaces share the same lookup index b) both logical interfaces have uRPF enabled c) these 2 different logical interfaces belong to 2 different FPCs d) at least one of the logical interfaces belongs to a DPC (ICHIP based) type FPC The lookup index is calculated by taking the lower 16 bits of the logical interface index (also called the IFL index). In other words $\text{lookup index} = \text{IFL index} \text{ MOD } 65536$. It is normal, valid and expected to have logical interfaces which share the same lookup index. The problem described in this PR is not the fact that the lookup indexes are the same. Here is an example of 2 different logical interfaces on 2 different FPCs which share the same lookup index: Interface ge-0/1/0.945 has an logical interface index of 1774 and a lookup index 1774:

```

user@router-re1> show interfaces ge-0/1/0.945 Logical interface ge-0/1/0.945 (Index 1774) (SNMP ifIndex 1635) ^^^^^^^^^^ Flags: Device-Down SNMP-Traps 0x4000
VLAN-Tag [ 0x8100.945 ] Encapsulation: ENET2 Input packets : 0 Output packets: 0
Protocol inet, MTU: 4462 Flags: Sendbcast-pkt-to-re, uRPF, uRPF-loose Addresses,
Flags: Dest-route-down Is-Preferred Is-Primary Destination: 52.3.168.216/29, Local:
52.3.168.217, Broadcast: 52.3.168.223 Protocol multiservice, MTU: Unlimited And
interface xe-2/2/0.0 has an logical interface index of 198382 and a lookup index of
198382 MOD 65536 = 1774: user@router-re1> show interfaces xe-2/2/0.0 Logical
interface xe-2/2/0.0 (Index 198382) (SNMP ifIndex 698) ^^^^^^^^^^^^^^ Flags:
SNMP-Traps 0x4004000 Encapsulation: ENET2 Input packets : 381 Output packets:
376 Protocol inet, MTU: 1500 Flags: Sendbcast-pkt-to-re, uRPF, uRPF-loose Addresses,
Flags: Is-Preferred Is-Primary Destination: 155.154.153.0/30, Local: 155.154.153.1,
Broadcast: 155.154.153.3 Protocol multiservice, MTU: Unlimited In the example above
if uRPF is disabled on ge-0/1/0.945 then xe-2/2/0.0 will start dropping all incoming
packets due to RPF failure. When this condition occurs the only way to recover is to
disable, commit and re-enable uRPF on the broken interface. When this is done the
following error messages are generated: Apr 15 16:02:53 router-re1 fpc2
rt_iff_generic_topo_handler: jtree error Not found for disconnect on iff-post-src Apr 15
16:02:54 router-re1 fpc2 RT(rt_rpf_jtree_drt_remove_ifl): Unable to remove logical
interface 198382 from drt(4) Apr 15 16:02:54 router-re1 fpc2
RT(rt_rpf_jtree_drt_remove_ifl): Unable to remove logical interface 198382 from loose(7)
PR873709: This issue has been resolved.

```
- This is a regression issue introduced by the fix of PR801982, which causes DOM MIB values for SFP+ "rx power" related statistics to be incorrect. Please note that XFP is not affected. [PR878843: This issue has been resolved.](#)
- Deactive/delete AE interface when route is flapping might cause Trio-based Packet Forwarding Engine to crash. [PR884837: This issue has been resolved.](#)
- In the case when the STFPC is on the CE side and trio fpc on the core side, ppp-ccc l2circuit connection will drop the small packets with Ethernet length error. [PR887098: This issue has been resolved.](#)

Routing Protocols

- Junos OS checks for mask-length mismatch for OSPF P2P-over-LAN interfaces, but skips the check if an interface has /32 mask configured. In a scenario with OSPF configured between Juniper platform and other vendor's platform, if a /32 mask IP address is configured on P2P-over-LAN OSPF interface of Juniper platform and a non /32 mask IP address is configured on the peer, the OSPF neighbor can establish but Kernel Routing Table (KRT) queue gets stuck. [PR840122: This issue has been resolved.](#)
- In subscriber management environment, routing protocol daemon (rpd) may crash and dump core due to snmpwalk fails at mplsL3VpnVrfRtInetCidrDestType when a subscriber access-internal route in a VRF has a datalink nexthop (such as when DHCP subscriber connects into a VRF). When issue happens, the following behaviors could be observed: user@router> show snmp mib walk ascii mplsL3VpnVrfRtInetCidr | no-more Request failed: Could not resolve 'mplsL3VpnVrfRtInetCidr' to an OID user@router> show snmp mib walk ascii mplsL3VpnVrfRtInetCidrDest | no-more Request failed: General error [PR840323: This issue has been resolved.](#)
- In IS-IS scenario, with graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) enabled, after Routing Engine switchover, in very rare case, the routing protocol process(rpd) might crash and create a core file on the new master (old backup) Routing Engine. This crash happens upon IS-IS lsp generation due to memory corruption. [PR841558: This issue has been resolved.](#)
- Under certain conditions moving a link that has BFD clients can cause stale BFD entry for the old link. [PR846981: This issue has been resolved.](#)
- Upstream interface of multicast rpf not matching multicast route in Inter-AS PIM. [PR847370: This issue has been resolved.](#)
- In multicast environment with PIM configured, in RP-on-a-stick scenario (aka one-legged RP), if the rendezvous point (RP) receives multicast traffic but there are no receivers, RP's kernel will keep sending resolve requests to the routing protocol process (rpd). These resolve requests might get stuck in resolve queue delaying other (S,G) resolves and thereby multicast traffic will be blackholed. [PR851210: This issue has been resolved.](#)
- If an invalid PIM-SSM multicast group is configured on the routing device, then when you issue the "commit" or "commit check" command, a routing protocol daemon (rpd) core file is created. There is no traffic impact because the main rpd process spawns another rpd process to parse the corresponding configuration changes, and the new rpd process crashes and creates a core file. When this problem occurs, you might see the following messages: user@router# commit check error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86) error: configuration check-out failed user@router# commit error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86) error: configuration check-out failed [PR856925: This issue has been resolved.](#)
- Routing protocol daemon (rpd) crashes and dumps core files when non-bgp routes (e.g. static route) being advertised as add-path route. [PR859307: This issue has been resolved.](#)
- RPD cored. [PR863148: This issue has been resolved.](#)

- In Release 12.1 MPLS OAM programs BFD, it does not provide the source address (no change in behavior). In BFD before programming PPMD it queries kernel for the source address matching the prefix of the destination address on a interface. BFD programs PPMD with this source address. PPMD will construct BFD packet with BFD provided source address in the IP header. [PR870421: This issue has been resolved.](#)
- In inter-AS Option-B L2VPN scenario, the ASBR might create a L2VPN cloned transit route incorrectly due to a cloned route is a Juniper specific mpls.0 route which Junos OS creates on the penultimate hop router. Then in a rare case, routing protocol daemon (rpd) tries to delete the L2VPN cloned transit route (in mpls.0 table) multiple times. After this, routing protocol daemon (rpd) crashes and dumps core. [PR878437: This issue has been resolved.](#)
- RPD CPU utilization keeps 100% due to "BGP resync" task when BGP is configured with no neighbor and NSR is configured. `id@router> show configure routing-options nonstop-routing; id@router> show configure protocols bgp { group bgp-group { type internal; inactive: neighbor 1.0.0.1; } }` [PR884602: This issue has been resolved.](#)

Services Applications

- ICMP echo request/reply packets are dropped when no reply is received for more than 64 echo-requests. This traffic must pass through MS-100/MS-DPC performing NAT/sfw to be impacted. [PR822596: This issue has been resolved.](#)
- The `jnxNatSrcNumPortInuse` counter is not refreshing when polling the `jnxNatSrcNumPortInuse` OID via SNMP after RSP switchover. [PR829778: This issue has been resolved.](#)
- MAC Flow-control asserted and MS-DPC reboot is needed [PR835341: This issue has been resolved.](#)
- 1) corrected the log to state 4 bundles per tunnel to have been exhausted. 2) change the log level from INFO to DEBUG 3) Add more context to previous log: New IPSec SA install time 1356027092 is less than old IPSec SA install time 1356027092 new log = Tunnel:<tunnel-id> <Local_gw, Remote_gw>: <local-gw-ip-addr, remote-gw-ip-addr> New IPSec SA install time 1356027092 is less than old IPSec SA install time 1356027092 4) added more context to previous log: SA to be deleted with index 3 is not present new log = SA to be deleted with index 3 is not present <Local_gw, Remote_gw>: <local-gw-ip-addr, remote-gw-ip-addr> 5) added a counter to show the number of times each of these messages occur per tunnel. [PR843172: This issue has been resolved.](#)
- This PR fixed syslog is not sent to remote host when rsp interface is used. [PR849995: This issue has been resolved.](#)
- `jnxNatSrcNumSessions` SNMP OID is broken in 11.4R6-S1 release [PR851989: This issue has been resolved.](#)
- Defining an application with destination-port range starting at 0 can cause TCP handshake to fail through NAT. As a workaround, specify the application with destination-port range starting at 1 instead of 0. [PR854645: This issue has been resolved.](#)
- The number of terms per NAT rule cannot exceed 200 for the inline-service si- interface. This constraint check is not applicable for other types of service interfaces like sp-, AMS and ms- etc. Following error message will be displayed when there are more than

200 terms per NAT rule: regress@aria# commit [edit services] 'service-set ss8' NAT rule rule_8 with more than 200 terms is disallowed for si-0/0/0.8 error: configuration check-out failed [PR855683: This issue has been resolved.](#)

- Using "destination-address 0::0/0" in SFWv6 presents a commit warning [PR857106: This issue has been resolved.](#)
- MS-DPC may crash in certain scenarios when using CGNAT PBA and junos-rsh, junos-rlogin, junos-rpc-services-udp and junos-rpc-services-tcp ALGs (either one) in combination with EIM. [PR862756: This issue has been resolved.](#)
- Service PIC might crash in corner cases when SIP ALG media flows are deleted. [PR871638: This issue has been resolved.](#)
- The issue is seen because of receiving malformed LCP configure-request packet with bad option length from PPP client. In this case when router tries to generate configure-nak it crashed. As a fix, check is added to discard such malformed configure-request packets. [PR872289: This issue has been resolved.](#)
- In a CGNAT environment when sp- interfaces, which are underlying rsp- interface, are present in the configuration, sp- interfaces service-options may wrongly overwrite rsp- interfaces service-options. [PR881792: This issue has been resolved.](#)

VPNs

- When "multicast omit-wildcard-address" is configured on a route-reflector for the MVPN address families, Leaf-AD route NLRIs are not reflected correctly in the newer, standardized format. The Leaf-AD routes transmitted from the RR in the new format will have invalid Leaf-IP fields in the NLRI set to 0.0.0.0. As a result, ingress PEs may fail to properly identify all egress PEs and thus fail to update provider-tunnel state to deliver traffic to those egress PEs. [PR854096: This issue has been resolved.](#)
- While l2circuit/l2vpn is not configured, if user requests for PW object info through mib, L2circuit/l2vpn is creating invalid job, which can lead to rpd crash. [PR854416: This issue has been resolved.](#)
- In a NG-MVPN scenario, on an ingress PE, if a RP is learned after receiving the BGP Type-6 route from egress PE, the ingress PE doesn't create PIM (*G) entries. This is seen only with dynamically learned RPs. With static RPs, after a commit, MVPN flashes the table and triggers creation of PIM (*G) entries. [PR866962: This issue has been resolved.](#)
- If CE facing interface is aggregated Ethernet with multiple member ports (more than 2 members), BUM traffic from MPLS core will be replicated on all child links of aggregated Ethernet interface and BUM from CE will be replicated at sending out from MPLS core facing interface. The problem is specific to M10i router running with I chip based CFEB. [PR880422: This issue has been resolved.](#)

Previous Releases

Resolved Issues in Junos OS Release 12.1R6

Class of Service (CoS)

- A few memory leaks have been fixed in the class of service daemon. [PR811613](#): This issue has been resolved.
- This cosmetic issue is specific of 3D linecards, based on MX Series routers with MPCs and MICs. In these cards, the logical interfaces with family mpls do not have any EXP rewrite rule applied by default. In other words, EXP value is copied from the previous codepoints: for example, from IP Precedence in IPv4->MPLS next hops. However, the **show class-of-service interface** command still shows the exp-default rule as if it was applied (in fact, it is not):

```
user@router> show class-of-service interface ge-2/3/1.204 | match rewrite Rewrite
exp-default exp (mpls-any) 33
```

[PR824791](#): This issue has been resolved.

- When rate limit is enabled and disabled on port cos scheduler configuration leaves rate limit configuration on queues in effect. This causes the rate limit feature in effect even after rate limit is removed. This PR addresses this issue in lieu with [PR843603](#). [PR833431](#): This issue has been resolved.
- Traffic-control-profile-remaining is not working for the logical interface in interface-set. [PR835933](#): This issue has been resolved.
- This seems to be hard to reproduce and noticed only once after GRES. When the cosd restarts (due to the GRES test you performed), cosd reconciles the configurations pushed to the Packet Forwarding Engine with configuration read from CLI and tries to reuse the object ID. In this case, it was trying to insert the same ID twice. [PR848666](#): This issue has been resolved.
- Commit throws an error "Invalid rewrite rule rule-name for logical interface <logical interface-name>. physical interface <physical interface-name> is not capable to rewrite inner vlan tag 802.1p bits" even though there is no rewrite configuration related to inner-vlan tag. [PR849710](#): This issue has been resolved.

Forwarding and Sampling

- There is always a chance to see this issue if any daemon adds a blob size which comes closer to 65520 (after IDR encoding). [PR700635](#): This issue has been resolved.
- With more than four archive-sites configured under the **[system archival configuration archive-sites]** hierarchy, after committing the configuration changes, the pfd process crashes and generates a core file due to memory corruption or double free. The core files could be seen by executing the **show system core-dumps** command. [PR849465](#): This issue has been resolved.
- MPLS forwarding table filter (ftf) not getting linked in JTREE after router or FPC reboot. [PR851599](#): This issue has been resolved.

General Routing

- Prior to this change, the L2TP sessions with cos/ firewall attachments fail to come up when the L2TP Access Concentrator (LAC) is reachable over a unilist nexthop. [PR660208](#): This issue has been resolved.
- The L2ald process might crash during issue with the core file and will fail to start the process. [PR731147](#): This issue has been resolved.
- MPLS LDP traceroute does not work if you have a default route 0/0 pointing to discard on the egress router with DPC cards. [PR790935](#): This issue has been resolved.
- On T1600-FPC4-ES, T640-FPC3, T640-FPC3-E, and T640-FPC3-E2 platforms that have multiple Packet Forwarding Engines, with auto-bandwidth enabled on LSPs where CoS-based forwarding (CBF) is configured, auto-bandwidth might trigger minor changes on LSP next hops. After this, flapping corresponding interface or any next-hop changes might result in an FPC crash and generate the core file. The core files can be seen by executing the CLI command **show system core-dumps**. This issue will be seen with auto-bw configuration where there will be continuous minor/major changes on LSP next-hops based on traffic conditions. When this issue happens, the following logs could be seen:

```
fpc3 PDP(pdp_free): %PFE-3: Invalid PDP 0x4e01d7d0 fpc3 PDP(pdp_free): %PFE-3:
Error while removing PDP (0x4df4c068) fpc3 PDP(pdp_free): %PFE-3: Error while
removing PDP (0x525b3f78) fpc3 PDP(pdp_free): %PFE-3: Invalid PDP 0x4de522b0'
```

[PR818021](#): This issue has been resolved.

- In a race condition where multiple interrupts are asserted, timer tick might not get well handled and remain asserted. This caused panic and generated a core file. [PR828496](#): This issue has been resolved.
- When an MS-DPC PIC reboots due to a crash or manual intervention, it might get stuck in a booting loop if the MS-DPC up-time is more than 49 days and 17 hours. After 5 consecutive boot failures, the MS-DPC PIC will go offline automatically and gives the following error message:

```
[ 15:21:22.344 LOG: Err] ICHIP(0): SPI4 Training failed while waiting for PLL to get locked,
ichip_sra_spi4_rx_snk_init_status_clk [ 15:21:22.344 LOG: Err] CMSPC: I-Chip(0) SPI4
Rx Sink init status clock failed, cmsdpc_spi4_init [ 15:21:22.344 LOG: Err] CMX: I(0)
ASIC SPI4 init failed [ 15:21:22.379 LOG: Err] Node for service control ifl 68, is already
present [ 15:21:23.207 LOG: Err] ASER0 SPI-4 XLR source core OOF did not go low in
20ms. [ 15:21:23.208 LOG: Err] ASER/XLR0 spi4 stop src train failed! [ 15:21:23.208
LOG: Err] ASER0 XLR SPI-4 sink core DPA incomplete in 20ms. [ 15:21:23.208 LOG:
Err] ASER/XLR0 spi4 sink core init failed! [ 15:21:24.465 LOG: Err] ICHIP(0): SPI4 Stats
Unexpected 2'b11 Error, isra_spi4_parse_panic_errors [ 15:21:24.465 LOG: Err] ICHIP(0):
SPI4 Tx Lost Sync Error, isra_spi4_parse_panic_errors
```

In order to recover from this state the whole MS-DPC needs to be rebooted. [PR828649](#): This issue has been resolved.

- In a PPPoE subscriber management environment, with "dynamic-profiles/routing-instances/\$junos-routing-instance/routing-options /access-internal/route" configured for PPPoE subscribers, in some rare conditions, the dynamic-profile database gets into a state where the routing protocol process (rpd) cannot read or access, causing PPPoE sessions to not be established due to the

rpdd process being unable to make calls to the dynamic-profile to add the access-internal routes. When this issue happens, the following logs could be seen:

```
NACK received for profile request with id=ac3d694 from rpdd daemon: Get dynamic
profiles failed: 212 retry FALSE rpdd[1382]: RPD_DYN_CFG_GET_PROFILE_FAILED: Get
dynamic profiles failed: 212
```

[PR830779](#): This issue has been resolved.

- When an FPC goes bad due to the hardware failure is stuck in a boot mode, it might affect the Routing Engine-Packet Forwarding Engine communication on other FPCs. [PR831233](#): This issue has been resolved.
- When the transit traceroute packets with ttl=1 are received on the LSI interface, you might retrieve the Source Address from the LSI interface to reply ICMP. As LSI does not have any IFA, it will use first the IFA in routing-instance to reply. So Source Address used was the first IFA added in VPN routing-instance. As a workaround, if the incoming interface is LSI, then retrieve the Source Address from the logical interface which is having the Destination IP Address. This will make sure we reply with Source Address from the CE facing the logical interface. [PR839920](#): This issue has been resolved.
- Reception of Ethernet packets containing an invalid Ether-Type can cause congestion within a Host Notification Queue on a router with M Series routers with FPCs, such as the M120 and M320 multiservice edge routers, and the M10i/M7i CFEB-E. This in turn can lead to the dropping of valid protocol traffic from reaching the Routing Engine and impact certain time-sensitive protocols such as LACP. Once the invalid Ethernet traffic stops, the queues drain and normal operation will continue without intervention. Refer to KB25385 for a mapping of chipset type to the Packet Forwarding Engine module. Refer to PSN-2013-04-916 for more information. [PR847603](#): This issue has been resolved.
- Distributed protocol adjacencies (LFM/BFD/etc) might experience a delay in keepalives transmission and/or processing due to a prolonged CPU usage on the FPC microkernel on T4000 Type 5-3D FPCs. The delay in keepalive transmission/processing can result in a mis-diagnosis of a link fault by the peer devices. The issue is seen several seconds after an Routing Engine mastership switch with NSR enabled, and the fault condition will clear after a couple of minutes. [PR849148](#): This issue has been resolved.
- FPC/PICs usually have high response time when they are loaded with high traffic. Kernel generates a core file when a PIC/FPC stops responding to the kernel after a new connection or a reconnection. [PR853296](#): This issue has been resolved.

High Availability (HA) and Resiliency

- The PR fix attempts to synchronize configuration again after a connection to the master succeeds if the configuration sync had failed earlier. [PR783832](#): This issue has been resolved.

IPv6

- With the fix IPv6 VRRP will not inter-op across new release and old release "with version 3 disabled", because of correction in checksum calculation. [PR826734](#)

Infrastructure

- If a router is configured with a POSIX compliant time-zone string, it does not update the time zone correctly. The problem can be observed in system logs and CLI commands when date/time zone is referenced. set groups re0 system time-zone EST5EDT,M10.3.0/2,M2.3.0/2 The router will incorrectly reference the previously configured time zone. After time zone configuration modifications, run the **commit full** command. [PR785946](#): This issue has been resolved.
- A kernel crash might occur on the routers running 10.4 or higher (which does not have a fix for this PR), with **targeted-broadcast** knob configured on a broadcast interface. If this knob is configured, the MAC address will be learned for subnet broadcast IP (configured on that interface). When this ARP table entry gets timed out, it corrupts an internal data structure, leading to kernel crash. This MAC learning will happen with one of the following :
 1. Mismatched IP subnet is configured on one of the connected devices
 2. A malformed packet (ARP request to subnet broadcast IP) is received on that interface



NOTE: MAC address learned for the subnet broadcast IP can't be seen using **show arp** command. This issue is platform independent.

[PR814507](#): This issue has been resolved.

Interfaces and Chassis

- The password recovery process does not work on some MX80 routers. [PR585092](#): This issue has been resolved.
- Under certain circumstances, MX80 might crash when using the **request system snapshot** command. [PR603468](#): This issue has been resolved.
- A vmcore is seen when file system corruption occurs during Junos OS upgrade. [PR683554](#): This issue has been resolved.
- Kernel can cache a high incorrect value for stats and is rejecting the correct subsequently statistics coming from the PIC. The fix consists of checking if the difference of what is cached in kernel and what is reported by the PIC is less than an acceptable value. If

the answer is no, the kernel does not get stuck permanently and recovers while fetching statistics next time. [PR806015](#): This issue has been resolved.

- On MX80 platform, a broadcast storm on fxp0 (out-of-band interface) will cause the process "irq32:tsec2" to consume enough CPU causing the Routing Engine to lose the connection with Forwarding Engine Processor (TFEB). Then the Routing Engine declares TFEB unreachable, then the chassisd process shuts down causing all interfaces to be removed and traffic loss. A few moments later the Routing Engine is able to re-establish connectivity to the TFEB, and the Packet Forwarding Engine components begin to get re-initialized. But if the broadcast storm on fxp0 still exists, this issue will happen again. If issue happens, the following logs could be seen:

```
/kernel: Interrupt storm detected on "irq32:"; throttling interrupt source /kernel:
peer_inputs:3690 VKS0 closing connection peer type 17 idx 0 err 5 /kernel:
pfe_send_failed(index 0, type 17), err=32 /kernel: pfe_listener_disconnect: conn dropped:
listener idx=0, tnpaddr=0x80000032, reason: none chassisd[1204]:
CHASSISD_SHUTDOWN_NOTICE: Shutdown reason: TFEB connection lost
chassisd[1204]: CHASSISD_IFDEV_DETACH_FPC: ifdev_detach_fpc(0) chassisd[1204]:
CHASSISD_SNMP_TRAP10: SNMP trap generated: Fru Offline (jnxFruContentsIndex
7, jnxFruL1Index 1, jnxFruL2Index 0, jnxFruL3Index 0, jnxFruName FPC @ 0/*/*;
jnxFruType 3, jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0,
jnxFruLastPowerOn 0) chassisd[1204]: CHASSISD_IFDEV_DETACH_FPC:
ifdev_detach_fpc(1) chassisd[1204]: CHASSISD_SNMP_TRAP10: SNMP trap generated:
Fru Offline (jnxFruContentsIndex 7, jnxFruL1Index 2, jnxFruL2Index 0, jnxFruL3Index 0,
jnxFruName FPC @ 1/*/*; jnxFruType 3, jnxFruSlot 1, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 0) chassisd[1204]:
CHASSISD_IFDEV_DETACH_ALL_PSEUDO: ifdev_detach(pseudo devices: all)
alarmd[1205]: shutting down chassisd connection: chassisd ipc pipe read error
alarmd[1205]: chassisd alarmd[1205]: connecton succeeded after 0 retries
alarmd[1205]: resending alarm state craftd[1206]: craftd_user_conn_shutdown: socket
5, errno = 0 craftd[1206]: chassisd connection succeeded after 0 retries chassisd[1204]:
CHASSISD_SNMP_TRAP7: SNMP trap generated: FRU insertion (jnxFruContentsIndex
6, jnxFruL1Index 1, jnxFruL2Index 0, jnxFruL3Index 0, jnxFruName TFEB, jnxFruType 5,
jnxFruSlot 0)
```

The cpu usage of process "irq32:tsec2" could be observed by following command:

```
user@router> show system processes extensive | match "aver|PID|irq32" last pid: 1380;
load averages: 0.76, 0.61, 0.36 up 0+00:23:47 09:46:40 PID USERNAME THR PRI NICE
SIZE RES STATE TIME WCPU COMMAND 31 root 1 -68 -187 OK 16K WAIT 1:03 0.63%
irq32: tsec2
```

[PR816253](#): This issue has been resolved.

- The **show interfaces redundancy** command might display secondary as down upon following sequence:
 - deactivate R.I.(that contains entire mfr logical interfaces)
 - restart fpc (that holds secondary MS pic)
 - activate the R.I. back

[PR816595](#): This issue has been resolved.

- Warning message added is syslog when external sync is not supported. [PR817049](#): This issue has been resolved.

- Prior to this PR, the speed of a GE interface capable of working at FE speeds was set to 'auto' in the Packet Forwarding Engine level. This causes a problem when manually setting the speed on the Routing Engine. Now the behavior is to set the speed to '1 g' in the Packet Forwarding Engine. For automatic speed detection the interface should be set to 'speed auto' in the configuration. [PR821512](#): This issue has been resolved.
- MX's chassis-control interrupt storm may be falsely reported when a Field Replaceable Unit (FRU) is removed, inserted, or FPM button pushed. A FRU may not be recognized/booted, resulting in chassis operational failure. [PR823969](#): This issue has been resolved.
- IEEE 802.3 ah LFM stats counter "OAM current frame error event information" is not cleared correctly by CLI operation. [PR827270](#): This issue has been resolved.
- When you try to delete a physical interface configured with per-unit-scheduler along with its logical interfaces in one single commit, the ksyncd process might generate a core file in the backup Routing Engine which will cause the GRES to malfunction. [PR827772](#): This issue has been resolved.
- If we receive an MAC Move event or an L2 logical interface change event, we don't immediately remove the next hops. The Backup Routing Engine has to delete the NH first and then it gets deleted from Master. During this phase if pointer is stale to the NHs as in the pointer pointing to the NH is valid, but the Nh has already been deleted then you will run into this condition. [PR829093](#): This issue has been resolved.
- Currently, no SNMP trap generated when FPC crashes. This PR is meant to enable the SNMP trap for such failure. [PR835112](#): This issue has been resolved.
- Currently, no SNMP trap sent when backup SPMB failure happened. This PR is meant to enable the SNMP trap for such failure. [PR835167](#): This issue has been resolved.
- Although physical interface is disabled, reseating 1GbE SFP on MPC/MIC restores its output optical power, hence the opposite router interface turns Up(Near-end interface is still down). Only 1g-SFP on MPC/MIC has the problem, but 1g-SFP on DPC/MX, EX series and 10G-XFP on DPC/MX don't have the problem. When the sfp is reseated, then the sfp periodic is going ahead and enabling the laser irrespective of the fact that interface has been enabled or disabled. Driver needs to store the state for each sfp link and enable laser based on that. This software problem is fixed in 11.4R7, 12.1R6, 12.2R4, 12.3R2 and later release. [PR836604](#): This issue has been resolved.
- Configuring 100-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning. The "optics-options" alarm/warning "low-light"; the syslog action was not taking effect on T1600 and T4k for 100 GE PICs. This was fixed as part of this PR. [PR836709](#): This issue has been resolved.
- It is possible that when a DPC boots up and link-training fails for all the links between fabric planes and the FPC, the interface is still brought up online and traffic blackholing will occur. [PR839076](#): This issue has been resolved.

- The logical interfaces are marked with 0 (null) after deactivate system commit synchronize and deactivate chassis redundancy which result the backup Routing Engine to generate a core file. [PR840167](#): This issue has been resolved.
- In PPPoE subscriber management environment, PPPoE daemon may crash and dump core in following two scenarios: 1 - Firewall Filter/Policer is not configured on Broadband Remote Access Server (BRAS) side, and AAA pushes the filter name in "Ingress Policy Name/Egress Policy Name" which will expire the lockout timer waiting to create required dynamic interface, and eventually causes pppoe process crash. 2 - When IPv6 only capable modem is trying to connect and the configuration does not contain IPv6 dynamic configuration; i.e. under PPPoE dynamic profile/family inet6 stanza; PPPoE dynamic profile/protocols/router-advertisement, this will again expires lockout timer waiting for dynamic interface creation, which crashes pppoe process. [PR859000](#): This issue has been resolved.

Layer 2 Ethernet Services

- It can happen that when changing an interface framing from lan-phy (default) to wan-phy and back a few times, the interface doesn't show up any more in **show interfaces terse**. [PR836382](#): This issue has been resolved.
- DHCPv6 relay terminates the client if DHCPv6-REPLY message from server contains status-code option [PR845365](#): This issue has been resolved.

Multiprotocol Label Switching (MPLS)

- The RPD process might crash when executing the command **clear mpls lsp name <lspname>** or **monitor label-switched-path <lspname>** [PR756551](#): This issue has been resolved.
- If the current configuration is as follows:

```
label-switched-path lsp1
{
  to XX.XX.XXX.XX;
  primary path1;
  secondary path2
{
  standby;
}
}
```

If the following configuration change is made (delete the LSP, reconfigure the LSP and make path2 as primary and path1 as standby) - delete protocol mpls label-switched-path lsp1 set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX primary path2 set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX secondary path1 standby commit It will result in a stale 'path2' standby. Later if another configuration change happens for 'path2', it can point to the stale entry and result in the assertion failure and core. The workaround is to do a 'commit' after deleting the LSP in the above configuration. Thus the configuration steps become - delete protocol mpls label-switched-path lsp1 commit set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX primary path2 set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX secondary path1 standby commit. [PR847038](#): This issue has been resolved.

Network Management and Monitoring

- The issue has been fixed where incorrect query can cause the Routing Engine CPU to go high. [PR771867](#): This issue has been resolved.
- The default maximum log file size depends on the platform type for TX Matrix or TX Matrix Plus routers. It is expected to be 10 MB. However, due to a software defect, this file size was only 1 MB. [PR823143](#): This issue has been resolved.
- Expand the buffer size and set break point to allow sending out large snmp messages due to ospf down event. [PR827660](#): This issue has been resolved.
- On a router with interfaces with Frame Relay encapsulation a SNMP WALK operation will cause a MIB daemon (mib2d) crash and will generate a mib2d core-dump. The crash itself does not cause any impact on the router as the MIB daemon is restarted automatically. The only effect is that a SNMP WALK will never complete successfully.

```

user@router-re1> show snmp mib walk 1 | no-more sysDescr.0 = Juniper Networks, Inc.
mx480 internet router, kernel JUNOS 11.4R6.5 #0: 2012-11-28 21:57:12 UTC
builder@evenath.juniper.net:/volume/build/junos/11.4/release/11.4R6.5/obj-i
386/bsd/kernels/JUNIPER/kernel Build date: 2012-11-28 21:39:15 UTC Copyright (c
sysObjectID.0 = jnxProductNameMX480 sysUpTime.0 = 339594 sysContact.0 <
..... > dot3OutPauseFrames.942 = 0 dot3OutPauseFrames.943 = 0
dot3OutPauseFrames.953 = 0 dot3OutPauseFrames.954 = 0 frDlcmilfIndex.153 =
153 frDlcmilfIndex.512 = 512 frDlcmilfIndex.513 = 513 frDlcmiState.153 = 6 Request
failed: General error user@router-re1> show log messages Dec 20 09:23:20 router-re1
clear-log[8240]: logfile cleared Dec 20 09:23:38.683 router-re1 /kernel:
%KERN-3-BAD_PAGE_FAULT: pid 7382 (mib2d), uid 0: pc 0x810fe09 got a read fault
at 0x7c, x86 fault flags = 0x4 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3:
Trapframe Register Dump: Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: eax:
00000000 ecx: bfbeda88 edx: 00000000 ebx: bfbeda7c Dec 20 09:23:38.683
router-re1 /kernel: %KERN-3: esp: bfbeda60 ebp: bfbeda98 esi: 089de834 edi:
089fb680 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: eip: 0810fe09 eflags:
00010297 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: cs: 0033 ss: 003b ds:
bfb003bes: 003b Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: fs: 003b trapno:
0000000c err: 00000004 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Page
table info for PC address 0x810fe09: PDE = 0x42e60067, PTE = 5290c425 Dec 20
09:23:38.683 router-re1 /kernel: %KERN-3: Dumping 16 bytes starting at PC address
0x810fe09: Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: 8b 40 7c 89 04 24 e8
5a 3f 2f 00 89 45 ec 8b 55 Dec 20 09:23:40.787 router-re1 init: %AUTH-3: mib-process
(PID 7382) terminated by signal number 11. Core dumped! Dec 20 09:23:40.787
router-re1 init: %AUTH-6: mib-process (PID 8247) started Dec 20 09:23:40.809
router-re1 mib2d[8247]: %DAEMON-5-LIBSNMP_SA_IPC_REG_ROWS:
ns_subagent_register_mibs: registering 88 rows Dec 20 09:23:41.595 router-re1
mib2d[8247]: %DAEMON-6-LIBSNMP_NS_LOG_INFO: INFO:
ns_subagent_open_session: NET-SNMP version 5.3.1 AgentX subagent connected Dec
20 09:23:43.533 router-re1 dumpd: %USER-5: Core and context for mib2d saved in
/var/tmp/mib2d.core-tarball.0.tgz Dec 20 09:23:43.793 router-re1 mib2d[8247]:
%DAEMON-6-SNMP_TRAP_LINK_UP: ifIndex 5, ifAdminStatus up(1), ifOperStatus
up(1), ifName dsc < ..... > user@router-re1> show system core-dumps
/var/crash/*core*: No such file or directory -rw----- 1 root field 680417 Dec 20 09:23
/var/tmp/mib2d.core-tarball.0.tgz /var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory /tftpboot/corefiles/*core*: No such file
or directory total 1

```

[PR835722](#): This issue has been resolved.

- Under certain conditions, duplicate SNMP index might be assigned to different interfaces by kernel to mib2d (Management Information Base II daemon). This might cause mib2d and other daemon such as lacpd (LACP daemon) to crash, which needs to be known about the SNMP index of an interface. [PR836823](#): This issue has been resolved.

Platform and Infrastructure

- Fixed the issue where `\n` (\ and n as two separate characters) if present in the database, get displayed as `\\n` to the user. [PR705067](#): This issue has been resolved.
- When cscript data memory limit is exceeded when executing op, event or commit script, the cscript process could reset and leave a core file when failing to allocate memory past its limits. The script would not succeed to run, in the case of a commit script the commit would not succeed. [PR722161](#): This issue has been resolved.
- The output of the following commands: `- "cli > show route forwarding-table vpn <vpn-name> interface-name <interface-name>" - rtinfo` might show up negative values for Ipkts and Opkts counters due to invalid format of output for the variables. eg. `user@test> show route forwarding-table vpn test-vpn interface-name ge-1/0/8`
Name Mtu Network Address Ipkts Ierr Opkts Oerr Coll ge-1/0/8 1522 <Link>
00.1d.b5.27.48.ad -891806008 0 -1176087381 0 0 The fix also corrects the output format for input and output bytes fields. [PR798999](#): This issue has been resolved.
- On TX/TXP multi-chassis systems, time synchronization between SCC/SFC chassis and the LCC chassis is not maintained. [PR811480](#): This issue has been resolved.
- When using `configure private` with large group definition and high number of groups the commit process can spend a lot of time to merge the configuration change with the global configuration. [PR828005](#): This issue has been resolved.
- NPC core generated at `ns16550_write,system_console_nputs,console_putc_polled,print_string`.issue is fixed now. [PR835759](#): This issue has been resolved.
- The **deny** commands might not work for the **show route community-name** command. [PR836624](#): This issue has been resolved.
- This applies to all Juniper M, MX, and T Series routers. In certain GRES scenarios, the backup Routing Engine might not have the complete state of the NH database from the active Routing Engine and might send duplicate NH add messages to the Packet Forwarding Engine with same NH IDs when it becomes active. This could potentially cause undesirable behavior in forwarding resulting in broken forwarding state and/or FPC cores. To limit the effect of these duplicate NH add messages, only certain duplicate NH adds messages which can be handled gracefully are allowed and all other duplicate add messages are rejected. There is no work-around for this problem. [PR843907](#): This issue has been resolved.
- On MX Series routers and T4000 router, when output Filter-Based Forwarding (FBF) destined to a routing-instance is configured, the packets matched by the FBF filter might be discarded or sent to unintended Packet Forwarding Engine. [PR845700](#): This issue has been resolved.

Routing Protocols

- The routing protocol process (rpd) might crash when doing multiple GRES in combination with bgp peer flapping with large number of dampened routes. This is observed only when certain sequence criteria are met, but may not be exposed under all switchover conditions. [PR793875](#): This issue has been resolved.
- Due to duplication of the traffic, assert will be triggered. *G and S,G assert is not handled properly hence few assert entries will not be deleted due to the Routing Engine switchover which result in core. HW type of chassis/linecard/RE. "ALL" Suspected software feature combination. Multicast feature Describe if any behavior/ change to existing function - Handle the *G and S,G assert properly. [PR809338](#): This issue has been resolved.
- Changes to add-path prefix-policy do not get absorbed automatically, and require a manual soft-clearing of the BGP session [PR818789](#): This issue has been resolved.
- In subscriber management environment, with IGMP enabled for about 1k subscribers, routing protocol daemon (rpd) might crash and dump core while receiving IGMP join/leaves over these subscribers continuously. [PR828533](#): This issue has been resolved.
- Changing static route with qualified-next-hop and order option to next-hop option results in static route missing from route table. We need to restart routing process to see the route again [PR830634](#): This issue has been resolved.
- If LDP-SYNC <hold-down> timer is configured under ISIS interfaces after configuration change the ISIS interfaces can go to <hold-down> state. [PR831871](#): This issue has been resolved.
- ISIS reports prefix-export-limit exceeded even though the number of exported routes is smaller than the configured value of prefix-export-limit. [PR844224](#): This issue has been resolved.
- In scenarios that use BGP to distribute traffic flow specifications, if the received flow-spec Network Layer Reachability Information (NLRI) contains invalid argument (such as dscp is larger than 63), routing protocol daemon (rpd) will generate flow-spec routes and install them in the routing table for these NLRIs; but these flow routes with invalid match conditions are rejected by dynamic firewall daemon (dfwd) from being added to the flowspec filters. When issue happens, the following errors could be seen:
krt_flow_trans_match_config: Failed defining match conditions
10.0.1.1,1.0.0.1,proto=6,dscp=81 krt_flow_trans_term_add: Failed adding term
10.0.1.1,1.0.0.1,proto=6,dscp=81 to filter 0x9504000 - Unknown error: 0
krt_flow_trans_filter_add: Failed sending transaction (ADD FILTER SINGLE TERM) for filter 0x9504000 __flowspec_default_inet__ to add term 10.0.1.1,1.0.0.1,proto=6,dscp=81
- Invalid argument When the bgp peer withdraws these flow routes, they will only be deleted but not freed, hence cause memory leak. [PR845039](#): This issue has been resolved.
- In BGP scenario with multipath configured, if a static route which has table nexthop (such as inet.0) is configured in the same routing-instance as BGP, when an interconnect link between BGP peers is brought down or flapping, the corresponding BGP session takes 90 seconds to timeout. During this period routes received over the BGP session will stay there. For a multipath transit route received from both BGP sessions, initially

both paths are resolved over the interconnect links directly. When one of the interconnect links is brought down or flapping, that path will be resolved over the static default route which has table nexthop (such as inet.0). So now, one path is resolved over a router nexthop and the other path is resolved over a table nexthop. This will cause routing protocol daemon (rpd) crash and create a core file. This issue usually occurs in BGP/L3VPN environment. The core files could be seen by executing CLI command "show system core-dumps". [PR851807](#): This issue has been resolved.

Services Applications

- SIP ALG was not allowing SIP 603 decline message. [PR822679](#): This issue has been resolved.
- When MX Series uses MS-DPC to provide the tunnelling service for flow-tap traffic, if there is SCU/DCU configured on the same slot of the flow-tap traffic ingress interface, all the flow-taped sampled packets will be dropped. It is caused by the wrong next-hop linking when DCU is configured. [PR825958](#): This issue has been resolved.
- In the case of a stateful proxy, two SIP users behind the NAT device (so-called SIP hairpinning) will be unable to signal the call. [PR832364](#): This issue has been resolved.
- With RTSP ALG enabled, RTSP keep-alive packets might be dropped if it's already Ack'ed by the receiver. [PR834198](#): This issue has been resolved.
- In Carrier Grade NAT (CGNAT) scenario, without any configuration change, under some conditions, MS-DPC PIC may crash and core dumped when encountering unknown flow-type. Service will be impacted during the period. When issue happens, the following logs could be seen:

```
chassisd[1477]: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power off
(jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName
PIC: MS-DPC PIC @ 5/1/*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 8,
jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 33404122) chassisd[1477]:
CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex
8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @
5/1/*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 2, jnxFruLastPowerOff
192338801, jnxFruLastPowerOn 192338924)
```

[PR834899](#): This issue has been resolved.

- In the case of a transparent proxy, two SIP users behind the NAT device (so-called SIP hairpinning) will be able to signal the call, but the RTP voice flow *may* be unidirectional if one side started to send his RTP traffic before the port was opened on the other end, causing an ICMP unreachable which confuses the NAT device. [PR834933](#): This issue has been resolved.
- In scenarios which use sp interface, such as IPSec VPN, multiservice process (mspd) will memory leak during sp interface flapping. The memory usage of mspd process can be checked by following CLI command:

```
user@router> show system processes extensive | match "PID | mspd" (Note: The "RES"
field means "Current amount of resident memory, in kilobytes") PID USERNAME THR
PRI NICE SIZE RES STATE TIME WCPU COMMAND 2048 root 1 96 0 36216K 34820K
select 0:10 0.00% mspd When the memory usage of mspd process increases to system
limit (about 131072KB), the following logs could be seen: /kernel: %KERN-5: Process
```


(2048,mspd) attempted to exceed RLIMIT_DATA: attempted 131076 KB Max 131072 KB

[PR836735](#): This issue has been resolved.

- The **hot-standby** CLI knob under **[edit interfaces <RSP-interface-name> redundancy-options]** is made hidden for the Redundant Service PIC (RSP). [PR838762](#): This issue has been resolved.
- Service PIC might crash under certain race conditions when receiving sip invite packets. [PR843047](#): This issue has been resolved.
- Service PIC might crash in corner cases when receiving specific SIP REGISTER. [PR843479](#): This issue has been resolved.
- Service PIC might crash in corner cases when EIM is enabled for SIP ALG. [PR847124](#): This issue has been resolved.
- When allocate the memory from shared memory for bitmaps used in port blocks , Junos OS requests as many bytes as the size of the block. If customers assign like 10K block size for deterministic NAT or PBA, then Junos OS allocates 10K bytes for that bitmap. However, it only needs 10K/8 bytes as one byte can represent 8 ports. These huge allocations are leading to memory depletion when many source addresses are behind the NAT, and port blocks are big. [PR851724](#): This issue has been resolved.
- spd core generated during switchover when CGAT configuration is there. Issue is well understood now and has been fixed in later releases. [PR854206](#): This issue has been resolved.

Subscriber Access Management

- You cannot police IPv4 and IPv6 traffic for DHCP dual-clients to one rate when a logical interface policer is configured under a single policer that is referenced by two different family firewall filters. For example, inet, and inet6. [PR749912](#): This issue has been resolved.
- Snmpwalk requests sent to MX Series router returns multiple duplicate records for jnxUserAAAAccessPool. [PR840640](#): This issue has been resolved.

User Interface and Configuration

- Before this fix, if you used the Junos OS XML API to configure a password, the password was encrypted using an older algorithm than that used when configuring a password through the CLI. This older algorithm did not allow certain characters including commas. Any characters entered after the disallowed characters were ignored. As a workaround, configure the password from the CLI. [PR744595](#): This issue has been resolved.
- If a commit sync error occurs for a commit performed in **edit private** mode and later it is followed by another commit in global mode (without private or exclusive mode), the configuration file may remain unzipped after the global commit is complete. [PR823555](#): This issue has been resolved.

VPNs

- In a scaled Multicast VPN setup, where many selective provider tunnels are used, and the MVPN instance is deleted, RPD can sometimes crash. [PR801667](#): This issue has been resolved.
- In BGP-MVPN, when the number of multicast routes falls below the threshold, the earlier suppressed MVPN multicast routes because of limit are not added back again. For MVPN, there was no mechanism to trigger the processing of cmcast entries that were not added earlier. The fix is to queue the cmcast entries that are suppressed for multicast route addition in a new list. When the reuse limit is reached, this list is walked and used to add back the entries. [PR841105](#): This issue has been resolved.
- In L2circuit (Martini l2vpn) scenarios where a backup neighbor is being defined along the 'standby' knob, after deleting this backup neighbor from configuration, its associated vc-route is not being eliminated. Later, if user deletes the l2circuit neighbor or restarts routing protocol daemon (rpd), rpd process will crash and core dumped. [PR841522](#): This issue has been resolved.
- Deleted logical interfaces might not be freed due to references in MVPN. [PR851265](#): This issue has been resolved.

Resolved Issues in Junos OS Release 12.1R5

Class of Service (CoS)

- On IQ2E PIC, if the logical interfaces under two different interface-sets have their own scheduler-map configured, when the same iflset shaping parameters are applied on these two interface-sets, the iflset shaping profile is not shared between the interface-sets, hence, the number of iflset supported is reduced. There is no workaround. [[PR804158](#): This issue has been resolved.]
- When you configure the following commands on MX480 Router and execute **show class-of-service scheduler-map** command, Cosd might crash:


```
set class-of-service interfaces ge-2/0/* scheduler-map-chassis derived
set class-of-service interfaces ge-2/0/1 scheduler-map-chassis p0
set class-of-service interfaces ge-2/1/* scheduler-map-chassis derived
set class-of-service interfaces ge-2/1/1 scheduler-map-chassis p0
```

Scheduler-map-chassis for ge-2/0/1 and ge-2/1/1 has to be overridden because of the more specific configuration. This is happening correctly, but the 'derived' scheduler-map previously allocated for these interfaces is not getting freed. Cosd crashes when trying to display these scheduler maps since they are not fully populated. [[PR807593](#): This issue has been resolved.]
- When configured "scheduler-map-chassis derived" and auto bandwidth computation is not success, sosd might crash after the "show class-of-service scheduler-map" command. [[PR811586](#): This issue has been resolved.]

Forwarding and Sampling

- IPv6 supports a chain of optional internet-layer fields called extension headers (ref. Section 4 of RFC 2460). The current implementation of the Junos OS stateless firewall filter criterion 'next-header' will only match on IPv6 protocol (TCP or UDP) if extension

headers are not present in the packet. To resolve the issue of matching IPv6 protocol on packets containing one or more extension headers, a new 'payload-protocol' attribute has been added to the match clause of the ipv6 filter family hierarchy:

```

firewall {
  family inet6 {
    filter match-tcp6 {
      term t1 {
        from {
          payload-protocol tcp;
        }
        then count v6_tcp;
      }
    }
  }
}

```

This new attribute performs a similar function to the **next-header <protocol>** filter match attribute and allows matching on the payload protocol (TCP, UDP) whether extension headers are present or not. This may be useful when the packet header's Next Header field is not the protocol payload (eg. Authentication or Encapsulating Security Payload, etc.). For example:

IPv6 header	Auth header	ESP header	TCP header + data
Next Header = Authd	Next Header = ESP	Next Header = TCP	

This enhancement is currently only available for the Trio Packet Forwarding Engine chipset. Refer to KB25385 for a mapping of chipset type to Packet Forwarding Engine module. [[PR788180](#): This issue has been resolved.]

- On MPC sometimes the policer configuration does not get correctly programmed in Packet Forwarding Engine. The issue can affect Trio-based platforms with Junos OS 11.1 and above. It can be triggered by the following steps:
 1. An interface-specific filter that contains at least a policer is already applied on an interface.
 2. The parameters of the policer that is included by this filter are changed. Note that the policer type change will not trigger this issue, e.g., from term-specific policer to filter-specific policer.
 3. In the instances of the interface-specific filter created after policer parameter change, the policer might not work as expected.

As a workaround, after 2) and before 3), make any change on the interface-specific filters that contain the changed policer. [[PR819465](#): This issue has been resolved.]

- It is possible for the CLI command **show pfe statistics traffic** to not display any output and time out. `root@router > show pfe statistics traffic` error: the mib-process subsystem is not responding to management requests This is due to MIB2D process unable to query the statistics information maintained by PFED process due to transient errors on the socket between the 2 processes used for communication, despite both processes

running on the system. This fix addresses this issue. [[PR826086](#): This issue has been resolved.]

- The command **show interface aeXXX detail** might not have correct link information and some statistics shown as 0 when there is traffic. [[PR828155](#): This issue has been resolved.]

General Routing

- The output of the command 'show lldp neighbors' displays contents of Port-Description TLV instead of displaying contents of mandatory PortId TLV. Also the Port Description TLV gets generated using Port Name instead of using configured Port Description. [[PR550544](#): This issue has been resolved.]
- In l3vpn setup, customer facing interface fails to forward traffic, If RPF and localization is enabled in sequence. [[PR752540](#): This issue has been resolved.]
- In rare circumstances, events ranging from noise chatter on the XAUI interface to unwanted electrical signal transition on the PHY interface might cause the ports on the PD-5-10XGE-SFPP PIC to get stuck in the down or up state, effectively dropping all the traffic on that port. [[PR754344](#): This issue has been resolved.]
- In scaled VPLS scenario, with GRES and NSR enabled, with end-to-end traffic running, as part of the Routing Engine switchover on CE/PE router, when Packet Forwarding Engine receives request from the Routing Engine requesting whether the bunch of MAC addresses should be aged and if the route for one of the MAC address from the list (not last) is not found, then Packet Forwarding Engine replies to the Routing Engine with the subset of MAC addresses to be aged containing zero MAC address (for which route was not found) which can lead to vm core, leading to router going to db prompt. [[PR783099](#): This issue has been resolved.]
- After the Routing Engine switchover or a reboot, MLFR interfaces may not send MVPN data to all downstream receivers [[PR787168](#): This issue has been resolved.]
- Junos OS 11.4 introduced PTP feature and this in turn introduced Periodic reading of clocking chip registers over SPI interface. This is a known thing since day 1 of 11.4. With newer chassis (MX80/midRangius) program (MX80-T, MX40-T, MX10-T, MX5-T) H/w is reworked to implement SPI reads in h/w and it has significantly reduces clksyncd CPU usage to around 1-2% of CPU. [[PR789804](#): This issue has been resolved.]
- On an EX8200 Virtual Chassis or a switch with redundant Routing Engines, when you swap the members of a link aggregation group (LAG), a vmcore or ksyncd core file might be created on the backup Routing Engine. [[PR793778](#): This issue has been resolved.]
- This issue depends on FPC type. There are two types of FPC supported in a T4000 router, Enhanced Scaling FPC types, which are designated by the "-ES" suffix in their description and the T4000 FPC5 type FPC, which is designated by the "-3D" suffix in their description. The issue manifests in two different ways: a) Multicast traffic entering the router on an "-3D" FPC and leaving the router on an "-ES" FPC, will experience packet loss at specific packet sizes. b) If output sampling is performed on a service PIC located on an "-ES" FPC, where the ingress and egress FPCs of the traffic being sampled are distinct "-3D" FPCs, these samples can be discarded at specific packet

sizes. This means flow collectors will register fewer flow records than expected. For both of the above cases, not all packet sizes are affected. Packets less than 128 bytes in size are not affected, whilst packets above 128 bytes in size are affected at different packet size boundaries. For both of the above cases, messages similar to the following will be reported by the "-ES" FPCs and logged in the system messages file. Sep 26 14:36:47 routername fpc7 SRCHIP(0): 71024 Bad packets on p1 Sep 26 14:36:47 routername fpc7 SRCHIP(0): 71815 SONN errors on p1 Sep 26 14:36:47 routername fpc7 SRCHIP(1): 71056 Bad packets on p1 Sep 26 14:36:47 routername fpc7 SRCHIP(1): 71826 SONN errors on p1 [[PR794978](#): This issue has been resolved.]

- This PR provides fix for a Packet Forwarding Engine micro-kernel memory leak, applicable only to certain types of firewall filter configurations on MX Series router with MPCs and DPCs. The concerned leak can occur with dynamic-profiles configuration with firewall filters (typical for subscriber services), or when protocols like Ethernet OAM are configured. The leak occurs only under a certain timing condition (so far observed for about 2-5% of interface delete operations), and only on interface delete operations. Leaked memory is of order of 1KB per incidence of the memory leak. Hence the issue will be of material impact in and only in deployment scenarios which involve many thousands of interface delete operations with either dynamic profiles configuration or Ethernet OAM configuration. [[PR797790](#): This issue has been resolved.]
- On M and T Series platforms, in L3VPN scenario with l3vpn-composite-nexthop enabled, while PE receiving packet with DF bit set and packet size is larger than the core-facing interface's MTU, FPC/FEB may crash and core dump due to software defect of handling composite nexthop which need packet fragment. [[PR800155](#): This issue has been resolved.]
- After a software upgrade to Junos OS Release 11.4R3 or later, a "xe" interface may report "L2 channel errors" when issue "show interface extensive". We have seen this issue when the xe interface is part of an "ae" bundle. [[PR800634](#): This issue has been resolved.]
- Customer might encounter transient SLchip error when link flap event occurs. [[PR812092](#): This issue has been resolved.]
- The ATM MIC 2xOC12/8xOC3 CC-CE might show 0 pps values while executing "monitor interface at-x/x/x" even though there is input traffic on the interface's logical units. This not a service impacting issue. [[PR815632](#): This issue has been resolved.]
- VC-Subscriber license change not updated to all the Routing Engine's in back up members [[PR835039](#): This issue has been resolved.]

Infrastructure

- When the delete command is issued on an unnumbered Ethernet user route (static route with a qualified next hop), the destination route created as a part of this user route does not get deleted. This results in duplicate ARP entries for the same address. [[PR752163](#): This issue has been resolved.]
- Multicast traffic drops can occur if one of child links of ae bundle failed. [[PR779238](#): This issue has been resolved.]

- If we flap the interface used as next hop in the forwarding table for the IPv6 remote router loopback address used for IPv6 BGP sessions, the session flaps although there is another valid route over the other interface. [[PR791881](#): This issue has been resolved.]
- Filter Based Forwarding is not working for IPv6 traffic. [[PR795730](#): This issue has been resolved.]
- If the Routing Engine running in low memory environment, while receiving IP fragmented packets destined to the Routing Engine, system may fail to defragment the packets and will free the memory of these fragmented packets. But the router will try to free the already freed memory again which is incorrect, and in rare conditions, the freed memory may be allocated before the double free process occurs. Then an access to the freed memory will cause memory corruption and kernel crash. [[PR810434](#): This issue has been resolved.]
- Over time when routing-table instances is added and removed, when the global index allocated to a route-table reaches around 36730, the default route-table (inet.0), having index 0 always, could be wrongly allocated to a newly added instance. After this corruption to the default routing-instance whenever this default route-table (inet.0) is accessed or modified this could result in kernel panic. The fix addresses this corruption by adding the default route-table inet.0 (index 0) to the structure tracking all default tables in the system, thereby ensuring that this index never gets re-allocated and not resulting in this kernel core. The possibility of running into this issue depends solely on the rate of adding and removing routing-instances. [[PR829412](#): This issue has been resolved.]
- The logical interface inet6 protocol may be stuck at down state because of either external loopback or duplicate inet6 address detected. DAD will not run after this inet6 protocol-down event. [[PR834027](#): This issue has been resolved.]

Interfaces and Chassis

- After a MX Series Routing Engine reboot, the internal em interface may be "link down". The fix for this is to monitor the link state of the em interfaces. When an em link is discovered to be harddown, re-initialize the em interface. This fix has been limited to address only MX Series platforms. [[PR611081](#): This issue has been resolved.]
- On E1 interface, when interface flaps on CE side of connection interface will flap a second time on the PE side [[PR690403](#): This issue has been resolved.]
- Multiple inbound/outbound IPSEC tunnels seen for a single security association upon certain conditions such as router reboot. [[PR730174](#): This issue has been resolved.]
- Issue is seen only when the following steps are followed: 1. Enable IRB MAC Sync feature. 2. Deactivate BD/MAC Sync/Service ID on the higher MAC node. 3. Activate virtual switch that configures MCAE under the virtual switch. The result: IRB MAC Sync happens even though the feature is not enabled. [[PR793889](#): This issue has been resolved.]
- An operation in order of PIC offline then deactivate ci member interface will have the deactivated member interface join to the ci upon PIC online. And it will cause unexpected forwarding issue. [[PR803817](#): This issue has been resolved.]

- When a hold-time is configured for xe interfaces, before the timer is started, the XFP alarm is checked before declaring the link status. As a result, it is possible for the link to remain down if there are XFP alarms/warnings as reported by the "show interfaces diagnostics optics" command. [[PR804315](#): This issue has been resolved.]
- In subscriber management environment, with either 'pppoe-underlying-options' or 'advisory-options' or both of them configured for ethernet interface, after each login and logout of a DHCP/PPPOE subscriber, it has some memory leak in kernel. After some time, BRAS starts to reject new sessions and needs to be restarted to restore its operation. The memory leak occurs in the Routing Engine kernel can be observed by following commands:

```
user@router> show system virtual-memory | match iflog
iflogical 1471864 281335K - 17671778 16,32,64,256,4096,16384,32768,262144,524288
```

Or,

```
user@router> show chassis routing-engine
Routing Engine status: Slot 0: Current state Master
Election priority Backup Temperature 36 degrees C / 96 degrees F
CPU temperature 42 degrees C / 107 degrees F
DRAM 3584 MB Memory utilization 81 percent < < -----
```

[[PR807838](#): This issue has been resolved.]

- When a PEM module is inserted into a chassis, no message is currently logged in /var/log/inventory file. This PR addresses this requirement. On inserting a PEM module into a chassis a message in the following format will be logged in file /var/log/inventory : <Current-Date> PEM <pem-slot-number> - part number <part-number>, serial number <serial-number> eg. Oct 22 14:44:21 PEM 0 - part number 740-123456, serial number VK11111 [[PR808450](#): This issue has been resolved.]
- Once a GRES Mastership switch is performed and at some point later the FPC is rebooted the aggregate interface flag is set to 'down' once FPC comes back online. Any traffic which enters this FPC and is send to the aggregate interface via an ECMP path will get dropped. ECMP paths without aggregate members are not affected. This symptom is applicable on PTX Series platforms, T4000 systems with T4000-FPC5-3D and on MX Series platforms with FPCs from Junos OS Release 11.4R3 or later. To clear the condition, flapping the AE interface in 'Harddown' state such as 'deactivate' and 'activate interface ae' [[PR809383](#): This issue has been resolved.]
- Interface damping is not working properly on MPC Type 2 3D [[PR810159](#): This issue has been resolved.]
- If LACP is configured in distributed mode, which is default mode, LACP packets are not counted in input statistics of interface [[PR821874](#): This issue has been resolved.]
- When downgrading a Junos OS version, if a particular feature is not supported in the target Junos OS Version, then those unsupported features must be deactivated before downgrading the Junos OS version. [[PR836448](#): This issue has been resolved.]

Layer 2 Features

- If an vpls interface is moved to an vpls aggregate bundle or changed to any other interface family different than vpls, on the next GRES/NSR Mastership switch, this interface will have the CCC-Down flag set and will not process any traffic. The interface needs to be deactivated and activated via configuration change to continue forwarding traffic. [[PR788631](#): This issue has been resolved.]
- After a physical interface that is configured for CoS on a logical tunnel interface flaps, the MAC address of the peering logical interface goes missing from the kernel. To resolve this issue, deactivate/activate the peering logical interface after the flap. [[PR790559](#): This issue has been resolved.]
- For L2/vpls configs running on a trio based card, when the trio chipset encounters ppe traps (e.g. xtxn errors or ppe timeouts), in certain cases it is possible that internal microcode state related to mac learning does not get cleaned up correctly. This could cause mac learning to stop. [[PR798622](#): This issue has been resolved.]
- "monitor traffic interface irb" fails to capture outgoing packets. [[PR802605](#): This issue has been resolved.]
- The issue is due to another PR-686399, where the DA mac entries on AE are getting aged out though the traffic is ingressing on different Packet Forwarding Engines of the same AE. [[PR802924](#): This issue has been resolved.]
- Routing Engine CPU utilization may increase faster than expected when LDP neighbors are configured under a mesh group of a local BGP-VPLS routing instance when the LDP neighbors themselves are not provisioned for VPLS service on the remote side. [[PR808333](#): This issue has been resolved.]

Multiprotocol Label Switching (MPLS)

- Configuration changes to some attributes of the standby secondary path of an MPLS label-switched path (LSP) might cause the LSP to flap, which might result in packet loss. [[PR394184](#): This issue has been resolved.]
- In GRES (graceful Routing Engine switchover) mode and due to a quick status change of MPLS CCC nexthop, a mismatch of index value between master and backup Routing Engines might happen, causing Backup Routing Engine to panic, generate a core file, and trigger a live core dump from master Routing Engine. [[PR755473](#): This issue has been resolved.]
- If ldp-tunneling is configured on an LSP, and the router receives thousands of LDP routes through this LSP, and additionally the router has a large number of FPCs (especially in multi-chassis platform), during MPLS statistics collection, it is possible to see packet drops on the interface connect the Routing Engine with Packet Forwarding Engine. If auto-bandwidth is enabled for the LSP, then auto-bandwidth will work incorrectly. Below pfestats message type could be seen with high frequency (>800) when issue happens: /kernel: pfestat: receive req id 20286 type 21 reL_id 0 ipc type 6 subtype 30 uniq 200 peer 4 sendmask 0x46 And the packet drops could be seen by following command: user@router> show interfaces em0 extensive | match erro Input errors: Errors: 2083389, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:

0, Resource errors: 0 Output errors: Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0 [[PR785360](#): This issue has been resolved.]

- When an AE link along a path becomes oversubscribed due to failure of one or more member links of the AE link, the ingress of a LSP may continue to use the same path although there may be alternate path available with sufficient bandwidth. This is so because CSPF algorithm during optimization of an existing LSP will continue to see such an over-subscribed link acceptable with sufficient available bandwidth. However, if a new LSP with a bandwidth requirement is signaled over such a link, the LSP will not get signaled successfully. [[PR807670](#): This issue has been resolved.]
- With RSVP disabled, when a SNMP get/get-next is received for RSVP MIB, a Path State Block (PSB) search request is enqueued, this enqueue operation returns nothing but, the memory allocated for the search request is not freed and this results in a memory leak of routing protocols daemon (rpd). The memory leak could be observed by the following commands: user@router> show task memory detail | match "rsvp psb lookup req" ----- Allocator Memory Report ----- Name Size Alloc DTP Alloc Alloc MaxAlloc MaxAlloc Size Blocks Bytes Blocks Bytes RSVP PSB lookup req 176 180 T 110 19800 110 19800 user@router> show system processes extensive | match rpd PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND 1311 root 1 4 0 1529M 1479M kqread 75:25 0.44% rpd When the memory usage of rpd process increases to around 85% of system limit, the following logs could be seen: reO: /kernel: %KER-5:Process (1859,rpd) has exceeded 85% of RLIMIT_DATA: used 1835088 KB Max 2097152 KB [[PR811951](#): This issue has been resolved.]
- RSVP interface state will show as DOWN on unnumbered p2p interfaces. This will cause MPLS LSPs over those interfaces to stay down. [[PR814071](#): This issue has been resolved.]
- BFD session between PEs may not come up correctly after ppmmd crashes. [[PR826300](#): This issue has been resolved.]

Network Management and Monitoring

- There are compilation problems with following 3 MIBs: 1. mib-jnx-license 2. mib-jnx-sp-nat 3. mib-jnx-subscriber In 11.2 version of the JUNIPER-SMI, these three objects are defined, but they are missing in 12.1. Issue has been resolved in later releases. [[PR794327](#): This issue has been resolved.]
- 12.3-The security-name specified under 'target-parameters' should be used while sending v3 specific v1/v2 traps [[PR813430](#): This issue has been resolved.]
- This issue happens when a burst of traces is triggered by an event(commit in this case). If the max tracefile size is configured to a low value, then this leads to a race condition during log rotation. This issue can be prevented by increasing the max size of the traceoptions file. [[PR820322](#): This issue has been resolved.]
- The 'timestamp year msec' command in syslog (not using structured data) is intended to include year and msec details in local syslog messages stored in rotating files, but not to be included in messages sent to a remote syslog collector. This fix corrects a wrong behavior introduced in 11.3R1 where such details were also included in syslog messages sent to a remote host. [[PR820436](#): This issue has been resolved.]

Platform and Infrastructure

- Event scripts, part of Junos OS automation infrastructure, run at high priority compared to other system critical daemons. This can result in resource contention and high CPU readings. [[PR512315](#): This issue has been resolved.]
- In Trio chipset, the incoming interface index (iif) for family multiservice under enhanced-hash-key is being used by default. By default the iif-bridged should not be used, this is a software bug and is corrected in future releases [[PR701872](#): This issue has been resolved.]
- When Trio inline NAT translates a UDP pkt with UDP checksum 0x0000, it rewrites checksum to nonzero value. [[PR782927](#): This issue has been resolved.]
- Flow records obtained by using "inline-jflow" may contain incorrect AS value. [[PR788879](#): This issue has been resolved.]
- CLI cosmetic-errors are noticed while executing commands under "edit system class login" [[PR812022](#): This issue has been resolved.]
- HTTP-get probes fail when routing-instance is used for the probes. Removing routing-instance should work. [[PR814357](#): This issue has been resolved.]
- AIS scripts error "error: xsl:import : unable to load" [[PR815978](#): This issue has been resolved.]
- When two irb interfaces with the same layer 2 trunk interface within the bridge domain, multicast replication might be handled incorrectly. [[PR823435](#): This issue has been resolved.]
- The timestamp in the DDOS_VIOLATION report is not correct. The time the violation occurred is the time when the message appeared in the log. [[PR828085](#): This issue has been resolved.]
- When a Trio based card is receiving an IPv6 packet encapsulated in a MPLS stack formed by two explicit-null labels (0 ? ipv4 transport label, 2 ? ipv6 explicit-null label) is corrupting it. [[PR830209](#): This issue has been resolved.]
- With DHCP/BOOTP relay agent configured under [forwarding-options helpers] hierarchy, interface flapping will cause forwarding UDP daemon (fud) memory leak. The memory usage of fud process could be seen by following command: (SIZE: Total memory size of the process (text, data, and stack), in kilobytes) user@router> show system processes extensive | match "pid | fud" PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND 8436 root 2 0 2668K 1708K select 0:03 0.00% 0.00% fud user@router> show system processes extensive | match "pid | fud" PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND 8436 root 2 0 2676K 1716K select 0:03 0.00% 0.00% fud [[PR831965](#): This issue has been resolved.]

Routing Protocols

- The output of "show route summary" command might get delayed or timed out when the system is busy populating ~1M routes. [[PR433419](#): This issue has been resolved.]
- The routing protocol daemon (rpd) crashes if BGP traceoptions and 'precision-timer' are enabled. [[PR735580](#): This issue has been resolved.]

- If there are more than one peers in a BGP group, and have NSR configured, in a rare condition, BGP session init job fails on the last established peer in the group, while at the same time, there is another peer in establish-ack-wait state. Then after that, the peer in establish-ack-wait state becomes established after getting an establish-ack from the remote peer which will cause routing protocol daemon (rpd) crash and generates a core file. [[PR736198](#): This issue has been resolved.]
- Problem statement ***** rpd core at "rip_dc_retrans_callback_p2mp" while unconfiguring p2mp configuration. [[PR769487](#): This issue has been resolved.]
- In a scaled setup, with many routing-instances, if the FPCs are restarted, some pim encapsulation interfaces (pe) will be marked as down. [[PR770213](#): This issue has been resolved.]
- The current implementation of the random number generator grand() is not thread safe, as it utilizes a number of global variables and arrays. When bgp precision timers are enabled, this can cause crashes, e.g., when we jitter timers. The changes herein give each thread its own random number generator. [[PR777802](#): This issue has been resolved.]
- rpd cored during commit where NSR and GRES disabled but backup rpd is still running. [[PR794738](#): This issue has been resolved.]
- When using precision-timers, it is possible that BGP does not send all its advertisements from its buffer, until the BGP session needs to send another non-keepalive message. [[PR801037](#): This issue has been resolved.]
- Unicast RPF check not working properly after bouncing the RPF interface , when there are 2 BGP routes for same destination via 2 different next-hop RPF interfaces. [[PR814303](#): This issue has been resolved.]
- RPD in backup cored@rt_nexthops_free multiple times [[PR816754](#): This issue has been resolved.]
- In an l3vpn configuration, if no-vrf-propagate-ttl is configured under the vrf, rpd can core when deactivating or deleting the routing instance vrf with no-vrf-propagate-ttl. [[PR816851](#): This issue has been resolved.]
- Enabling "advertise-external" knob towards the Route-Reflectors or remote PEs can trigger RPD to core if: - "multipath vpn-unequal-cost equal-external-internal" enabled in a VRF - a multipath route exists with contributors from iBGP - an external path with lower preference (local-preference, as-path, etc) exists in the same VRF [[PR823844](#): This issue has been resolved.]
- During ISSU, LACP timeout may occur when time between fpc-upgrade and the Routing Engine switchover takes more than 90sec due to large scale config. [[PR826984](#): This issue has been resolved.]
- Changing the static route config from next-hop to qualified-next-hop will result in static route getting missed from routing table. Restarting routing process can bring back the routes but with RPD core. [[PR827727](#): This issue has been resolved.]
- Any protocol (MVPN,PIM,etc) which does install multicast nexthops and goes through a forwarding-table export policy with an install-nexthop policy action will fail to install

forwarding nexthop and multicast traffic is not forwarded. [[PR830448](#): This issue has been resolved.]

Services Applications

- MSDPC pic sending DPD triggers for tunnels that does not have IPSEC SAs. This result in kmd daemon on the Routing Engine continuously trying to send initiate IPSEC phase 1 negotiation since there are no active IKE/IPsec SAs. [[PR754461](#): This issue has been resolved.]
- Memory leak on FPC/FEB when service next-hops are deleted on LNS for l2tp sessions on non-MX Series platforms. For each l2tp session teardown, there are 240 bytes of memory leaks. After some time this leak will lead to FPC/FEB memory exhaustion leading to unpredictable FPC/FEB reset. [[PR770903](#): This issue has been resolved.]
- In a carrier-grade NAT (CGN) configuration that includes the "address-pooling paired" statement(APP feature) the service PIC or MS-DPC PIC might reset unexpectedly. This behavior occurs only in certain corner case scenarios in which the service PIC software times out in an APP mapping but a new flow is created within the same mapping. There is no workaround. The following logs could be seen before the crash:
pic offline req, pic 0, fpc 2, reason 0 CHASSISD_IFDEV_DETACH_PIC:
ifdev_detach_pic(2/0) LS: PIC DELETE: Deleting physical interfaces(#0) on pic 0 slot 2
ls_ifdev_detach ifd sp-2/0/0 marked as gone ifd gr-2/0/0 marked as gone
if_pseudo_detach_cleanup: Clearing physical interface present flag for gr-2/0/0 from ifd_pseudo_pim ignored. ifm_desc:83 [[PR800241](#): This issue has been resolved.]
- NAPT:port range start from 512,should start from 1024,this PR fixed this issue. [[PR804598](#): This issue has been resolved.]
- MS-DPC PIC crash in fwnat_natpool_release_port_from_pblock [[PR809139](#): This issue has been resolved.]
- When adding/deleting/changing the address ranges configured under NAT pool the Service PIC may run into an inconsistent state and restart [[PR810994](#): This issue has been resolved.]
- Please refer the AT, for the list of new commit constraint checks imposed. [[PR815053](#): This issue has been resolved.]
- Internally, each term is treated as a rule. Multiple copies of NAT pool is created one for each rule(or term). When address-allocation round-robin is configured, the NAT ip is read from separate copy of the NAT pool thus we see only 10 NAT IPs allocated for 20 different hosts matching two different terms. After the fix , all terms now can share same NAT pool. [[PR815147](#): This issue has been resolved.]
- Adaptive Service PIC (Service PIC II) is not supported by 12.1R4. Upon getting its configuration, the Service PIC reboots. A fix is planned in 12.1R4-S1. [[PR819833](#): This issue has been resolved.]
- In common NAPT44 translation scenario, without Address-Pooling Paired (APP) and Endpoint-Independent Mapping (EIM) configured, if NAT pool mapping for a private address is being deleted (e.g. mapping-timeout and inactivity-timeout expire), and at the same time there is a flow being created for the same private address, then service

PIC or MS-DPC might crash and core dumped in rare cases. [[PR821037](#): This issue has been resolved.]

- On an MX480 router, the FTP operations fail on a multiservices DPC. [[PR825355](#): This issue has been resolved.]
- Making SDG1 from standby to Master with all the flows already in sync will result in cli hang for sometime & flows will get cleared & re-sync in SDG1 during this cli hang timeframe. [[PR829950](#): This issue has been resolved.]
- Service PIC might crash in routing loops scenarios because of SIP ALG [[PR830070](#): This issue has been resolved.]

VPNs

- Junos OS 11.4x26.1 upgrade may fail if PIM MVPN instance does not have the vpn-group-address configured. [[PR753863](#): This issue has been resolved.]
- The issue is seen when multicast traffic is stopped and PIM states are allowed to clear. It was seen that some of (S, G) states on the PE did not clear. As a result data streams for the affected groups do not reach the intended receivers. However it was seen that the fix for this issue caused another issue to appear as documented in PR# 823884 [In NG-MVPN RPT-SPT mode, failure and then recovery of the primary path to the RP, causes the PE routers to forward traffic on (*, G) even though the (S, G) join is pruned] [[PR779786](#): This issue has been resolved.]
- When the egress PE receives type 4 route before discovering the ingress PE, it queues all the type 4 routes for later processing. Due to a defect in the corresponding code, the replaying of type-4 was not happening always, thereby causing the ingress PE to not send the corresponding stream to the egress PE. [[PR801437](#): This issue has been resolved.]
- With "vpn-apply-export", per-nexthop label and a "bgp policy overridden nexthop rather than the bgp local-address", labeled route for L3VPN didn't install in standby Routing Engine even with NSR configured. BGP on standby Routing Engine thinks it's not nexthop self and does not create labeled route on standby Routing Engine. This will cause label/nexthop change after switchover, therefore traffic loss. [[PR807285](#): This issue has been resolved.]
- For Static PW if user configures "static send-oam", PW will go down (state: Vc-Down) [[PR829666](#): This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R4

Class of Service (CoS)

- A memory leak in cosd can occur due to one of the cosd 8011p rewrite functions not releasing memory after use. This results in COSD memory running out of memory space, thereby resulting in continuous COSD core files. [[PR782728](#): This issue has been resolved.]
- When CoS configuration for chassis schedulers is applied to wildcard interfaces on a non-queuing DPC, an error should occur. But sometimes this configuration gets applied and causes issues such as tail-drops. [[PR784607](#): This issue has been resolved.]

- During addition/deletion or just deletion of an interface configured for a shared scheduler, some portion of memory is not reclaimed. Continuous addition/deletion of these interfaces results in memory depletion, causing packet loss and other issues. [PR803939: This issue has been resolved.]
- On the IQ2E PIC, if the logical interfaces under two different interface sets have their own scheduler map configured, when the same iflset shaping parameters are applied on these two interface sets, the iflset shaping profile is not shared between the interface sets. Hence, the number of iflset supported is reduced. There is no workaround. [PR804158: This issue has been resolved.]

Forwarding and Sampling

- This issue is caused by attempting to gather interface statistics simultaneously on too many interfaces or at too short an interval. Intervals of less than 5 minutes are not recommended. At this time, gathering statistics on more than 500 interfaces using "accounting-options interface-profile" can lead to inaccurate statistics output. [PR753960: This issue has been resolved.]
- In a heavily scaled setup when dfwd is busy processing the filter configuration, the ppm daemon would wait (approximately 2 minutes) for the firewall daemon to process the message it sent. The wait will happen for the nth message sent, since there is a limited buffer between the two daemons. [PR769452: This issue has been resolved.]
- DCU statistics can be broken on a sub-interface under certain circumstances when another sub-interface having DCU enabled gets modified (add/delete/change). If DCU statistics are broken, it could mean that the value gets corrupted, and continues to count from there properly or always remain in broken. The circumstances when this issue could be triggered are: - when 2 sub-interfaces have DCU enabled - when the last 16 bits of the 20 bit sub-interface index are the same - configuration change is made on one of these sub-interfaces - other sub-interface could have DCU statistics broken. keyword: sub-interface = logical interface = IFL. When this happens, we can see messages like the following being logged on the FPC syslog or messages log file. Aug 6 11:12:32.649 faraday-re1 fpc4 RT(rt_dest_class_stats_read): unable to read stats (ifl 196619, IPv4) - Not found Aug 6 11:12:32.723 faraday-re1 fpc1 RT(rt_dest_class_stats_read): unable to read stats (ifl 196619, IPv4) - Not found Aug 6 11:12:32.768 faraday-re1 fpc4 RT(rt_dest_class_stats_read): unable to read stats (ifl 196619, IPv4) - Not found Aug 6 11:12:32.843 faraday-re1 fpc1 RT(rt_dest_class_stats_read): unable to read stats (ifl 196619, IPv4) - Not found [PR794116: This issue has been resolved.]
- The message 'ifmon: Failed to find shmlog_dolog8' is logged when the **monitor interface <logical interface>** command is executed. [PR799849: This issue has been resolved.]

General Routing

- FPC crashes with the **show route ip table index # hardware statistics** command. [PR710756: This issue has been resolved.]
- With the configuration of the VPLS interface, the NH topology has a default filter-class being associated with the logical interface in the VPLS next-hop topology chain. During the PIC Offline/Interface deactivate event, the first message received from the kernel

by the Packet Forwarding Engine is to delete the VPLS family from the interface. This triggers a topology change wherein the filter-class gets deleted from the topology tree. During this topology change, the key buffer pointer adjustment fails thereby causing the increment in truncated key error counters for the Packet Forwarding Engine. The increment in these error counters triggers a chassis alarm associated with the R-chip error counters. This error counters stop incrementing once the interface completely goes down as a part of the PIC Offline/Interface deactivation event. [PR718591: This issue has been resolved.]

- When large CLI output is displayed, the CLI CPU might go high and remain high. [PR731647: This issue has been resolved.]
- A PCI read error can sometimes cause the 8xOC3-2xOC12 MIC bootup to fail and the MIC to stay offline with "Hardware Error" status. This fix addresses that problem. [PR733520: This issue has been resolved.]
- The crash is triggered by an AE link flap. The prerequisite for the crash is a forwarding topology with composite next hop pointing to unilist pointing to an aggregated interface (For example, VPLS pseudowire configured on MPLS LSP with node-link protection enabled). [PR746509: This issue has been resolved.]
- When configuring FIB localization, if a Packet Forwarding Engine is configured with FIB-remote, heap memory leak occurs on the Packet Forwarding Engine while installing and removing some prefixes using no-route-localize policy. [PR756787: This issue has been resolved.]
- DPC might randomly crash during ISSU. It will be kept offline after the ISSU period. [PR773960: This issue has been resolved.]
- On T Series Core routers and on TX Matrix and TX Matrix Plus routers, packets marked with the error flag, for example, due to DA (Destination Address) reject, are also counted as L3 incomplete, which is not correct and is misleading. Packets marked as errors from the PIC are now counted via the 'show pfe statistics error' command. [PR782070: This issue has been resolved.]
- On LMNR and Stoli FPCs, when a transit packet's (300 bytes or more packet size) ingress and egress interfaces are in the same Packet Forwarding Engine (with scaled egress NHs) and if a notification is sent to the Routing Engine for that packet, FPC might reset. The following list provides some scenarios where a notification is sent to the Routing Engine:
 1. IP options packet is received.
 2. TTL expired packet is received.
 3. Sampling is configured and a packet is sampled.[PR785143: This issue has been resolved.]
- In Junos OS Release 11.4R4, ETH-DM packets greater than 994 bytes fail the DMM test. The size of the default ETH-DM packets is much smaller, and bigger ETH-DM packets are used only with optional data payload size. For packet sizes less than 994, the functionality works fine. [PR790040: This issue has been resolved.]

- On an MX Series router, an error occurs while deleting "protect protocols l2circuit" from a Virtual Chassis configuration. [[PR794782](#): This issue has been resolved.]
- This issue depends on FPC type. There are two types of FPC supported in a T4000 router:
 1. Enhanced Scaling FPC types, which are designated by the "-ES" suffix in their description.
 2. The T4000 FPC5 type FPC, which is designated by the "-3D" suffix in their description.

The issue manifests in two different ways:

- a. Multicast traffic entering the router on a "-3D" FPC and leaving the router on an "-ES" FPC will experience packet loss at specific packet sizes.
- b. If output sampling is performed on a service PIC located on an "-ES" FPC, where the ingress and egress FPCs of the traffic being sampled are distinct "-3D" FPCs, these samples can be discarded at specific packet sizes. This means flow collectors will register fewer flow records than expected.

For both of the above cases, not all packet sizes are affected. Packets less than 128 bytes in size are not affected, while packets above 128 bytes in size are affected at different packet size boundaries. For both of the above cases, messages similar to the following will be reported by the "-ES" FPCs and logged in the system messages file.
 Sep 26 14:36:47 routename fpc7 SRCHIP(0): 71024 Bad packets on p1
 Sep 26 14:36:47 routename fpc7 SRCHIP(0): 71815 SONN errors on p1
 Sep 26 14:36:47 routename fpc7 SRCHIP(1): 71056 Bad packets on p1
 Sep 26 14:36:47 routename fpc7 SRCHIP(1): 71826 SONN errors on p1 [[PR794978](#): This issue has been resolved.]

- It has been observed in the test lab that when there is a prefix and pointing to a multipath BGP next hop (8 in number), and in turn each of this next-hops are pointing to multiple MPLS LSPs, the convergence number for 450 k routes were in the order of 15 to 20 minutes. It is also observed that any change in next hop, such as the interface flapping or neighbor flapping had a significant impact on convergence time computing the convergence for this 450 k prefix routes * 64 nexthop (8 paths for each of next hop). [[PR798771](#): This issue has been resolved.]
- When there are multiple recursive routes are available for a prefix such as BGP route pointing to a indirect (8) and in turn pointing to a unilist (32) and when this is going over an AE of 8 links, we saw the software going through a large computation for each prefix and this makes it worse proportional to number of prefixes. At this instance, the system crashes because of the large computation. [[PR800157](#): This issue has been resolved.]
- [MPC] Non-QX cards do not transmit PPP hellos on wire. [[PR801565](#): This issue has been resolved.]
- If MPC3 is equipped with a 10x10GE MIC or 2x40GE MIC in the MIC slot 0 and 20x1GE MIC in the MIC slot 1, the links will not come up for MIC in MIC slot 0. [[PR803613](#): This issue has been resolved.]
- The KSYNCD core followed by kernel live core is observed very rarely after the Routing Engine switchover. This issue can be detected when a ksyncd core is observed along

with the following log message in /var/log/messages. "Aug 27 01:28:03 indiranagar1 ksyncd[2506]: KSYNCD: resync error, issu_state[0], type Generic config subtype 8 : File exists". As a workaround, reboot the backup Routing Engine. [[PR810787](#): This issue has been resolved.]

Infrastructure

- The top utility with the "ores" options was not sorting the output based on resident memory size. This has been now fixed. [[PR507675](#): This issue has been resolved.]
- Due to a defect introduced earlier this year, a timing issue in the ttymodem() internal I/O processing routine can cause the Junos OS kernel to crash. The crash can be triggered by simple remote access (eg. telnet, SSH) to the router. [[PR755448](#): This issue has been resolved.]
- If we flap the interface used as the next hop in the forwarding table for the IPv6 remote router loopback address used for IPv6 BGP sessions, the session flaps although there is another valid route over the other interface. [[PR791881](#): This issue has been resolved.]
- In an IPv6 scenario, when "ipv6-duplicate-addr-detection-transmits" is configured with a value of zero, IPv6 Neighbor Discovery might not function properly. [[PR805837](#): This issue has been resolved.]
- Router might reboot while running 'show system core-dump core-file-info' command. This command uses /tmp and while uncompressing the core file, the tmp file system might be exhausted. /tmp in turn uses a swap device only. MFS (Memory File System) and the rest of the OS share the same swap space. Consuming more swap spaces might lead to out of memory and a swap situation, which could eventually bring down the system. [[PR808243](#): This issue has been resolved.]

Interfaces and Chassis

- After a MS-DPC core file, it might result in ICHIP "stream blocked detected". Traffic flow will be dropped and can only be restored by restarting the DPC. This is due to an SG FPGA soft reset issue. [[PR743262](#): This issue has been resolved.]
- When aggregated Ethernet next hops in discard state are replicated to the backup Routing Engine, the backup Routing Engine might crash and generate core files. In a multichassis system, both (SFC backup and LCC master work as a backup Routing Engine) the core files are generated on both SFC and LCC chassis. [[PR748436](#): This issue has been resolved.]
- There can be a mismatch between the ifIndex value on IF-MIB-ifName and the ifIndex value on SONET-APS-MIB-apsMapGroupName and apsMapEntry. [[PR771877](#): This issue has been resolved.]
- Junos OS Release 11.4X27 does not support unified in-service software upgrades (unified ISSU) for configurations that include interface sets. [PR-779377] [[PR779377](#): This issue has been resolved.]
- Load average values collected via SNMP do not show the correct values of the other Routing Engine. This can be verified by using the following commands: show snmp mib walk jnxOperatingEntry | match LoadAvg.9.1.0.0 show snmp mib walk jnxOperatingEntry | match LoadAvg.9.2.0.0 [[PR782817](#): This issue has been resolved.]

- Local SNMP walk for VRRP MIBs might loop continuously when the IRB interface is deleted. [[PR785582](#): This issue has been resolved.]
- The prefer-status-control-active configuration knob at the [edit interfaces aeX aggregated-ether-options mc-ae events iccp-peer-down] hierarchy requires the configuration knob to be active at the [edit interfaces aex aggregated-ether-options mc-ae status-control] hierarchy. When this is not present prefer-status-control-active has no impact and its presence in the configuration knob incorrectly implies that the current node is preferred active. [[PR785930](#): This issue has been resolved.]
- In the environment of composite-next-hop with FMBB (fast make-before-break) function (for example, VPLS, multicast, P2MP LSP), if the system is configured with a feature that contains an interface having active/standby links (for example, RLSQ/AMS/RMS), the Packet Forwarding Engine having a standby interface might be incorrectly involved when building the Packet Forwarding Engine flooding tree. This results in traffic blackhole. [[PR786007](#): This issue has been resolved.]
- 'monitor ethernet delay-measurement' command does not time out when CFM adjacency is down and/or all DMM frames are sent. As a result, ethdm binary does not close normally leading to an increase in resource consumption. [[PR787985](#): This issue has been resolved.]
- FPC might crash due to page fault. [[PR791195](#): This issue has been resolved.]
- T640 frame relay interface status is shown as up/up with mismatched lmi-type. [[PR791501](#): This issue has been resolved.]
- Configuring fxp0 with Speed 10m and Full-Duplex generates log message "fxp0: Full duplex link mode is not supported with speed 10M, Hence Speed will default to 100M" [[PR791777](#): This issue has been resolved.]
- **show chassis hardware** output for some optics is not correct sometimes. [[PR792704](#): This issue has been resolved.]
- On MX Series routers and PTX Series switches, a change to the 'oam lfm pdu holdtime' on an interface is not updated correctly. This results in an incorrect LFM state, which should be reported as Adjacency Lost. As a workaround, issue the clear oam ethernet link-fault-management state command from the CLI to correctly update the 'pdu holdtimer.' [[PR792763](#): This issue has been resolved.]
- When you configure the untagged GE interface or untagged aggregate Ethernet with a link member on the IQ PIC with a per-unit-scheduler, this might cause the failure of interface statistics on this interface. As a result an error will be reported in the log message and on "show interface extensive" outputs. [[PR794975](#): This issue has been resolved.]
- When you configure deterministic port block allocation preserve-range settings and traffic is first sent for preserve-range ports (1023 or below) and then the non-privileged regular ports (1024-65535), packets are discarded at slow path and drop flows are created. [[PR795609](#): This issue has been resolved.]

- When upgrading to 11.4R4, links that are using tunable DWDM XFP are not working anymore and are reporting a different wavelength than the configured one. [[PR796330](#): This issue has been resolved.]
- In 11.4R4, 12.1R3, 12.2R1, and later, the option of 'routing-engine' under "> request system snapshot" was mistakenly removed. [[PR809321](#): This issue has been resolved.]

Layer 2 Features

- With VPLS using IRB, the line card might crash when large changes to interfaces and nexthops are processed. It is not determined which conditions will trigger the crash. [[PR752378](#): This issue has been resolved.]
- When upgrading from 12.1R2.9 to 12.1R3.5, the validation might fail due to some configuration under class-of-service forwarding-policy. [[PR807365](#): This issue has been resolved.]

Layer 2 Ethernet Services

- LACP status disagreement after the Routing Engine switchover. [[PR751745](#): This issue has been resolved.]
- The AFTR information from RADIUS is advertised by MX Series router to the client via DHCPv6. [[PR779679](#): This issue has been resolved.]
- Old dhcp session still can be renewed by the client after the client moves to another vlan. [[PR784951](#): This issue has been resolved.]
- In 11.4R2, 12.1R1, and 12.2R1 and then subsequent builds on those releases, there might be a false alarm of a hardware problem from a DPC. For example: "fpc0 EZ: %PFE-3: ezchip_periodic_check_free_rfd_buffer[4245] XETH(0/3) : Rx RFD buffers exhausted" This can be ignored, unless traffic impact is seen. [[PR796824](#): This issue has been resolved.]
- DHCP relay does not forward ACK to the client from the backup DHCP server after primary DHCP server failure. [[PR799090](#): This issue has been resolved.]

Multiprotocol Label Switching (MPLS)

- The issue appears to happen when MVPN tries to add a vt-interface for an egress tunnel. RSVP would try to find the flood next hop for the route installed for the label for the branch LSPs in the P2MP LSP to add the vt-interface. When RSVP does not find the flood next hop for the label route for a branch LSP, it triggers an assertion failure. [[PR770538](#): This issue has been resolved.]
- A commit for configuration change that simultaneously disables RSVP and a point-to-point interface (like so, t1, atm) might generate an rpd core file. To solve this issue, do not commit a configuration change that simultaneously disables RSVP and a point-to-point interface. Rather disable RSVP and point-to-point interfaces in separate configuration commits. [[PR782174](#): This issue has been resolved.]
- With the fix of this PR, at the end of the adjust-interval operation, the max_average counter does not reset to zero and maintains the old value until it receives a new sample. When the new sample is received, the max_average counter is updated to the new

sample value. This is just a display counter fix and there is no operational impact. The auto-bandwidth functionality works as it is. [[PR799155](#): This issue has been resolved.]

- Point-to-multipoint inclusive tunnels over MVPN might not come up because the RSVP state for the "vt" logical interface appears as down. [[PR802344](#): This issue has been resolved.]
- Some implementations of cSPF allow zero-bw LSPs (LSPs, which are requesting bandwidth of 0bps) to be calculated via links which have 0 bps of available bandwidth. In some cases implementation of RSVP in Junos OS allows AvailBW on link to become negative. This might happen, for example, in case of failure of one of links in ae bundle. As it's impossible to include negative values in OpaqLSA, in this case Junos OS announces 0 bps of AvailableBW on such links. Zero-bw LSPs will not be established via such link, because transit node with negative AvailBW fails check for bandwidth availability on egress interface. HeadEnd will constantly try to signal LSP and every time it will receive PathErr: BW Unavailable from node which has link with negative AvailBW. cSPF calculation will not resolve this, because according to TED on HEnd we have 0 bps of AvailBW on this link, not negative. [[PR802995](#): This issue has been resolved.]

Network Management and Monitoring

- After a Routing Engine switchover, LACP and MIB process (mib2d) core files might be created. [[PR790966](#): This issue has been resolved.]

Platform and Infrastructure

- You might see zombie processes increment while doing commit each time. [[PR692382](#): This issue has been resolved.]
- Under very special rare conditions, the MPC CPU might stop processing and will be reset due to the Level3/Level 2 watchdog expiration timer. Potential exposure is a high load of traffic send to the Host. The following syslog message will be reported in the syslog once MPC reboots. "fpc[x] MPC: Reset reason (0xc): Level3 watchdog, Level2 watchdog" [[PR717899](#): This issue has been resolved.]
- PTSP and AACL services do not work with AMS interfaces. [[PR727588](#): This issue has been resolved.]
- In a scenario where the telnet session is disconnected ungracefully while accessing "load merge terminal" prompt problem can be exhibited with other CLI users unable to access configuration mode. [[PR745280](#): This issue has been resolved.]
- Memory exhaustion on the Packet Forwarding Engine ukern heap causing FPC to generate a core file. [[PR777609](#): This issue has been resolved.]
- Need to add 'start-time', 'stop-time' and 'timezone' attributes in TACACS+ accounting packets. Solution: ----- Added a configuration knob to allow enabling these attributes to be framed into the accounting packets. Default behavior is to not include these attributes in the packet to maintain backward compatibility. The new knob is [system tacplus-options timestamp-and-timezone]. You will have to enable this knob and then check for the 'start-time', 'stop-time' and 'timezone' attributes in the accounting packets. [[PR780484](#): This issue has been resolved.]

- When reconfiguring an interface from a native VLAN to another tagged VLAN, the logical interface mapping on the Packet Forwarding Engine might get corrupted. In case traffic is being received on this interface, it can lead to LU congestion and wedge. [[PR792633](#): This issue has been resolved.]
- Committing a Q-in-Q configuration results in an FPC crash in these conditions: 1. Core facing interface is configured this way: flexible-vlan-tagging; encapsulation flexible-ethernet-services; unit xxx { encapsulation vlan-bridge; vlan-tags outer xxx inner-range 1-4094; } 2. Core-facing interface is an aggregate interface. 3. Core-facing interface is on MPC card. [[PR793429](#): This issue has been resolved.]
- On MX Series routers with MPCs/MICs (in releases 11.4R4+, 12.1R3+, 12.2R1+), when the first large sized (>1500) transit packet hits a resolve route on the Packet Forwarding Engine it might cause memory to leak on the Packet Forwarding Engine micro kernel. [[PR802051](#): This issue has been resolved.]
- In 11.2 Layer3 services over MC-LAG are supported through IRB only, and thus family inet is not supported directly on MC-LAG interface. However, an appropriate commit check is missing in 11.2R3 and later maintenance releases of 11.2. [[PR802938](#): This issue has been resolved.]
- With inline sampling, when there are multiple flow servers being configured or multiple equal cost paths exist for a single collector, the flow record packet might trigger the following trap message from the Packet Forwarding Engine which causes a drop for the flow record packet. PPE Sync XTXN Err Trap: Count 1659, PC 45f, 0x045f: balanced_multi_nh_use_cp_index [[PR805061](#): This issue has been resolved.]
- IPv6 traffic on MX Series routers with MPCs/MICs might cause IPv4 SCU/DCU counters to increment. [[PR805257](#): This issue has been resolved.]

Routing Protocols

- Dynamically signaled routes might flap when a "commit full" is performed. [[PR672838](#): This issue has been resolved.]
- The rpd might generate a core file after deleting or renaming a non-forwarding instance. Specifically, issue occurs when: 1. the interface is configured in a non-forwarding instances (that is, routing instances xxx with no instance-type). and any one of the following: 2a. igmp is configured with all interfaces (For example, protocols igmp interface all) 2b. igmp is configured on the specific interface (For example, protocols igmp interface ge-0/0/0.1) 2c. mld is configured with all interfaces (For example, protocols mld interface all) 2d. mld is configured on the specific interface (For example, protocols mld interface ge-0/0/0.1) 2e. pim is configured in the master instance with all interfaces (For example, protocols pim interface all) 2f. pim is configured in the master instance on the specific interface (For example, protocols pim interface ge-0/0/0.1) and any one of the following actions are subsequently committed: 3a. delete the non-forwarding instance (For example, delete routing-instances xxx) 3b. rename the non-forwarding instance (For example, replace pattern xxx with yyy) [[PR704699](#): This issue has been resolved.]
- If the network is configured for PIM nonstop active routing (NSR), a core file might be triggered in an upstream router because of high churn in unicast routes or a continuous

clearing of PIM join-distribution in the downstream router. To prevent this possibility, disable PIM-NSR. [[PR707900](#): This issue has been resolved.]

- The rpd process is reinitialized when you commit a configuration change. When multiple reinitializations occur while OSPF is running on the router, the periodic refresh of OSPF router LSAs might stop. If the LSAs are not refreshed, the router no longer participates in the OSPF routing domain. You can issue the **show ospf database router advertising-router router-id extensive | match timer** command to see evidence of the issue. In the error state, the output does not include the Gen timer field. [[PR744280](#): This issue has been resolved.]
- If a routing instance is configured to add static routes to its instance-specific routing table using both the routing-options static route stanza and the routing-options rib <instance specific table name> static route stanza and a configuration event changes something else in the routing-options rib <instance specific table name> stanza such as modifying the maximum-paths value, the static routes in the instance table specific section can be deleted. A commit full can be used to recover or only use one of the two mechanisms for defining the static routes for that instance. [[PR755558](#): This issue has been resolved.]
- The rpd process crashes when the customer-specific BGP configuration is unconfigured. [[PR782816](#): This issue has been resolved.]
- During extended stress testing of the PIM protocol, a malformed PIM Hello message triggered an rpd crash. While the crash was caused by a malformed PIM message, simply replaying the crafted packet alone does not lead to the crash. This issue affects both IPv4 PIM and IPv6 PIM. Refer to PSN-2012-10-732 for additional information. [[PR792334](#): This issue has been resolved.]
- An MX Series router that has only some bridge-domains configured for igmp-snooping might discard traffic in bridge-domains without igmp-snooping enabled. [[PR795781](#): This issue has been resolved.]
- The rpd process crashes after making changes to policy-statement related to VPN. [[PR807357](#): This issue has been resolved.]
- The following log messages are generated when a commit is performed.
"task_set_option_internal: task ICMP socket 103 option GroupAdd(23) interface ae12.0: Address already in use." The error leading to this log is handled properly. Since these logs can be misleading, the logs have been suppressed with the fix. [[PR809472](#): This issue has been resolved.]

Services Applications

- When a TX Series router is configured with a manual OSPF ipsec-sa for authentication, something like the following cosmetic messages will be logged: Feb 16 16:27:40 flame-sfc-re1 lcc0-master kmd[17194]: KMD_RTsock_ERROR: Error adding inbound SA OSPF3_AH_SHA1_96 spi=1024 proto=AH to kernel: No such file or directory Feb 16 16:27:40 flame-sfc-re1 lcc0-master kmd[17194]: KMD_RTsock_ERROR: Error adding outbound SA OSPF3_AH_SHA1_96 spi=1024 proto=AH to kernel: No such file or directory If there's a service PIC, there will be these additional cosmetic log entries: Feb 16 16:27:46 flame-sfc-re1 lcc1-master kmd[16853]: KMD_INTERNAL_ERROR: Failed to connect PIC, ERR: Failed to connect PIC, ERR: F Feb 16 16:27:46 flame-sfc-re1 lcc1-master kmd[16853]: KMD_INTERNAL_ERROR: Unable to connect PIC sp-8/3/0; Feb 16 16:27:46 flame-sfc-re1 lcc1-master kmd[16853]: KMD_INTERNAL_ERROR: Couldn't request PIC: sp-8/3/0 to send sa state [[PR738736](#): This issue has been resolved.]
- This crash was observed during a mixed traffic test for 7 hours on below traffic profile. Traffic profile used HTTP 0.8m HTTPS 0.15m FTP 0.1m RTSP 0.08m UDP 8.87m (IMIX traffic) [[PR769322](#): This issue has been resolved.]
- MS-PIC might crash when releasing port block for a flow with SIP ALG enabled. SIP flows that can trigger this are the ones that have SDP media address as 127.0.0.1. [[PR774589](#): This issue has been resolved.]
- In the deterministic NAT block-size 0 case, the block size per user is automatically calculated by the following formula:
$$(64512 * \text{Total number of IP addresses in the NAT pool}) / \text{Total number of subscribers}$$
where 64512 is derived from (65535 - 1023), because the assignment of regular ports begins at 1024. [[PR776247](#): This issue has been resolved.]
- RTSP streaming is not working in a laptop when moving or fast forwarding the video. [[PR786085](#): This issue has been resolved.]
- MX LNS does not support CLI command services l2tp session user filter option. [[PR792239](#): This issue has been resolved.]
- The clear services l2tp session user <> command accepts any arbitrary alphanumeric characters in place of a user name, and the CLI command will drop all l2tp subscribers. [[PR792631](#): This issue has been resolved.]
- [Deterministic NAT]ports allocation overlapped [[PR797457](#): This issue has been resolved.]
- The kmd process is running high with key chain configuration. Need to modify the necessary changes so that kmd does not run wherever is it not required. [[PR798030](#): This issue has been resolved.]
- deNAT:wrong mapping between nat-port-block and internal-host. [[PR799947](#): This issue has been resolved.]
- Certain changes to NAT (Network Address Translation) PBA (Port Block Allocation) configuration require a reboot of the services PIC in order for the changes to take effect on the PIC. A warning should be issued during the commit informing you of the need to reboot. Because the warning is not currently being issued, the documentation should include the following warnings: For the PBA or Deterministic NAT configuration changes

to get effect on the service PIC, you must reboot the Service PIC when the following NAT configuration changes are made under the hierarchy:

- [edit services nat pool]
- port range
- address/address-range
- block-size
- max-blocks-per-address
- active-block-timeout

[[PR807350](#) . This issue has been resolved.]

VPNs

- An optimization has been implemented with BGP-MVPN next hop infrastructure that will improve scalability in some multi-dimensional scaling scenarios with aggregate interfaces. [[PR690690](#): This issue has been resolved.]
- When you disable protocols in a Layer 2 circuit with egress protection, rpd generates a core file if no routes are found in context routing table. [[PR735789](#): This issue has been resolved.]
- UMH selection should select the highest IP address as the Upstream PE. However, in the code the highest IP address is selected by comparing the lowest order byte of the IP address first. In this case between IP address 10.233.38.34 and IP address 10.233.32.46 - 10.233.32.46 gets chosen as the upstream PE because its lowest order byte (46) is more than the lowest order byte of 10.233.38.34. This is because code does not account for the endianness of the machine it is run on. Fix - convert the IP address to network order before comparing. [[PR754114](#): This issue has been resolved.]
- In NG-MVPN, if PE X loses connectivity to source PE Z and learns a route via PE Y to reach PE Z, it sends a type 7 source tree join route to PE Y. When the topology restores, PE X relearns the original route to the source PE Z and withdraws the type 7 route from PE Y. While withdrawing the route, the routing protocol process (rpd) on PE Y will crash because the selective PMSI is bound to PE Z, not PE X. [[PR755077](#): This issue has been resolved.]
- In L2VPN scenario with at least one L2VPN connection in up state, during SNMP walk on jnxVpnPwAssociatedInterface.bgpL2Vpn (OID: .1.3.6.1.4.1.2636.3.26.1.4.1.6.3.9), an infinite loop occurs due to software defect and causes routing protocol daemon (rpd) to stuck at 97% indefinitely until the rpd process is restarted. While in this state all protocol adjacency will expire and the router will stop forwarding traffic. [[PR782654](#): This issue has been resolved.]
- In scaling scenarios, buffer overwrites might occur when there are multiple sources for the same group. This might lead rpd to generate a core file when MVPN traceoptions are enabled. [[PR783615](#): This issue has been resolved.]
- When the Source PE is rebooted, S,G,RPT state is not cleaned up. The change is to clean up this S,G RPT state if we learn the source address is remote. [[PR784627](#): This issue has been resolved.]

- In Rosen MVPN scenario, after performing Routing Engine switchover, some of the Rosen VRFs don't have a tunnel interface (mt-) assigned for incoming, and therefore will not join the transport group sourced with the remote PE loopback address. [[PR791333](#): This issue has been resolved.]
- The rpd incorrectly sets the PWE3 Control Word flag for local switching circuits. PWE3 Control Word is needed for Layer 2 VPN OAM packets to get pushed to the Routing Engine. On MX Series routers with MPCs/MICs, the traffic payload is examined and if it matches the first nibble being 0001, it sends the traffic to the Routing Engine for further processing. Once Layer 2 VPN traffic is set to IPv4, it would test against the IPV4 ID field. [[PR793751](#): This issue has been resolved.]
- This issue was experienced in Junos OS 11.4R4 code with NG-MVPN RPT-SPT mode. When a link failure causes the route to the source and RP via the backup path, the PE in the backup path fails to forward multicast traffic to the receiver. This is documented in PR 794222 and upgrading to Junos OS 11.4R4-S2 fixes the above mentioned issue. [[PR794222](#): This issue has been resolved.]
- This change would allow customers to use the less restrictive CLI knob 'vrf-advertise-selective', which now accepts a null list. If no family is configured under vrf-advertise-selective, then no MVPN routes are advertised to the neighbor. [[PR795108](#): This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R3

- [Class of Service \(CoS\)](#)
- [Forwarding and Sampling](#)
- [General Routing](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Layer 2 Ethernet Services](#)
- [Multiprotocol Label Switching \(MPLS\)](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Services Applications](#)
- [User Interface and Configuration](#)
- [VPNs](#)

Class of Service (CoS)

- **Show class-of-service classifier name** "classifier name with spaces" does not work for classifiers that have spaces in their name. [[PR535967](#): This issue has been resolved.]
- The following configuration: "set class-of-service interfaces all unit * classifiers exp <name>", causes the exp classifier exp-default to be attached to every Label Switched Interface (LSI) of the router. This is regardless of the <name> specified, and also regardless of the configuration at [edit class-of-service routing-instances] level. [[PR710427](#): This issue has been resolved.]
- RED drop counters are not displayed correctly in the **show interfaces queue** command. [[PR735610](#): This issue has been resolved.]
- Packet Forwarding Engines configured with egress-shaping-overhead value of -14 reports 128 byte packet as a 248 byte packet. There is no work around for this issue. [[PR775454](#): This issue has been resolved.]
- The in-service-software-upgrade to Junos OS Release 12.1R2 might fail on Queuing DPCs in Mx Series platform, when IFLSETS are defined and the IFLSETS have traffic control profile attached. The traffic flow will be as per configuration. However fetching of queue statistics after ISSU for few logical interfaces might fail due to incorrect software state of the system post ISSU. The issue is not seen for the logical interfaces which are member of an IFLSET with traffic-control-profile attached, but there is atleast one member logical interface in the same IFLSET, that does not have a traffic-control-profile attached. This statistics problem is seen only for a few logical interfaces. [[PR779375](#): This issue has been resolved.]

Forwarding and Sampling

- When the configuration archiving FTP process stalls during file transfer, it can result in the PFED process stalling as well. Once the master PFED process is restarted, it results in the inability to commit certain new configuration changes. Ensuring that the configuration archiving and FTP server are correctly configured and working will avoid this problem. [[PR528653](#): This issue has been resolved.]
- Sampled memory increases when interfaces bounce and BGP is running. [[PR594509](#): This issue has been resolved.]
- The Junos OS firewall process (dfwd) may generate core files and restart during a unified in-service software upgrade when the configuration includes Ascend-Data-filters. [[PR746128](#): This issue has been resolved.]
- The customer might see the following firewall logs while running VRRPv3 over VPLS:
04:39:14 pfe A irb.111 VRRP fe80::2 ff02::e42a:1:e:a0c 04:39:13 pfe A irb.111 ICMPv6 fe80::2 ff02::1 04:39:13 pfe A irb.111 ICMPv6 fe80::1 ff02::8:0:1d:44f2 [[PR748826](#): This issue has been resolved.]
- On ADPC cards Output Layer-2 policer drops the packets when configured on a interface with vpls encapsulation. [[PR749141](#): This issue has been resolved.]
- Any change to the last member of a service-filter chain, can lead to the loss of layer 3 connectivity over the interface. [[PR750957](#): This issue has been resolved.]

- In a heavily scaled setup when dfwd is busy processing the filter configuration, pcmd daemon would wait (approx 2 min) for firewall daemon to process the message it sent. The wait will happen for the nth message sent, since there is a limited buffer b/w the two daemons. [[PR769452](#): This issue has been resolved.]
- This issue can occur for daemons that connect to pfed for statistics information. If the daemon starts before pfed or has problems making a connection to pfed then that daemon could experience a crash. This can also occur using cli commands like **show interface statistics** that invoke ifinfo. This problem was introduced by the fix for [PR743135](#) and only exists in the specific versions that were fixed by that PR. There is no known workaround for this problem. However, you may use a Junos OS Release with the fix. [[PR770766](#): This issue has been resolved.]
- This PR eliminates an erroneous error message which would appear in syslog while pfed was checking the syntax of a configuration containing firewall filters that reference accounting counters. This problem only affected the syntax check prior to the commit. The actual configuration on the router was correctly committed. [[PR772463](#): This issue has been resolved.]
- The system archival feature configured that configuration is backup at an archival site periodically. This may leave behind files in /var/tmp when the connection to remote site fails. [[PR778962](#): This issue has been resolved.]
- When IPv4 and IPv6 classifiers are configured on a logical interface that already has a Layer 2 hierarchical policer configured on its physical interface, the rate of the policer deviates by more than 5 percent. To prevent this issue, attach the Layer2 policer and the classifiers in a single commit. If the physical interface already has a policer configured, deactivate it first. Then in a single commit, attach the classifiers and reactivate the policer. [[PR779357](#): This issue has been resolved.]
- Sampled process might core when both Origin AS and Peer AS are configured for the Routing Engine-based sampling and when there is a **commit** command issued to apply the configuration changes. [[PR779620](#): This issue has been resolved.]
- IPFIX issue in Junos OS Release 11.4R2 on MP4C-3D-16XGE-SFPP- Unable to see IPFIX (IPv4) statistics while issuing the **show** command **show services accounting flow inline-jflow fpc-slot 0**. [[PR787487](#): This issue has been resolved.]
- **show firewall detail** command will not display all the firewall policer counters if the "enhanced-policer" chassis knob is set and if the configured filters contain more than 10 policer counters. [[PR789889](#): This issue has been resolved.]

General Routing

- Once a child interface of an aggregate bundle is in down state (for example: CCC-Down of the logical member link interfaces), the next-hop of the control channel is not correctly programmed. LACP packets received are not dropped but processed and pointing to invalid NH entries, which might yield to such errors as below or a combination of all: fpc5 LUCHIP(0) IDMEM[0x000433ba] Read Uninitialized Memory Error fpc5 LUCHIP(0) PLCT INT_STAT 0x00000001 Illegal PL Uninitialized EDMEM Read 0x6db6db6d6db6db6d @ 0x1cf30001 XTXN 0xa8cd87 BULK 0x005c0094 FN 0 sync PPE 14 CNTX 1 fpc5 LUCHIP(0) RMC 2 Uninitialized EDMEM[0x1001c0] Read (0x6db6db6d6db6db6d) fpc5 LUCHIP(0) PPE_6 Errors sync xtxn error thread timeout error fpc5 PPE Sync DMEM WP Trap: Count 103, PC 620b, 0x620b: nat46_loop 0x620b: nat44_loop fpc5 PPE Sync XTXN Err Trap: Count 980053, PC 2f9, 0x02f9: nh_ret_simple_last fpc5 PPE Thread Timeout Trap: Count 2840, PC 4c6, 0x04c6: set_lif_inc_ifl_cnt fpc5 PPE PPE Stack Err Trap: Count 20347, PC 310, 0x0310: add_default_layer1_overhead fpc5 PPE PPE HW Fault Trap: Count 529, PC 373, 0x0373: inner_rewrite There is no operational impact other than the filling up of error messages in the system log. [[PR703245](#): This issue has been resolved.]
- *IF* you have DCU statistics configured in conjunction with the copy-plp class of service knob & a output firewall filter you may encounter a situation where DCU stats are no longer working. Check first that both ingress & egress ports for the flows you are counting are *NOT* on different Packet Forwarding Engines the same fpc. *IF* this is the case then removing the copy-plp knob will restart the DCU statistics collection. [[PR707834](#): This issue has been resolved.]
- Provided that an MX Series router has only one interface that uses auto-configure to set up a VLAN interface for PPPoE/IPoE subscriber sessions, the system's auto-configure process is killed when you delete the auto-configure command under that interface. But the MX Series router does not clear the corresponding VLAN interface. Thus the subscribers can still dial in. [[PR709911](#): This issue has been resolved.]
- On MX Series routers with DPC/FPC, M320 with E3-FPC, M120 and M7i/M10i with E-CFEB, under larger scale environments or negative conditions, core file might be generated on the Packet Forwarding Engine, which is caused by memory corruption. [[PR710853](#): This issue has been resolved.]
- On Systems platforms M320 E3FPC/M120/M7i(10i) CFEB-E with l2vpn or l2circuit, using a control-word and the mpls payload is corrupted in a certain way, the interface might stop forwarding traffic. To recover from this condition a FPC reboot is needed. Only Junos OS Release 10.0 or later is exposed with non-cookie based PICs. MX Series platforms with DPC are not exposed. [[PR720523](#): This issue has been resolved.]
- When receiving large bytes of PIM Join/Prune refreshes at a very rapid rate, it might exhaust ukernel buffer memory on the Packet Forwarding Engine, and the PIM Join/Prune packets will be lost. [[PR720966](#): This issue has been resolved.]
- An I2C read error is seen sometimes on a MIC bootup, causing the MIC to fail to boot up. This fix addresses that error. [[PR733067](#): This issue has been resolved.]
- RPD might generate a core file when a community named in the policy options configuration is changed and that community name is used in a show route CLI

command before the changes effectively take place. [[PR740427](#): This issue has been resolved.]

- With GRES enabled, ksyncd core is observed during switchover when one member of aggregated Ethernet uplink is flapped when router is in steady state. [[PR735437](#): This issue has been resolved.]
- When "delete", or "deactivate" "interface <unit>:family inet:accounting" configuration; FPCs which have the configuration removed may be reset unexpectedly. [[PR743442](#): This issue has been resolved.]
- If IPv6 traffic needs to trigger icmpv6 MTU exceeded message to the source and the source is resolved via next-table next-hop it might leak packet memory on the FPC. [[PR745988](#): This issue has been resolved.]
- Under certain timing conditions, the Packet Forwarding Engine might install an uninitialized firewall filter next hop in the forwarding hardware memory. When a forwarding chipset encounters the uninitialized next hop during packet processing, it triggers an exception error condition. [[PR751088](#): This issue has been resolved.]
- When there are at least three routes to a specific destination (that is, two destination routes and one clone route), deleting and re-adding one of the logical interfaces (that is, board replacement) might trigger a kernel crash due to a timing issue with route deletion. This is triggered in the specific topologies such as an OSPFv3 next hop, which is connected to a different vendor device: lab@shark-re0> show route forwarding-table destination fe80::21f:9eff:fea9:c140 Routing table: default.inet6 Internet6: Destination Type RtRef Next hop Type Index NhRef Netif fe80::21f:9eff:fea9:c140/128 dest 0 0:1f:9e:a9:c1:40 ucst 966 2 ae2.70 fe80::21f:9eff:fea9:c140/128 dest 0 0:1f:9e:a9:c1:40 ucst 968 2 ae4.90 This type of next-hop topology was not seen when a Juniper device established an OSPFv3 adjacency to another Juniper device. [[PR753849](#): This issue has been resolved.]
- On an FPC offline event, mcdiagd logs all the correlated sub-LSPs in the system log. This behavior is independent of having traceoptions enabled or disabled. However, with traceoptions disabled, the logging was happening only in some of the correlated sub-LSPs. In this case, mcdiagd was providing the log messages to syslogd at a faster rate than that could be consumed and displayed by syslogd. The fix throttles the rate of mcdiagd sending messages to syslogd, and also adds an additional event timestamp to the log message to help identify the sub-LSPs that got correlated due to the FPC offline event. [[PR754834](#): This issue has been resolved.]
- Summary: ----- This PR fixes JTREE corruption seen with per-prefix load balancing when the route-memory-enhanced knob is enabled. Affected hardware: ----- Stoli FPCs (FPCs that show up with a 'ES' suffix in the Description column of the 'show chassis hardware' output) and I-chip FPCs (ADPCs and FPCs) are exposed to this. Description: ----- As soon as the forwarding table starts building up on the FPCs, all the affected FPCs will start reporting the following JTREE errors which is an indication of this issue. Apr 29 13:45:54 lab-router fpc5 JTREE(jt_nh_get_reachable_nh32): Not reachable 0x00000000:0x082d1782 for seg 1 (rt_jtree_build_nh) Apr 29 13:45:54 lab-router fpc5 RT: Failed prefix add IPv4 - 1.0.0/24 (jtree nh build failed) on FE 0 Apr 29 13:45:54 lab-router fpc5 RT: IPv4:0 - 1.0.0/24 (add rt entry into jtree failed) These messages will be seen as soon as the forwarding table starts building up, even if there

is no traffic. When traffic starts flowing, the FPCs may crash as a result of this corrupted JTREE. Necessary conditions: ----- 1) route-memory-enhanced knob configured 2) Per-prefix load balancing (which is the default unless per-packet is explicitly enabled) for some or all of the prefixes 3) Routing table with either or both of the following a) IPv4 prefixes with the first octet in the 0 to 127 range; b) IPv6 prefixes beginning with 7fff:: or lower 4) Above routes being resolved over an indirect next-hop (like above routes being received over BGP) 5) The above next-hop is reachable through a load balanced path over Aggregated Ethernet (AEs) or MPLS (LDP / RSVP) LSPs. In case of load balancing over AEs, at least one of the AE links being used in the load-balance must have link IP address a) equal to or greater than 128.X.X.X - if IPv4 b) equal to or greater than 8000:: - if IPv6 Note: ----- - This issue IS NOT seen if per-packet load balancing is configured on the router for all prefixes. [[PR756464](#): This issue has been resolved.]

- Memory leak on a router crashes the ADPC card and generates a core file. [[PR768772](#): This issue has been resolved.]
- When an FPC restart is performed some of the PICs and physical interfaces were unable to be created by chassisd due to EBUSY error returned by kernel. Kernel was unable to process the new requests until the previous states of the same object (PIC, physical interface in our case) was consumed by all peers interested in this. The enhancement addresses the design which makes sure new state changes which could have been processed by the faster peers are not blocked due to these slower peers. [[PR769632](#): This issue has been resolved.]
- The ASN.1 buffered I/O functions in OpenSSL before 0.9.8v do not properly interpret integer data, which allows remote attackers to conduct buffer overflow attacks, and cause a denial of service (memory corruption). J-Web is explicitly not affected by this vulnerability, since J-Web is a server and this is a client-side vulnerability. However, many other functions in Junos OS use these buffered I/O routines and can trigger fetches of untrusted X.509 certificates. Refer to PSN-2012-07-645 for more information. [[PR770702](#): This issue has been resolved.]
- During bulkget application, the Packet Forwarding Engine might generate some corrupted packet. pfd crashes and cores as a result. Fix: Added defensive mechanism to protect the Routing Engine and avoid pfd to crash. [[PR771108](#): This issue has been resolved.]
- T4000-FPC5 forwarding performance is degraded for interfaces with sampling enabled. [[PR771379](#): This issue has been resolved.]
- Certain hardware data structure used for replicating packets on a single the Packet Forwarding Engine stream (SSM list) is not getting updated when the corresponding nexthops get modified, resulting in use of stale data for multicast replication. 2. Impact: All applications that depend on packet replication (IP multicast, P2MP, VPLS BUM traffic) can get impacted. Packets will either be sent out on wrong logical interfaces or can get dropped. However this will happen only if the nexthops used for packet replication get modified. 3. This can affect line cards using the I/J Chipset in MX, TX and M Series chassis. [[PR776149](#): This issue has been resolved.]
- Forwarding performance on T4000-FPC5-3D is degraded when filters with multiple terms are attached to an interface. [[PR777997](#): This issue has been resolved.]

- RN to be added here by ssivaraj [[PR779279](#): This issue has been resolved.]
- Upon Major version ISSU on EQ-DPC with scaling configuration, longer packet loss duration will be seen. [[PR780657](#): This issue has been resolved.]
- On a VPLS circuit with no-tunnel-services configured, when a provider edge (PE) device receives the MPLS frame from the core, it might be dropped by the PE as "L2 mismatch timeout" interface error. [[PR781782](#): This issue has been resolved.]
- While installing image at the validation phase, configuration error is reported even for valid CoS and filter configuration due to this issue. Even after the issue is fixed, image upgrades from Junos OS Release 12.1R1 to a later release fail because the base package (12.1R1) does not have the fix. The following workarounds can be used to overcome this problem: 1. When doing an image upgrade from 12.1R1 to a later release, remove the CoS and Firewall configuration before image installation. The configuration can then be re-applied after the reboot with the new SW version. OR 2. Use 'no-validate' option for the **request system software add** command. [[PR782585](#): This issue has been resolved.]
- 'rms' interfaces are unhidden in the following show commands: show interfaces redundancy request interface revert request interface switchover configured rms interfaces will be shown as an option using "?" or "tab". [[PR784678](#): This issue has been resolved.]

High Availability (HA) and Resiliency

- During high routing churn, a flapping interface can in some rare circumstances result in the replicated (backup) kernel to panic with reason "<interface-name>: bitstring index 14 not empty for <mac-address>". [[PR698608](#): This issue has been resolved.]
- Determining readiness for graceful Routing Engine switchover in an MX Series Virtual Chassis (MX240, MX480, and MX960 routers with MPC/MIC interfaces)--You can use the new check option for the request virtual-chassis routing-engine master switch command to determine whether the member routers in an MX Series Virtual Chassis configuration are ready for a global graceful Routing Engine switchover (GRES) operation from a database synchronization perspective. A global GRES changes the mastership in an MX Series Virtual Chassis by switching the global roles of the master router and backup router in the Virtual Chassis configuration. Depending on the router configuration, a variable amount of time is required before a router is ready to perform a GRES. Attempting a GRES operation before the router is ready can cause system errors and unexpected behavior. Using the **request virtual-chassis routing-engine master switch check** command before you initiate the GRES operation ensures that the subscriber management and kernel databases on both member routers in an MX Series Virtual Chassis are synchronized and ready for the GRES operation. The **request virtual-chassis routing-engine master switch check** command, which you must issue from the Virtual Chassis master router (VC-Mm), checks various system and database components to determine whether they are ready for GRES, but does not initiate the global GRES operation itself. The readiness check includes ensuring that a system timer, which expires after 300 seconds, has completed before the global GRES operation can begin. If the member routers in an MX Series Virtual Chassis are ready for GRES from a database perspective, the **request virtual-chassis routing-engine master switch check** command returns the command prompt and displays no output. If the

member routers are not ready for GRES, the command displays information about the readiness of the system. [Junos OS High Availability Configuration Guide] [[PR722016](#): This issue has been resolved.]

- When performing ISSU on 10GE DPC, peer device will see link flap. [[PR777798](#): This issue has been resolved.]
- When you downgrade Junos OS from Release 11.4R4 to 11.3, a core file is created. This occurs because the interface information on the Gigabit Ethernet IQ PIC gets deleted when the FPC on which the PIC is installed on goes offline during the downgrade. To prevent this error, deactivate the physical interface configuration on the IQ PIC before the downgrade. You can reactivate the configuration after the downgrade. [[PR784291](#): This issue has been resolved.]

Infrastructure

- icmpv6 error codes are not handled correctly and displayed on the source router even when destination router send back correct reply to the source. [[PR718122](#): This issue has been resolved.]
- Certain system resources might become exhausted during Routing Engine switchover under heavy load, causing the system to restart. After restart, the router will operate as expected. [[PR733555](#): This issue has been resolved.]
- Ethernet driver for the internal Ethernet interface on the Routing Engine causes the kernel to crash and Routing Engine reboot. This problem only could happen on Routing Engine models that use "bcm" type of Ethernet interfaces for internal communication. To identify if the Routing Engine is using this type of interface, use the following command:

```
user@router-re1> show interfaces terse
Interface Admin Link Proto Local Remote
[.....]
bcm0 up up <-----
bcm0.0 up up inet 10.0.0.1/8
          10.0.0.5/8
          128.0.0.1/2
          128.0.0.5/2
          inet6 fe80::201:ff:fe00:5/64
          fec0::a:0:0:5/64
          tnp 0x5
[.....]
lsi up up
mtun up up
pimd up up
pime up up
tap up up
```

[[PR734419](#): This issue has been resolved.]

- Fetching ppX interface statistics leaks in pfestat_table are leading to "pfestat_req_add: pfestat table out of ids" error logs. When in this state it is NOT possible to fetch any interface statistics. To recover from this issue, reload the Routing Engine. Products affected by this are non-MX Series products which offer PPPoE services. [[PR751366](#): This issue has been resolved.]

- When the **delete** command is issued on an unnumbered Ethernet user route (static route with a qualified next hop), the destination route created as a part of this user route does not get deleted. This results in duplicate ARP entries for the same address. [[PR752163](#): This issue has been resolved.]
- Address Resolution Protocol (ARP) entries are not flushed out after disabling interface with the purging, aging-timer configured on the local router. [[PR753268](#): This issue has been resolved.]
- The core is due to a null pointer dereference in ND6 code in kernel and this bug was introduced when new feature HFRR was added. An IPv6 route which points to a discard next-hop will not require ND6 cache entry, and this check has been coded in to fix this issue. [[PR755066](#): This issue has been resolved.]
- Processing of a neighbor advertisement can get into an infinite loop in the kernel, given a special set of events with regard to the Neighbor cache entry state and the incoming neighbor advertisement. [[PR756656](#): This issue has been resolved.]
- In scenario where "family inet" is configured in pppoe dynamic profile, if ARP request is received on pppoe interface kernel crash is exhibited. [[PR769646](#): This issue has been resolved.]
- CLI Telnet hangs while executing commands like **show chassis hardware** with inet mtu set to "1500". The Physical MTU is 1600.fix in progress [[PR770454](#): This issue has been resolved.]

Interfaces and Chassis

- This error message is seen continuously on the MX80 router: "fru_is_present: out of range slot -1 for CB". [[PR540868](#): This issue has been resolved.]
- "HS Link FIFO underflow" errors may occur as traffic egresses a PIC when the ingress interface is on another PIC in the same MX-FPC. The speed of the interfaces and the traffic pattern are relevant to this problem. [[PR687905](#): This issue has been resolved.]
- The issue will be seen if the following conditions are satisfied. 1. VPLS routing instance is created with configuration "protocol vpls connectivity-type permanane" 2. LSI interface for the vpls routing-instance has a Primary/secondary MPLS LSP present on the ADPC physical interface which is the DUT. Now on just rebooting the ADPC, these logs are seen on the DPC console. [[PR693066](#): This issue has been resolved.]
- Source MAC address in DHCPv6 packets is set to 00:00:00:00:00:00 when VLAN-demux interface is used to terminate DHCPv6 client [[PR718672](#): This issue has been resolved.]
- In scaled configuration having 1000 vrrp groups, user can see messages like `ppman_vrrpman_change_xmit_status` on the backup vrrp router with the repeated Routing Engine switchover. [[PR719560](#): This issue has been resolved.]
- NTP-related **show** commands, such as **show ntp status** might display incorrect output. [[PR722528](#): This issue has been resolved.]
- MPC2 may reboot when swapping MIC cards in the same MPC [[PR728095](#): This issue has been resolved.]

- On T Series ES type of FPC, BFD sessions might get flapped when other PIC on the same FPC is brought online. This is caused by the fact that the PIC drivers take long time to do initialization when being brought up which might cause the BFD thread to lose chances of processing the keepalive packets and hence drop the sessions. [PR733657: This issue has been resolved.]
- The issue is present in Mx Series platform. MIPs must place their own MAC address in the Egress Identifier TLV in CFM Linktrace Messages that they process. They are incorrectly leaving this value unchanged. [PR735419: This issue has been resolved.]
- The failure of the BFD sessions is caused by a CPU hog in the kernel that is addressed as part of 11.4R3. There was insufficient time to get this patch into 11.4R2. The problem can be circumvented by increasing the BFD session timeout in scaled AE configurations. [PR744882: This issue has been resolved.]
- In an Active / Active MC-LAG scenario, traffic might get dropped if: a. upstream and downstream interfaces are MC-AE interfaces b. you have routing protocols running over the IRB c. traffic crosses the ICL [PR746055: This issue has been resolved.]
- Workaround : Deactivate and activate the interface solves the issue [PR747522: This issue has been resolved.]
- In a scenario where all PPPoE sessions attempt to reconnect simultaneously at a high rate, typically following a reboot or maintenance, it has been observed that the PPPoE daemon might get into a situation where it runs out of memory and fails to create new sessions. The system would be under a lot of stress due to reconnect attempts leading to this problem. A message similar to "allocateSession: no memory for dynamic allocation UIFL xe-5/0/1.1073806881" shows up in the log when PPPoE traceoptions are enabled. [PR747586: This issue has been resolved.]
- In Junos OS Releases earlier than 11.2, the BERT test results would report Error Bit/LOS sec for a newly confirmed E1 link in unframed mode. The issue would be seen on CHSTM1-IQ and CHE1T1-IQE pic. Due to known hardware limitation on CHSTM1-IQ pic, this issue persists on this pic type. However for CHE1T1-IQE pic it has been fixed on Junos OS Release 11.2R7 and later. [PR748175: This issue has been resolved.]
- If rlsq interfaces are part of a routing instance, upon deactivation and activation of routing instance, all rlsq interfaces were not brought up. This issue is fixed as part of this PR. [PR749760: This issue has been resolved.]
- A GE port with optic SFP-FX which has auto-negotiation disabled, it may claim up even though no cable connected and have issue with traffic forwarding on the interface. [PR751536: This issue has been resolved.]
- This issue is specific to a MX Series router with DPCs/FPCs and impacts all type of multicast traffic such as IP multicast packets, or L2 multicast/broadcast packets going through L2vpn/VPLS. An i-chip-based DPC/FPC will only forward multicast traffic to the first 1024 receivers of a multicast group if the total number of receivers on a particular PIC of the DPCE, for that group, is between 1025 and 1088 (1024+64). [PR752662: This issue has been resolved.]
- When "restart routing", "restart chassis-control" and "restart fpc" is done, sometimes DCD is not able to clear the pd-/pe- interfaces as a result fpc is struck in ready state and does not come up. [PR768928: This issue has been resolved.]

- When an F13 SIB is removed, multiple events are handled by chassisd/fabric management. They include marking the plane faulty and recording the SIB's removal. SIB absence is checked at low frequency (every 10 seconds) so it is usually processed later than other events. However, if the SIB removal event wins the race, then we need to mark a SIB faulty when it is already absent and this situation is not handled correctly. [[PR769769](#): This issue has been resolved.]
- For IGMP snooping over A/A MC-LAG, with Integrated routing and bridging enabled or with singly homed receivers, it is required to configure the ICL port as a multicast router interface. However, currently there is no way to configure a trunk port in the default routing instance as a multicast router interface. Hence when an MC-LAG is configured in the default routing instance and the ICL is over a trunk port, IGMP snooping over the MC-LAG might not work correctly with IRB or singly homed receivers. [[PR774556](#): This issue has been resolved.]
- LCP is terminated after LCP negotiation is finished successfully when you use static PPP interfaces, and the intention is to use static IP addressing on the subscriber CPE. To make this work, a configuration like below is needed: `interfaces { pp0 { unit 14 { pppoe-options { underlying-interface ge-5/2/0.612; server; } keepalives interval 10 up-count 3 down-count 3; family inet { address 1.1.9.1/32 { destination 0.0.0.0 ### Gives a hint for static addressing on subscriber side } } } }` Before this fix "destination 0.0.0.0" configuration knob is not taken into account. [[PR777042](#): This issue has been resolved.]
- 802.lag not working as expected in some scenario. [[PR777777](#): This issue has been resolved.]
- PPPoE sessions remain stuck in "init" state. [[PR779047](#): This issue has been resolved.]
- Configuring Multicast address (Inet6) on an interface results in RPD core (mc_ssm_add). [[PR780751](#): This issue has been resolved.]
- In a PPPoE environment, the router might not be able to process any PADI (PPPoE Active Discovery Initiation) packets. This is caused by PPPoED memory leak issue. [[PR781985](#): This issue has been resolved.]

Layer 2 Features

- In a router running a VPLS configuration, an administrator configuration change or a network event that causes the removal of an IFF from a VPLS instance could lead to a panic on the backup Routing Engine. [[PR750036](#): This issue has been resolved.]
- Routing Engine kernel crash was caused by a suspicious packet in the wrong system queue. Packet was classified as a TNP packet (ethertype: 0x8850). TNP is a Layer 3 protocol used for interprocess communication between the Routing Engine and the Packet Forwarding Engine. [[PR779079](#): This issue has been resolved.]
- With VPLS configuration present, when chassis daemon is restarted or if an FPC is rebooted, there is a chance that the dynamic logical interfaces created by rpd are not deleted from the VT interface. As a result, the VT interface will be down along with the VPLS instances. [[PR786263](#): This issue has been resolved.]

Layer 2 Ethernet Services

- DHCP Server Identifier (option 54) in a Request changes from server ip address to Giaddr address [[PR729833](#): This issue has been resolved.]
- The DHCPv6 server application may not properly process DHCPv6 packets when Option 18 (DHCPv6 Interface-ID) and Option 37 (DHCPv6 Relay Agent Remote-ID) are received in the same packet. [[PR774631](#): This issue has been resolved.]
- Extracting option-37 (remote id) might cause an overflow that corrupts other data. The result is a failure of RADIUS account message. [[PR777157](#): This issue has been resolved.]
- DHCP may fail to delete both routes for a DHCPv6 client, leaving the client stuck in the RELEASING state. When a client binds successfully after requesting both an address (IA_NA option) and a prefix (IA_PD option), two routes are added for the client. If only one of the routes is deleted and DHCP fails to properly retry the failed route deletion, the client remains in the RELEASING state. [[PR784977](#): This issue has been resolved.]

Multiprotocol Label Switching (MPLS)

- By design, family MPLS under the virtual-router type routing-instance does not get created without a corresponding "protocols:ldp". Hence, without MPLS family, the MPLS filter is not working on an interface configured under virtual-router type routing-instance. As a workaround, configure ldp in the instance, and disable it on all interfaces if not used. [[PR601989](#): This issue has been resolved.]
- MPLS forwarding broken when labeled BGP routes are distributed to LDP [[PR724658](#): This issue has been resolved.]
- LSP repair scaling improvement has been made by making RSVP more aggressive in sending out messages. [[PR737498](#): This issue has been resolved.]
- Routing protocol daemon may redundantly try to save the RSVP ERO object in the graceful restart database. This is applicable only for non-traffic engineered LSPs when graceful restart is configured. [[PR741694](#): This issue has been resolved.]
- Traffic fails to go through service output when it comes from MPLS core and is routed inside VRF without vrf-table-label configured. This should NOT work on all types of FPCs except MPCs on MX. This PR fixes the problem on MPCs. [[PR749661](#): This issue has been resolved.]
- When LSP configured with auto-bandwidth switches from primary path to secondary path, bandwidth estimation on the secondary path may be under-estimated. Due to under-estimation, overflow sample count may get reset. [[PR752777](#): This issue has been resolved.]
- The kernel may crash at tag_mtu_calc when the Routing Engine attempts to send a packet larger than the configured MPLS MTU, warranting fragmentation (over a LSP) using a l3vpn-composite-nexthop. For the issue to occur both must be true: 1) l3-composite-nexthop knob must be turned on. 2) MPLS MTU must be manually configured by the user. [[PR755950](#): This issue has been resolved.]
- If bfd packets are blocked and mpls ping is succeeding then BFD will not come up after bfd packets are unblocked. [[PR770203](#): This issue has been resolved.]

- 12.1R2: For point-to-point interfaces, RSVP interface state might show as DOWN if interface MTU is changed after the interface was already running RSVP. To avoid running into this issue, make any changes to point-to-point interface configuration before enabling RSVP on the interface. [[PR772807](#): This issue has been resolved.]
- On TRIO platforms, when switch L2 MPLS packets on egress PE routers, the inner MPLS label TTL value is checked and if valid decreased by 1. During the egress process, the TTL value is rechecked. If the value is 1 at this point, the packet is sent to the Routing Engine instead of being forwarded out the interface. [[PR776203](#): This issue has been resolved.]
- Non Juniper routers may send RESV messages with the peak rate value different than the one received in the PATH message. Even though this does not have any functional impact the Junos OS will still log a warning message. With this PR the "RPD_RSVP_INCORRECT_FLOWSPEC: Bandwidth in PATH Tspec greater than RESV flowspec for Session" messages are not logged anymore on peak rate mismatch. [[PR780697](#): This issue has been resolved.]
- A-----B-----C X(primary) X(protector) Y(protector) Y(primary) Given a topology with primary and protector context IDs like above, B seems to have a problem installing LDP label routes in mpls.0 and the forwarding table. [[PR782499](#): This issue has been resolved.]
- RPD coredumps are generated on backup Routing Engine when RSVP+NSR+GRES and "routing-instance provider-tunnel rsvp-te" is configured. [[PR782969](#): This issue has been resolved.]

Network Management and Monitoring

- In Junos OS Release 12.1R1, there may be an issue where multiple SNMP queries for large volumes of information may cause Mib2d to grow in size and eventually create a core file. Mib2d will restart, possibly multiple times, but should recover by itself. [[PR742186](#): This issue has been resolved.]

Platform and Infrastructure

- Packets exchanged between logical routers within the same physical router over logical tunnel (LT) interfaces will not have their TTL decremented. [[PR685639](#): This issue has been resolved.]
- Under very special race conditions, the MPC CPU might stop processing and will be reset due to the Level3/Level 2 watchdog expiration timer. Potential exposure are high load of traffic sent to the Host. The following syslog message will be reported in the syslog once MPC reboots. "fpc[x] MPC: Reset reason (0xc): Level3 watchdog, Level2 watchdog" [[PR717899](#): This issue has been resolved.]
- In case of trio cards, for incoming vlan tagged packets, when some vlan tags are pushed or popped on the ingress pfe and the packet is sent to the host from the egress pfe, the packets received will not be correct. This can typically happen for CFM/OAM packets and cause some CFM/OAM functionality to break. [[PR731639](#): This issue has been resolved.]

- When installing AI Scripts (part of Service Now product) on a device with a very large configuration (>100K lines), the cscript daemon may crash resulting in a core file. [PR736138: This issue has been resolved.]
- When queue rate-limit is configured for interfaces on MPCs, the rate of rate-limit drops reported in 'show interfaces queue' may not be accurate. The value fluctuates between 0 and the actual rate. However the total count of dropped packets and bytes is displayed correctly. [PR740750: This issue has been resolved.]
- Enabling of Dynamic Profile versioning is not supported if dynamic profiles have been already configured on the router. If you deactivate existing dynamic profiles in order to enable and commit dynamic profile versioning, profile version numbers are not subsequently incremented. As a workaround, you must delete all existing dynamic profiles before you enable profile versioning, and then reconfigure the dynamic profiles. [PR741001: This issue has been resolved.]
- Information from other fields in the PR are also shared with customer, namely: Product, Platform, Configuration, State, Affected-Releases, Resolved-In [PR745556: This issue has been resolved.]
- By default, Junos OS will load configured scripts files from /var/db/scripts. If the "load-scripts-from-flash" is configured as below, the Junos OS will load the scripts file from /config/scripts. However, with such configuration, when Junos OS upgrading is performed, the new Junos OS will fail on the configuration validation with messages "mgd error: could not open commit script" etc. Even scripts files are located under both flash and harddisk. router# show system scripts commit { file test-commit-script.xml; } load-scripts-from-flash; [PR746370: This issue has been resolved.]
- This is an issue during the passing of timestamp message from kernel to rmopd for 64 bit Junos OS. Fix will be checked in soon. [PR746428: This issue has been resolved.]
- If a filter contains multiple prefix actions and the filter is applied, changing one prefix action referenced by this filter might crash NPCs. The change on a prefix action could be direct or indirect (For example, changing the policer reference by this prefix action). The workaround is detaching all such filters before changing a prefix action and then applying the filters back after the change. [PR750370: This issue has been resolved.]
- When CoS rewrite is configured for an IRB interface, and the IRB interface participates in L2 multicast, the copies sent over the physical interface will not have the CoS rewrite applied. This issue is applicable only when the chassis is configured in the "enhanced-ip" mode. [PR754720: This issue has been resolved.]
- A MPC-* FPC installed in a MX240/480/960 router or the integrated TFEB of a MX5/10/40/80 router may crash and reboot when the unsupported command **show route hw nhs** is executed from the FPC cli. This command is unsupported and should not be used without the explicit instructions of JTAC. It is not needed for the normal operation of a Juniper router. [PR772413: This issue has been resolved.]
- If "source-filtering" is turned on under an interface, packets with multicast destination mac address will get dropped. Such packets are used by applications like CFM. The multicast mac addresses cannot be explicitly added to be accepted using the CLI. [PR772611: This issue has been resolved.]

- When an event policy, such as those use in AI Scripts, is configured for so-called unstructured syslog messages and also implements a dampening policy, only the first unstructured syslog messages will be processed by the system during the dampening interval. Other unstructured messages, regardless of difference, are also suppressed from being expressed via syslog. This reduces the effectiveness of AI Scripts, but also, due to <https://prsearch.juniper.net/PR612498> suppresses unstructured events from any other logging process. [[PR773712](#): This issue has been resolved.]
- Traffic-control-profile applied on LT logical interface used to terminate a vpls instance has no effect and the logical interface is not shaped. [[PR773764](#): This issue has been resolved.]
- There is known issue with ICMPv6 packet header. This issue is seen only if service-set is enabled on interface. ICMPv6 message generated by FPC will have IPv4 source address instead of IPv6 address of underlying interface on which service-set is enabled. [[PR773828](#): This issue has been resolved.]
- Customers using Junos OS Release 11.4R3.6 code on MX Series routers with MPC 3D 16x 10GE line cards may experience issues with interfaces on these line cards. Some interfaces on the MPC 3D 16x 10GE line cards may be reported as UP ("Enabled" and "Physical Link UP") in the **show interfaces <interface>** command. However, **show interfaces <interface> terse?** command for the same interface reports that interface as DOWN (Admin - UP and Link Protocol - DOWN). The link lights at both ends of the link will be GREEN ? thereby indicating connectivity. However, no traffic passes through the affected interfaces. This issue was seen on interfaces that were part of Aggregated Ethernet (AE) bundle as well as on interfaces that were NOT part of the AE bundle. In addition, "Wedge Detected" messages may be seen in the syslogs and in the telnet/ssh session to the router. This behavior was NOT seen with DPC hardware. This is documented in PR 776727 and upgrading to Junos OS Release 11.4R3.7 fixes the above mentioned issues. [[PR776727](#): This issue has been resolved.]
- Input bridge filter configured on MX Series router with MPCs does NOT work without chassis reboot. [[PR778321](#): This issue has been resolved.]
- On Trio based platforms, LU chip might get into wedge condition which is caused by global Packet Processor Engine (PPE) timeout unrestored after ttrace (Trinity Trace). [[PR785695](#): This issue has been resolved.]

Routing Policy and Firewall Filters

- Changes to the values of MAC Policer for Gigabit Ethernet Interfaces under hierarchy [edit interfaces interface-name gether-options ethernet-switch-profile ethernet-policer-profile policer cos-policer-name], do not reflect on the interface traffic rate after the configuration commit. [[PR739764](#): This issue has been resolved.]

Routing Protocols

- RPD Slow to update change in next-hop index [[PR693808](#): This issue has been resolved.]
- There is a race condition when family route-target is configured on a scaled system. During an IGP change that makes BGP VPN routes inactive, and those routes are

monitored by family route-target, it is possible that RPD may core at a later time. [PR710117: This issue has been resolved.]

- This issue happens when Virtual Router, VRF, and/or VPLS routing-instances are significantly reconfigured and have a large number of route entries. [PR737474: This issue has been resolved.]
- ISO/CLNS prefixes with more than /152 VPN prefix length when advertised by BGP across VPN core causes BGP adjacency flap since the remote BGP rejects the same prefix as invalid address. This is because the ISOVPN draft allows only up to /152 prefixes. [PR742491: This issue has been resolved.]
- Pruned multicast traffic continues to flow from the source even when receiver leaves the multicast group for Junos OS Releases 10.4R8.5, 10.4R9, 10.4R9.2. [PR746474: This issue has been resolved.]
- Route advertisement stops for RT family enabled BGP peers after VRF is deactivated and activated. This issue is only seen with RT enabled peers and non-stop routing enabled. [PR749288: This issue has been resolved.]
- The multipath flash mechanism runs unnecessarily when BGP multipath is configured for inet-vpn routes. When large amounts of inet-vpn routes change, there is a noticeable delay in convergence for the inet-vpn routes. [PR751469: This issue has been resolved.]
- When GRES is done for around 10 times, the backup Routing Engine access freed multicast route and core generated. [PR751702: This issue has been resolved.]
- When RTF(route-target family) is enabled, the special route refresh logic for MVPN address family is skipped. [PR753900: This issue has been resolved.]
- When you are using IS-IS for forwarding only IPv6 traffic and IPv4 routing is not configured, if you perform an SNMP get/walk on an IS-IS routing database table, the RPD process might crash and restart, possibly causing a momentary traffic drop. [PR753936: This issue has been resolved.]
- If BFD authentication configured for IGP and use the meticulous-keyed-sha-1 algorithm, the sequence numbers will not get update every packet. There is no workaround once meticulous-keyed-sha1 is configured and hit the issue where BFD is stuck in INIT state. meticulous-keyed-md5 works as expected and can be used. [PR755303: This issue has been resolved.]
- If BGP receives an ISO-VPN prefix of length 248, i.e. ISO part of prefix contains NSEL-byte, BGP session will be reset. This is according to standards, but it would be good if BGP can handle this gracefully w/o resetting BGP-session. This PR makes BGP handle it gracefully, by ignoring the NSEL byte received in the ISO-VPN prefix. [PR771835: This issue has been resolved.]
- Customer needs these debug logs to be changed to severity LOG_DEBUG. "mcsn[91713]: krt_decode_nexthop: Try freeing: nh-handle: 0x0 nh-index: 1049040 fwdtype: 2" This was introduced as part 11.4 with severity set to "LOG_INFO" (will not be seen with earlier releases). This is used as a debug log and is harmless. "mcsn[91713]: Received MC_AE_OPTIONS TLV for intf device ae1; mc_ae_id 0, status 2" This was introduced as part of RLI 8857 in 10.0 (reference from PR-411614). This also has a severity of

"LOG_INFO" and is part of the rpd-infra that is used by mcsnoopd. [[PR772063](#): This issue has been resolved.]

- Routing protocol daemon (rpd) might dump a core file while processing malformed RIP or RIPng message from neighbor during adjacency establishment. [[PR772601](#): This issue has been resolved.]
- Advertise-external knob not working in C-EBGP scenario. [[PR775175](#): This issue has been resolved.]
- The control plane load balancing was not working when the keyword "JOIN load-balance active" was configured under protocols pim with 4000S and 1 G scenario. Also we saw an issue where the convergence time for such a scenario was four times the time needed as compared to 2000S and 1G. [[PR775707](#): This issue has been resolved.]
- Limited Support for multiple area TLVs in a single ISIS Hello message: When many area TLVs are found in a single IS-IS Hello packet, L1 adjacencies may not be formed correctly and can be stuck in the initializing state. Currently, there are no identified workarounds; however, this does not impact L2 adjacencies. [[PR775852](#): This issue has been resolved.]
- In scaled scenarios where BFD is created with transmit intervals of less than 100msec it is possible to end up in a state where we have duplicate stale entries left behind in the FPC for BFD sessions which do not exist any longer. This causes unnecessary BFD hello packet to be generated and transmitted for each of those entries to the BFD peering router despite not having a valid bfd session. This might result in PPMD and BFD processes on the remote peer routing-engine to report a constantly high CPU utilization. This issue is due to a race condition in PPMD which ends up duplicating transmission entries in FPC instead of cleaning up properly. The fix addresses this race condition. [[PR778813](#): This issue has been resolved.]
- Routers can RPD core in task_timer_delete with an assert "!BIT_ISSET(flags, TIMERF_PROCESSING)" when a rare timing issue occurs while a thread is trying to free a BGP keepalive timer. If this keepalive timer is currently being processed by the keepalive thread in RPD the core will occur. [[PR786842](#): This issue has been resolved.]
- On a dual-Routing Engine system with NSR enabled we could run into an issue with the following symptoms: - RPD on backup Routing Engine showing high CPU utilization - master Routing Engine reporting BGP spooling messages - OSPF neighborship possibly flapping This is due to a bug in RPD NSR infra on the backup-Routing Engine which leads to an increase in the collisions and ultimately the size of the hash-table structure. This results in RPD on backup Routing Engine spending a lot CPU cycles leading to the above symptoms. [[PR788394](#): This issue has been resolved.]

Services Applications

- Support for displaying logical system and routing instance for L2TP tunnels (MX Series 3D Universal Edge Routers)--When you issue the **show services l2tp tunnel** command with the detail or extensive option on either the LAC or LNS, the output now displays both the logical system and the routing instance in which the L2TP tunnel is brought up. [[PR581182](#): This issue has been resolved.]
- When sending traffic through IPSec tunnels for above 2.5Gbps on an MS-400 PIC, the Service-PIC might bounce due to prolonged flow control. [[PR705201](#): This issue has been resolved.]
- If the Service PIC processing DS-Lite packets receives packets from overlapping IPv4 addresses present behind different B4s at the same time then there is a possibility that the PIC will crash with a similar coredump. [[PR711307](#): This issue has been resolved.]
- "linerate-mode" may not be applied correctly to interfaces when first configured on a PIC which does not support "linerate-mode" and later on replace the first PIC with a second PIC which supports "linerate-mode" [[PR734887](#): This issue has been resolved.]
- * Issue here is observing DCD_CONFIG_WRITE_FAILED with "Device not configured" error when system is rebooted with rlsq configuration. * Reason why we are seeing this error, service-pic is unable to create control logical interface hence sp interface is down. Since device is not there and while trying to add lsq logical interfaces, observing the error. After some retries of booting sp pic (say 1 or 2 tries and came up), not seeing these errors and the bundles are UP. * Workaround here we can try deactivate interfaces -> request system reboot -> activate interfaces, instead of only request system reboot. [[PR741121](#): This issue has been resolved.]
- When using Junos OS maintenance releases 11.2R6, 11.4R2 or 12.1R1 (or any Junos OS service releases based on the mentioned maintenance releases) SNMP traps are not generated when the service PIC cpu usage exceeds 85% and/or the memory usage changes from one zone to another. The SNMP traps are not generated due to an internal Junos OS mis-programming. According to the Juniper SP-MIB definitions the following traps should be generated: a) jnxSpSvcSetCpuExceeded (OID 1.3.6.1.4.1.2636.4.10.0.3) - when service PIC CPU usage becomes bigger than 85% b) jnxSpSvcSetCpuOk (OID 1.3.6.1.4.1.2636.4.10.0.4) - when service PIC CPU usage becomes smaller than 85% c) jnxSpSvcSetZoneEntered (OID 1.3.6.1.4.1.2636.4.10.0.1) - when service PIC memory usage enters a specific zone (Yellow, Orange or Red) d) jnxSpSvcSetZoneExited (OID 1.3.6.1.4.1.2636.4.10.0.1) - when service PIC memory usage leaves a specific zone (Yellow, Orange or Red) In Junos OS releases 11.2R7, 11.4R3 or 12.1R2 and all later versions this specific Junos OS feature has been modified to send the SNMP traps correctly. [[PR745190](#): This issue has been resolved.]
- This is a memory leak in the IDPD daemon on the routing engine. It occurs when SNMP queries are done on the routing-engine. This leak is relatively slow and occurs over several days. When the size of the daemon reaches 512M, it dumps a core. [[PR748414](#): This issue has been resolved.]
- There is no logical binding of <flow-analysis-statistics-entry> to <flow-analysis-statistics-pic-info> in the output of "show services stateful-firewall flow-analysis | display xml". At the moment pairs of these tags are just put sequentially

[illegible]

- In some conditions EIM entries are not valid. EIF entries are used when a user tries to create a connection from internet to internal side. As the EIM entries are invalid MSDPC cores, code is fixed in the fixed versions of this PR. [[PR750284](#): This issue has been resolved.]
- Due to some fixes that went into 11.4 we are observing this performance drop. Still trying to figure out the fix causing this drops. The issue is mainly with PBA and with EIM/EIF/APP configuration [[PR751524](#): This issue has been resolved.]
- While trying to Allocate NAT ports for SIP headers, MS-PIC crashes [[PR769605](#): This issue has been resolved.]

- The router may experience a dynamic flow capture process (dfcd) core dump when a Routing Engine switchover occurs in a highly scaled (100,000 or more subscribers) configuration that also has subscriber secure policies enabled. The switchover can result from either GRES or unified ISSU. [PR776698: This issue has been resolved.]
- VC-Power up of VCMb chassis/line cards will cause L2tp LAC clients to terminate [PR777538: This issue has been resolved.]

User Interface and Configuration

- Using the **load** command to replace policy configuration could lead to a configuration corruption which causes RPD to crash upon commit. [PR704294: This issue has been resolved.]
- In the J-Web interface, we recommend that you do not use the Sample check box from the Other Actions section under the Actions tab of the Configure> Security> Filters> IPv4 Firewall Filters page to configure the sampling action modifier. If you save your settings with this check box selected, an error message is displayed. This problem occurs because the set firewall filter foo term bar then sample configuration command has been deprecated in the Junos OS CLI in Release 12.1 and later. [PR781753: This issue has been resolved.]

VPNs

- An optimization has been implemented with BGP-MVPN nexthop infrastructure which will improve scalability in some multi-dimensional scaling scenarios with aggregate interfaces. [PR690690: This issue has been resolved.]
- Under certain circumstances a vrf-import policy's term with the "accept" action that matches the BGP VPN route based on the criteria different than the target community can reject the matching route. [PR706064: This issue has been resolved.]
- Currently, MVPN Leaf-AD routes with IR provider tunnels are sent without the PMSI attributes. These routes should be sent with the PMSI attributes. The label will be the same label as advertised in the Type 1 route. [PR717451: This issue has been resolved.]
- For 12.1R2 continuous rpd restart can lead to rpd crash. problem is seen one in 5 or 6 successive rpd restarts. [PR734276: This issue has been resolved.]
- In NG-MVPN with a multihomed source attached to ingress PE, when original-DR goes down and then comes back to claim its role as DR, the other node will lose its intermediate DR-role and withdraw its type 5 AD-route. However the new DR which comes back will not advertise a type 5 AD route. As result of this misbehavior, neither the non-DR nor the DR will advertise a type 5 AD route in the re-convergence case and hence no egress-PE could join the source. [PR754222: This issue has been resolved.]
- The issue happens when the ingress PE receives the type-4 leaf AD route before discovering the egress PE as a neighbor using a type-1 route. PE ignores the type-4 leaf AD route as there is no nbr. When the ingress PE receives the type-1 route, it only processes inclusive p-tnl and since it did not add the unicast IR tunnel as a leaf to the spmsi tunnel, the egress PE does not receive the traffic. [PR755209: This issue has been resolved.]

- When the label for intra-AS AD route changes, it is not reflected in the intra-as AD route generated to the MVPN PE peers as a result the peers still use the old label information and results traffic drop. [[PR771059](#): This issue has been resolved.]
- In 11.4, PE discards IPv6 BSR messages that come in with a HopLimit != 1. This is causing problems while working with Junos OS Release 10.4 which sometimes sends a BSR message with the default Hop Limit of 64. The fix is to relax this requirement of a Hop Limit of 1. A hidden command is added to enable this strict Hop Limit/TTL checking in case somebody really wants it. [[PR779966](#): This issue has been resolved.]
- A direct route which is primary in another instance and imported into a VRF using rib-groups is not advertised when vrf-table-label is configured. The error message "BGP label allocation failure: Need a nexthop address on LAN" can be seen in "> show route advertising-protocol bgp". [[PR789054](#): This issue has been resolved.]

Resolved Issues in Junos OS Release 12.1R2

The following issues have been resolved since Junos OS Release 12.1R2. The identifier following the description is the tracking number in our bug database.

Class of Service (CoS)

- This issue is specific to traffic-control-profile and scheduler-map constructs in the class of service hierarchy. It is a trio specific issue. When you have a traffic control profile/scheduler-map that is bound to a physical interface hosted on one FPC and the same traffic-control profile is bound to a logical interface hosted on a different FPC, after you move the traffic control profile from the physical interface to its own logical interface, scheduling on this logical interface may not work as expected. You need to have a logical interface on another Packet Forwarding Engine complex referring the same traffic control profile for this issue to happen. Then moving the traffic control profile from logical interface to physical interface and then back again may cause this issue. As a workaround, instead of having only one traffic control profile and moving it from physical to logical interface, you can define two separate traffic control profiles - one exclusively for the physical interface and other for the logical interface. These two traffic control profiles instead of pointing to one scheduler-map would have two different scheduler-maps. The contents of traffic control profiles as well as scheduler maps would be the same. [[PR735870](#): This issue has been resolved.]
- Hierarchical Policer Support on Trio Platforms is available in 11.4R1 and 11.4R2 but is not available on MX-80. The support for MX-80 will be there in 11.4R3 [[PR737500](#): This issue has been resolved.]

Forwarding and Sampling

- Sampled memory increases when interfaces bounce and BGP is running. [[PR594509](#): This issue has been resolved.]
- In scenario where sampling is enabled under forwarding-options, pppoe subscribers churn will cause sampled memory allocations to continuously increase. [[PR741218](#): This issue has been resolved.]

- When the learnable number of MAC address is specified by interface-mac-limit command, but it is not reflected. The number of MAC address will stay at default value. [[PR743934](#): This issue has been resolved.]
- On MX Series Router, an issue is seen in Junos OS Releases such as 10.4S8, 10.4R9.2 and 11.2R1 or later. When **interface:accounting-profile** is configured. The rate of memory leaks depends on the number of interfaces with "accounting-profiles". In our lab testing with 1000 logical interfaces, we experienced up to 50Mbytes per 5 minutes interval. [[PR744537](#): This issue has been resolved.]
- During SNMP poll period 0 stats may be displayed [[PR745211](#): This issue has been resolved.]
- On ADPC cards Output Layer-2 policer drops the packets when configured on a interface with vpls encapsulation. [[PR749141](#): This issue has been resolved.]

General Routing

- During ISSU on MX Series platform, if an MPC has any non-ethernet mics plugged in slot-0, the ISSU procedure will not request offline confirmation for mics in slot-0. [[PR693863](#): This issue has been resolved.]
- In multicast vpn setup with one or more MLFR interface(s) in routing instance, and two MS pics configured in redundant mode, multicast packets with some dscp value are lost on offlining/onlining primary MS pic, followed by deactivating/activating the routing instance. [[PR696475](#): This issue has been resolved.]
- *IF* you have DCU statistics configured in conjunction with the copy-plp class of service knob & a output firewall filter you may encounter a situation where DCU stats are no longer working. Check first that both ingress & egress ports for the flows you are counting are *NOT* on different Packet Forwarding Engines the same fpc. *IF* this is the case then removing the copy-plp knob will restart the DCU statistics collection. [[PR707834](#): This issue has been resolved.]
- If we have an rlsq interface configured under a routing instance and after we perform deactivate/activate on the routing in certain scenarios the ingress traffic from the member media links are not sent to the correct LSQ pic and hence we see traffic black hole. [[PR708720](#): This issue has been resolved.]
- On Systems platforms M320 E3FPC/M120/M7i(10i) CFEB-E with l2vpn or l2circuit, using a control-word and the mpls payload is corrupted in a certain way, the interface might stop forwarding traffic. To recovery from this condition a FPC reboot is needed. Only Junos OS Release 10.0 or later are exposed with non-cookie based PICs. MX platforms with DPC are not exposed. [[PR720523](#): This issue has been resolved.]
- The fix for the PR will be enabling the throttling on the IQ2pics to handle the interrupt storm on Agave pics. [[PR722812](#): This issue has been resolved.]
- On T Series Enhanced Scaling type of FPC, when "route-memory-enhanced" and "l3vpn-composite-nexthop" are configured, Jtree segment 1 memory leak occurs after routing engine switchover. [[PR725623](#): This issue has been resolved.]
- The CLI command **show pfe statistics ip option** might not show proper statistics [[PR727277](#): This issue has been resolved.]

- This issue may be seen whenever a VRF is configured on the Box and the routing instance corresponding to the VRF is changed. "Multiple Free :jt_mem_free" messages will be seen on the fpc. This may result in JTREE memory corruption and in that case will affect forwarding non deterministically. [[PR730686](#): This issue has been resolved.]
- Local Loopback on SONET interface seems to be intermittent with hold-timers configured. [[PR730844](#): This issue has been resolved.]
- Changes related to CoS settings might affect interface byte counter accuracy. [[PR731948](#): This issue has been resolved.]
- In M320, the E3 FPCs could crash when L3VPN CNH, vpn-label-memory-enhanced, route-memory-enhanced knobs are used in conjunction. The fix/workaround is not available as of 11.4R2 [[PR738922](#): This issue has been resolved.]
- Jtree Memory full condition can lead to junk value being stored in composite nexthop structure. A change operation on composite nexthop can lead to crash due to junk value stored in its structure. [[PR739631](#): This issue has been resolved.]
- This issue is seen only where there are tunnel physical interfaces in a configuration on MPC1/2 (QX based) neo mpcs and an ISSU is attempted from a 11.2 based build. If attempted, the tunnel physical interfaces will stop forwarding after ISSU. To fix, deactivate and activate the tunnel physical interfaces (this can be done by deactivating and activating chassis fpc <> pic <> tunnel-services. [[PR739744](#): This issue has been resolved.]
- An I2C read error is seen sometimes on MIC bootup, causing MIC to fail to boot up. This fix addresses that error. [[PR733067](#): This issue has been resolved.]
- Under certain circumstances while doing MIB walks it is possible to get a mib2dcore. This will recover and normal operation will continue without any intervention. All other router functionality will continue as usual. [[PR740692](#): This issue has been resolved.]
- Messages like "Dispatching Synchronous Ethernet Capability TLV" are informational/debug messages and can be safely ignored. [[PR735516](#): This issue has been resolved.]
- When vpls neighbor is configured under user configured mesh group it leads to deletion of the flood route for LSI logical interface(would happen for VT case also). This is observed in M Series as flood route is maintained per logical interface in case of M Series. The problem could happen for any dynamic logical interface. [[PR736993](#): This issue has been resolved.]
- This issue is resolved after necessary software modification [[PR744192](#): This issue has been resolved.]
- The following release notes are applicable only for Junos OS Release 12.1R1. Release Note 1 ----- On the MX Series routers, when the CLI configuration statement "class-of-service interfaces <Physical Interface> shaping-rate" is removed for the 100GE physical interfaces on MPC3E, the throughput will be limited to 16Gbps. As a workaround, an user can explicitly set the shaping-rate to 100Gbps using the CLI configuration statement "set class-of-service interfaces <Physical Interface> 100g". Release Note 2 ----- On the MX Series routers, when the CLI configuration statement "chassis fpc <FPC Slot> pic <PIC Slot> tunnel-services bandwidth" is

removed for MPC3E, the tunnel's throughput will be limited to 16Gbps. As a workaround, a user can deactivate and activate the tunnel-services using the CLI configuration statement "<deactivate/activate> chassis fpc <FPC Slot> pic <PIC Slot> tunnel-services". [[PR744334](#): This issue has been resolved.]

- The allowed options under "set access radius-options request-rate ?" displays the request-rate in seconds [[PR745252](#): This issue has been resolved.]
- During trinity_rt_change, if the new rt-nh is same as old rt-nh then it would copy over stats from "same" old rt-nh to new rt while saving the old stats in route option. When the stats is sent to the Routing Engine, first the current stats is fetched from the shim layer and then the old data from route option is added. In this case the stats value will get doubled while sending it to the Routing Engine. Added a fix to avoid fetching stats from old rt-nh if the new rt-nh is same as old rt-nh. [PR746951](#)
- T4000-FPC5 forwarding performance is degraded for interfaces with sampling enabled. [[PR771379](#): This issue has been resolved.]
- 1. Issue: Certain hardware data structure used for replicating packets on a single Packet Forwarding Engine stream (SSM list) is not getting updated when the corresponding nexthops get modified, resulting in use of stale data for multicast replication. 2. Impact: All applications that depend on packet replication (IP multicast, P2MP, VPLS BUM traffic) can get impacted. Packets will either be sent out on wrong logical interfaces or can get dropped. However this will happen only if the nexthops used for packet replication get modified. 3. This can affect line cards using the I/J Chipset in MX, TX and M Series chassis. [[PR776149](#): This issue has been resolved.]
- Forwarding performance on T4000-FPC5-3D is degraded when filters with multiple terms are attached to an interface. [[PR777997](#): This issue has been resolved.]

High Availability (HA) and Resiliency

- If one or more Packet Forwarding Engine peers are slow in consuming ifstates, that results in slave Routing Engine not sending CP ack to the master Routing Engine within prescribed time. As a result of this, as of today, slave Routing Engine is assumed to be having a problem and hence the connection for slave Routing Engine peer is reset, so that ksyncd can cleanup the ifstates on the slave Routing Engine and resync with master Routing Engine again. With this fix, if slave CP ack is not arrived in prescribed time, if there is any Packet Forwarding Engine which is causing this delay, the same is logged and the CP ack timer is reset. If no peers are found to be causing the delay of slave CP ack, the behavior is retained to reset the slave Routing Engine connection. [[PR727344](#): This issue has been resolved.]
- If static route is configured to point to a broadcast address as next-hop of static route, it can lead to "KSYNCD resync error and core dumped" on backup Routing Engine. This indeed is a configuration issue, but the code is fixed to avoid replication error and core in such a configuration. [[PR732621](#): This issue has been resolved.]
- The MPC can core during ISSU. This issue is intermittent. [[PR744992](#): This issue has been resolved.]

Infrastructure

- Unified ISSU Failure on Devices Configured for Multicast: Because of certain changes in the Junos OS multicast infrastructure in Release 11.1 and later, unified ISSU is not supported in the following scenarios: - Upgrading from Junos OS Releases 10.x to 11.1 or 11.2 on M Series, T Series, and TX Series devices with multicast configuration. - Upgrading from Junos OS Releases 10.x to 11.2 or later on SRX and J-Series devices with multicast configuration - Upgrading from Junos OS Releases 11.1 to 11.2 or later on SRX and J-Series devices with multicast configuration Attempting a unified ISSU might cause undesirable results such as a kernel or routing protocol process crash in such scenarios. Prior to the fix for this PR, backup Routing Engine will crash when performing ISSU go thru above scenario. After this PR, backup Routing Engine will not crash, but it gives error messages. ISSU will fail regardless the fix for PR in above given scenarios. [[PR675582](#): This issue has been resolved.]
- Under certain rare conditions Kernel "devfs" may become locked, this may cause other processes that use /dev/filesystem to wait, eventually some processes start spawning until reaching the maximum limit, as a consequence the Kernel will crash. The following message logged by the kernel is an indication that the system is approaching the maximum number of active processes: /kernel: %KERN-2: nearing maxproc limit by uid 0, please see tuning(7) and login.conf(5). /kernel: %KERN-2: Process with Most Children- 1:init - Children - 365 [[PR678971](#): This issue has been resolved.]
- /mfs partition has only 64MB when only compact flash is available as storage media and for large configuration file is possible to not be able to hold it. So, when HDD is broken and /mfs partition is not large enough to store the configuration files the Routing Engine does not boot. [[PR720540](#): This issue has been resolved.]
- Ethernet driver for the internal ethernet interface on the Routing Engine causes kernel crash and Routing Engine reboot - this problem only could happen on Routing Engine models that use "bcm" type of ethernet interfaces for internal communication. How is to identify if the Routing Engine is using this type of interface: user@router-re1> show interfaces terse Interface Admin Link Proto Local Remote [.....] bcm0 up up <----- bcm0.0 up up inet 10.0.0.1/8 10.0.0.5/8 128.0.0.1/2 128.0.0.5/2 inet6 fe80::201:ff:fe00:5/64 fec0::a:0:0:5/64 tnp 0x5 [.....] lsi up up mtun up up pimd up up pime up up tap up up [[PR734419](#): This issue has been resolved.]
- Routing Engine cored due to Loss of soft watchdog after configuration change. [[PR736927](#): This issue has been resolved.]
- This problem is that there is a socket hole in the received sequence space on backup Routing Engine and that backup Routing Engine cannot handle the TCP SACK from master Routing Engine properly. When backup Routing Engine becomes new master Routing Engine by switchover, this potential sequence mismatch in the previous backup Routing Engine comes out on the new master Routing Engine, therefore, this message is generated on syslog. Once after the new master Routing Engine starts handling TCP SACK properly, this mismatch will be cleared and sooner or later this message stops. This is just a cosmetic issue. [[PR743382](#): This issue has been resolved.]

Interfaces and Chassis

- Backup Routing Engine crashes with panic: "rnh_index_alloc: nhindex" when the lsq bundle with atm member links flaps randomly because of reset/reboot of the remote end device. [[PR675650](#): This issue has been resolved.]
- Backup Routing Engine may go to db prompt if AE configured with multiple FastEthernet members. [[PR692664](#): This issue has been resolved.]
- Regarding the K2-RE (64-bit Routing Engine) when speed/link mode are statically configured on the router for the fxp0 interface, the driver for fxp0 accepts the configuration from DCD process, but does not propagate the setting to the hardware driver. Instead, the driver setting is forced to auto-negotiate. Thus, as the fxp0 interface is auto-negotiating, and the far end device is forced to 100/full, the auto-negotiation on fxp0 will detect the speed but not the duplex and defaults that duplex to half-duplex. [[PR704740](#): This issue has been resolved.]
- Router may crash while trying to bringup 52k pppoev4 subscriber with service-accounting at login in time. [[PR706495](#): This issue has been resolved.]
- COC12 interface stay in down state for extended period of time when deleting local loopback configuration [[PR726762](#): This issue has been resolved.]
- During Routing Engine switchover or connectivity-fault management (CFM) process (cfmd) restart, Connectivity Check Messages (CCM) received from remote Maintenance End Points (MEP) are not processed properly. All action profile events are ignored during the period, after CFMD restart, the events are not cleared. This may cause cfmd not to work properly. [[PR729490](#): This issue has been resolved.]
- If packets are dropped due to high scaling rates, some protocol buffers may be leaked, causing an increase in the memory utilization of the PPP subscriber services daemon that might eventually lead to the inability to login additional PPPoE clients. [[PR731963](#): This issue has been resolved.]
- On T Series ES type of FPC, BFD sessions might get flapped when other PIC on the same FPC is brought online. This is caused by the fact that the PIC drivers take long time to do initialization when being brought up which might cause the BFD thread to lose chances of processing the keepalive packets and hence drop the sessions. [[PR733657](#): This issue has been resolved.]
- See RCCA attached to this PR for Cause/Workaround/Impact. [[PR735122](#): This issue has been resolved.]
- When Ethernet OAM is enabled on member links of a lag group the peer address reported in show oam ethernet link-fault-management output shows the peer address as the lag group's MAC address. [[PR735436](#): This issue has been resolved.]
- When pppoe subscribers dial in and dial out, kernel task devbuf will have memory leak. [[PR735637](#): This issue has been resolved.]
- The issue is seen in the MX Series platform When Active/Active (A/A) mode is enabled and the AE interface name and ICL link start with the same digit, the AE bundle that is going to the MC-LAG device providing connectivity to the Core network does not come up. [[PR736012](#): This issue has been resolved.]

- Whenever there is any alarm(RDI-L, RDI-P or AIS-L) on the interface, on switchover in transmission, 10g port configured in WAN-PHY mode on DPC 4x 10GE R, with hold-time down time of less than 1 second, flaps even before the down hold time is expired. Code has been modified now to address this issue and the fix is available in 11.4R3 11.2R7 12.1R2 10.0R5 10.4R10 [[PR736477](#): This issue has been resolved.]
- Due to an incorrect calculation, memory heap utilization of a service PIC can go over 100% under the "show chassis pic" cli command. There is no service impact. [[PR737676](#): This issue has been resolved.]
- Chassisd core might be generated while performing Routing Engine switchover. It is caused by interrupt storm when master Routing Engine transition to backup and chassisd can be stalled. [[PR738084](#): This issue has been resolved.]
- PPPOE access-internal routes not cleaned up correctly. [[PR739869](#): This issue has been resolved.]
- If the configuration has routing instance(s) with the character "/" in its name and the routing instance has VPLS family configured under it, a configuration change at the top level causes IFFs to get deleted and added back, which results in flapping of the respective VPLS connections. This is a day one issue [[PR740950](#): This issue has been resolved.]
- On lchip platforms, when there are more than 16 ECMP routes exist, after some of the routes flapping (but the available ECMP routes still more than 16), the memory on FPC may be corrupted and error messages like "jtree memory free using incorrect value" may be prompted. [[PR743323](#): This issue has been resolved.]
- On VRRP backup router for a given irb unit, following a non-GRES switchover, ARP entries for this irb may be not learnt over MC-LAG [[PR743385](#): This issue has been resolved.]
- When pppoe subscribers dial in and dial out, kernel task socket will have memory leaks. [[PR745040](#): This issue has been resolved.]
- 'fru_is_present: out of range slot 0 for' logs are displayed in chassisd logs for all the routers where CIP fru is not present. This log is cosmetic in nature. [[PR746923](#): This issue has been resolved.]
- Workaround : Deactivate and activate the interface solves the issue [[PR747522](#): This issue has been resolved.]
- With 'epd-threshold' configured in scheduler-map for ATM2 interface, the EPD threshold (early packet discard) value applied to the interface might not be correct. As a consequence, packets might be unexpectedly tail dropped. This is caused by the defect code of edp-threshold value calculation. [[PR748864](#): This issue has been resolved.]
- If rlsq interfaces are part of a routing instance, upon deactivation and activation of routing instance, all rlsq interfaces were not brought up. This issue is fixed as part of this PR. [[PR749760](#): This issue has been resolved.]

J-Web

- The J-Web interfaces on the J Series and SRX Series devices will not be available on port 32768 or greater, despite the configuration [[PR462624](#): This issue has been resolved.]

Layer 2 Features

- The problem was because of re-using the logical interface index. In case of ccc routes, the logical interface index is used as the route prefix. RPD changes introduced as a part of 572780 fixes this issue. [[PR570168](#): This issue has been resolved.]

Layer 2 Ethernet Services

- With the configuration of STP/AE under IRB interface, you might see kernel panic on both master/backup Routing Engines after a multiple GRES switchover is done. [[PR742940](#): This issue has been resolved.]
- There is a limitation in the support of IRB interfaces used with DHCP such that: 1. If an LT interface is configured as the underlying interface for an IRB interface and a DHCP client requests a unicast response, then instead of rejecting the send operation a corrupt packet is sent. 2. If the underlying interface of an IRB interface has a different number of tags configured than the bridge domain of the IRB interface and a DHCP client requests a unicast response a malformed packet is sent. 3. Regardless of tag configuration on the bridge domain, if the packet needs to be relayed out a VPLS tunnel (an LSI or VT interface as underlying) a malformed packet is sent. [[PR751398](#): This issue has been resolved.]

Multiprotocol Label Switching (MPLS)

- By design, family MPLS under virtual-router type routing-instance does not get created without a corresponding "protocols:ldp". Hence, without MPLS family, MPLS filter is not working on an interface configured under virtual-router type routing-instance. As a workaround, configure ldp in the instance, and disable it on all interfaces if not used. [[PR601989](#): This issue has been resolved.]
- On Junos OS Release 9.6 no-decrement-ttl hides the internal network when trace route is done on a MPLS core running RSVP label-switched path. But the same fails when running 10.4 code, the bug was introduced due to new feature of no-propagate-ttl command which was introduced under vrf. [[PR725779](#): This issue has been resolved.]
- Statically configured P2MP LSPs for VPLS or MVPN flaps when there is subsequent configuration change after branch LSP is added to the P2MP LSP. The P2MP LSP will flap only once after a branch LSP is added to the P2MP LSP. [[PR729522](#): This issue has been resolved.]
- If l3vpn-composite-nexthops were configured under [edit routing-options] and at the same time an output filter was applied to lo0 unit 0 the router might have malformed the L3 part of locally generated MPLS packets (one or more labels). This was limited to only those packets that were originated from within a routing-instance and hence that had their lookup done in the <instance-name>.inet.0 (eg ping routing-instance <instance-name>) - packets generated for the main instance (inet.0) were not affected. Transit packets were unaffected in any case. [[PR734580](#): This issue has been resolved.]

- RPD_MPLS_INTF_MAX_LABELS_ERROR can be seen in log messages even if the maximum-labels option is configured under "family mpls". [[PR734680](#): This issue has been resolved.]
- With nonstop routing (NSR) enabled, under very rare case, the routing protocol daemon (rpd) may crash and dump core file. It is caused when rpd is very busy and there are many background jobs waiting to run. One of them is LDP resync write job. There is a window in which LDP is waiting for LDP resync write job to run to flush the data on session(a). Before the job runs, LDP on standby Routing Engine learns that session(a) is not operational anymore and it moves on to request session sync for session(b). LDP on master starts session sync for session(b) and creates a job LDP sync session. Then the old job LDP resync write runs and wanted to re-start session sync for session(b) by creating a job LDP sync session but finds out that it has already created (It should not have been created) and rpd cores. This window increases with busy rpd and more likelihood of this core to happen. [[PR735337](#): This issue has been resolved.]
- RPD sometimes cores when running snmp walk to mplsXCLSpId. RPD core seems to happen more in case the number of LSP (ingress/egress/transit) is increased. [[PR737147](#): This issue has been resolved.]
- Sometimes the MPLS Autobandwidth value "Max AvgBW util" displays an incorrect value. This was due to an error in the way that the value was being calculated. [[PR737922](#): This issue has been resolved.]
- If member links of an aggregate bundles or ECMP paths are located on MPC in slot 8 or higher l3vpn traffic using l3vpn-composite-nexthops might get dropped. [[PR740719](#): This issue has been resolved.]
- If a P2MP LSP is deleted after graceful-restart, rpd may crash during the deletion of the RSVP session. [[PR741418](#): This issue has been resolved.]
- l3vpn-composite-nexthop with MPC and MSDPC doing stateful-firewall with interface service-set will drop packets on service input direction. The interface where service input/output is configured has to be inside a VRF, and the destination for which service input should intercept the traffic and send it to service PIC in MSDPC should be reachable through MPLS backbone, so resolved through composite next-hop, in order to see this issue. [[PR747914](#): This issue has been resolved.]

Network Management and Monitoring

- When firewall filter counter names are changed, MIB2D assumes all changes are implemented immediately at commit. MIB2D queries the Packet Forwarding Engine before it is updated and therefore gets the incorrect list of counter names. Queries done via SNMP therefore reflect the old firewall counter names. [[PR703606](#): This issue has been resolved.]
- With large scaled of subscribers that are configured over vlan demux interface, high CPU utilization of Mib2d (Management information database of SNMP) might be observed after flapping subscribers. [[PR710573](#): This issue has been resolved.]
- In Junos OS Release 12.1R1, there may be an issue where multiple SNMP queries for large volumes of information may cause Mib2d to grow in size and eventually create

a core file. Mib2d will restart, possibly multiple times, but should recover by itself. [PR742186: This issue has been resolved.]

Platform and Infrastructure

- Random VLAN Ids associated to MAC-address learning from untagged frames on Trio cards. [PR594605: This issue has been resolved.]
- In MVPN scenarios with extranets or hub and spoke NG-MVPNs, it is possible to have more than one customer site using the same provider tunnel on one PE. In such deployments, at the receiver PE with multiple customer sites, if IPv6 multicast traffic is carried over the IPv4 provider tunnel, all customer sites may not receive the IPv6 multicast traffic; only the first site will receive IPv6 traffic. This issue is only for IPv6 over IPv4 tunneled traffic; IPv4 over IPv4 will work correctly. [PR711891: This issue has been resolved.]
- The allow-configuration-regexps statement at the [edit system login class] hierarchy level does not work exactly the same way as the deprecated allow-configuration statement at the same hierarchy level. [PR720013: This issue has been resolved.]
- When l3vpn-composite-nexthop knob is enabled, the VPN label is pushed in the ingress Packet Forwarding Engine. The Egress Packet Forwarding Engine pushes the LSP label. When EXP rewrite is configured on the outgoing interface, the egress Packet Forwarding Engine was performing the rewrite only on the LSP label. [PR723816: This issue has been resolved.]
- Ping with a size that exceeds MTU does not work, because ICMPv6 Path MTU packets to the originator are sent with an incorrect MTU value. [PR725695: This issue has been resolved.]
- Protecting an OSPF Hello with IPSEC AH may cause a cosmetic log entry like the following when "forwarding-options helpers bootp" is configured: Jan 12 13:17:29.946 SomeRouter /kernel: ipsec_is_inbound_pkt_valid(2004): Socket option not set with the addr=130.215.0.129, ifl=331, rtbl_idx=0, so=0xa53c31a0 [PR728779: This issue has been resolved.]
- When queue rate-limit is configured for interfaces on MPC Type 1 3D or MPC Type 2 3D (i.e. non-Q/EQ MPCs), the output of 'show interfaces queue' for such interfaces does not display the count of packets and bytes dropped due to rate-limit. [PR732274: This issue has been resolved.]
- A core dump is generated when a time conditional is applied to a configuration group in a configuration that includes a commit script. [PR732402: This issue has been resolved.]
- VRRP transit Traffic(DA MAC address of VRRP group mac-addresses) flowing on a node, specifically through an physical interface, which has VRRP configured locally, will dropped those packets, if the locally configured VRRP group is different from transit VRRP group. This is only on Trio platforms. [PR732516: This issue has been resolved.]
- When a tagged interface is part of a bridge domain or vpls routing instance with shared vlan learning enabled using "vlan-id all" configuration, and the interface has an explicit input vlan-map which pops all the incoming vlan-ids, the mac-addresses in the bridge

domain/vpls routing instance is learned on an incorrect vlan-id. [[PR733864](#): This issue has been resolved.]

- When tagged traffic is sent on an untagged interface, which is part of a bridge-domain or vpls instance with shared vlan learning (vlan-id all), traffic is learned on VLAN-ID 4096 instead of the incoming vlan-id. [[PR733877](#): This issue has been resolved.]
- If MIC offline is issued after MPC platform performs ISSU, it will potentially cause system crash due to PCIe accessing to IXCHIP being power down. This only happens with MICs having IXCHIP. [[PR735932](#): This issue has been resolved.]
- Port mirroring feature is not working when the ingress and egress(port mirror port) is on different Packet Forwarding Engine for Layer 2 traffic with PPPOE traffic underneath it in a MPC card. [[PR736145](#): This issue has been resolved.]
- The route record fields of certain flows can be incorrect, if these flows are received in an intrafce, while IIF lookup terminates in a non-leaf node. [[PR737472](#): This issue has been resolved.]
- If a GRES is performed before an NPC has finished its fabric training, its attempts to reconnect will not be recognized by the master Routing Engine. In this reconnect period if the FPC crashes chassisd will mark it for removal and cleanup. If the NPC comes on again before the cleanup is done chassisd will accept the connection but 6 minutes later will assert the nmi for this FPC. It will restart again normally. [[PR737774](#): This issue has been resolved.]
- The fix is verified and shud fix the issue [[PR743070](#): This issue has been resolved.]
- LU wedges under (1) scaled flow (> 1M flows), and (2) high export rate (> 1kpps). The workaround is to limit the flow-export-rate to 1kpps. [[PR744230](#): This issue has been resolved.]
- If interface is disabled and enabled cyclically, interface stop forwarding traffic. [[PR744824](#): This issue has been resolved.]
- Information from other fields in the PR are also shared with customer, namely: Product, Platform, Configuration, State, Affected-Releases, Resolved-In [[PR745556](#): This issue has been resolved.]
- Customers using Junos OS Release 11.4R3.6 code on MX Series routers with MPC 3D 16x 10GE line cards may experience issues with interfaces on these line cards. Some interfaces on the MPC 3D 16x 10GE line cards may be reported as UP ("Enabled" and "Physical Link UP") in the "show interfaces <interface>" command. However, "show interfaces <interface> terse?" command for the same interface reports that interface as DOWN (Admin - UP and Link Protocol - DOWN). The link lights at both ends of the link will be GREEN ? thereby indicating connectivity. However, no traffic passes through the affected interfaces. This issue was seen on interfaces that were part of Aggregated Ethernet (AE) bundle as well as on interfaces that were NOT part of the AE bundle. In addition, "Wedge Detected" messages may be seen in the syslogs and in the telnet/ssh session to the router. This behavior was NOT seen with DPC hardware. This is documented in PR 776727 and upgrading to Junos OS Release 11.4R3.7 fixes the above mentioned issues. [[PR776727](#): This issue has been resolved.]

Routing Protocols

- Prior to Junos OS Release 12.2, filtering BGP routes by community value does not work properly when issuing the following CLI command: `show route receive-protocol bgp <neighbor address> <community value>` ex. `{master} juniper@PR01.sjc1> show route receive-protocol bgp 128.241.219.125 1.0.4.0/22 community 2914:420 detail inet.0: 378385 destinations, 2919636 routes (377586 active, 2 holdddown, 1440 hidden) Restart Complete` However, this works properly when filtering based-on the name of the community: `{master} juniper@PR01.sjc1> show route receive-protocol bgp 128.241.219.125 1.0.4.0/22 community-name ALL detail inet.0: 378385 destinations, 2919632 routes (377586 active, 2 holdddown, 1440 hidden) Restart Complete 1.0.4.0/22 (6 entries, 1 announced) Accepted Nexthop: 128.241.219.125 MED: 5 AS path: 2914 4323 7545 7545 7545 7545 56203 | Communities: 2914:420 2914:1008 2914:2000 2914:3000. [PR67777: This issue has been resolved.]`
- The `show bgp group` output is updated to new multiline format in order to display the full name of table `bgp.rtarget.0`. [PR696476: This issue has been resolved.]
- When the traceoptions under the `[edit routing-options]` configuration hierarchy are deactivated, the "ipv6_ra_receive_advertisement" debug message (and others) may continue to be logged well after deactivating the configuration. [PR699797: This issue has been resolved.]
- In the NG-MVN scenario the "maximum-prefixes" knob configured under `<routing_instance>.mvpn.0 rib` does not take effect. Configuration example: `user@router# show routing-instances vrf routing-options rib vrf.mvpn.0 { maximum-prefixes 5 threshold 3; }` Fix is available in the 10.4R9, 11.4R3, 11.2R6 and newer releases. [PR712060: This issue has been resolved.]
- When there are more than one tunnel pic on the system and after bringing down one active tunnel pic, the pe/pd interfaces may fail to switch to another active one in scaled environment. [PR717158: This issue has been resolved.]
- Getting "Error: OID not increasing" while doing `snmpwalk` on `IPMROUTE-STD-MIB::ipMRouteNextHopState`. [PR717893: This issue has been resolved.]
- If BGP export policies run for the routes coming from restarting BGP router (graceful restarting router) during graceful restart, helper router may end up advertising less routes to remote BGP peers. One possible scenario is IGP graceful restart failure during BGP graceful restart. [PR723809: This issue has been resolved.]
- RPD crashes when a route is resolved through indirect next-hop which points to service next-hop, i.e. l2tp session (LNS setup with service PIC). The indirect next-hop in this case is created when having a bgp route that points to a multihop bgp peer reachable through l2tp session, or configuring resolve static route which resolves via l2tp session. [PR725800: This issue has been resolved.]
- In inter-AS option B scenario, unreachable NLRIs can be sent via MP-eBGP session for some prefixes after IGP topology changes. The issue can affect prefixes with multiple paths. [PR730950: This issue has been resolved.]

- On broadcast networks running ISIS, a RPD restart event on one ISIS router could result in the loss of ISIS routes on another router, which will remain in this state until the adjacency is cleared. This issue will not occur on ISIS point-to-point networks. [[PR734158](#): This issue has been resolved.]
- In NSR mode for multicast enabled, backup RPD might dump core at `rt_iflist_kref()`. This has been fixed in 12.1R2. [[PR734769](#): This issue has been resolved.]
- An error in the logic for nonstop routing switchover when family "inet6 labeled-unicast" is configured causes RPD to crash after a nonstop routing switchover. [[PR736669](#): This issue has been resolved.]
- When join-load-balance is not configured and multiple ECMP upstream interfaces flap, the routing protocol process might crash. [[PR739085](#): This issue has been resolved.]
- When enable the use of None-stop-Routing (NSR), auto-rd, or "route target filtering"; BGP's peer group creation will be deferred. During this window, if the command 'show bgp group' is executed it might trigger a rpd core dump. This can only happen during unstable bgp peer groups. [[PR741719](#): This issue has been resolved.]
- ISO/CLNS prefixes with more than /152 VPN prefix length when advertised by BGP across VPN core causes BGP adjacency flap since the remote BGP rejects the same prefix as invalidate address. This is because the ISOVPN draft allows only up to /152 prefixes. [[PR742491](#): This issue has been resolved.]
- With ISO-VPN setup, CLNS traffic outage might occur due to ISIS route loss. Issue might occur when sending a set of prefix via BGP the size of which is higher than the maximum IS-IS fragment size. It could also be seen by redistributing the same set of prefixes from static routes into IS-IS database. [[PR745969](#): This issue has been resolved.]
- The eBGP default behavior, that is the no-advertise-peer-as knob, which is responsible for not sending an update back to the same AS, did not properly filter the advertised update in some cases. [[PR748197](#): This issue has been resolved.]
- If PIM is configured with Dense or Dense Sparse mode, and there are more than 1500 sources for a group we will experience RPD_SCHED_SLIP, with IGMP running high on CPU Cycles. [[PR748420](#): This issue has been resolved.]
- if there are some link micro flapping, it may bring the BFD into a problematic state. as a result, for next event of BFD state down, it will not bring down the client sessions like OSPF, ISIS, BGP. [[PR749388](#): This issue has been resolved.]
- The routing protocol daemon (rpd) crashes and dump a core file after executing 'show ospf context-identifier area <area>' command which is given for an area that has not been configured. The issue is caused by insufficient check code. [[PR750914](#): This issue has been resolved.]
- Limited Support for multiple area TLVs in a single ISIS Hello message: When many area TLVs are found in a single IS-IS Hello packet, L1 adjacencies may not be formed correctly and can be stuck in the initializing state. Currently, there are no identified workarounds; however, this does not impact L2 adjacencies. [[PR775852](#): This issue has been resolved.]

Services Applications

- When you pump in more than 2.1 Million passive monitoring flows into Monitor-II PIC, the router might not send memory overload SNMP trap. [[PR677162](#): This issue has been resolved.]
- IP fragments that are let reaching the MS-PIC or MS-DPC, are not getting filtered out by a stateful-firewall rule configured to explicitly block them hence might hog the compute CPUs in the Service PIC/DPC. [[PR689364](#): This issue has been resolved.]
- In some circumstances, after turning the Routing Engine on/off, some MS-DPC might generate a core file because the Packet Forwarding Engine tries to connect to the Routing Engine after disconnection. [[PR698226](#): This issue has been resolved.]
- Additional session will be seen in the application system cache output. This will not cause any functional impact. [[PR733256](#): This issue has been resolved.]
- Due a software defect, the packets were not reaching the collector devices. This issue is fixed in latest Junos OS Release. [[PR738164](#): This issue has been resolved.]
- It is not possible to have 4000 stable mlppp bundles over l2tp sessions through MSPIC setup if the release used has the fix for 574756. [[PR739338](#): This issue has been resolved.]
- l2tp peer using non default port wont be able to connect to MX LNS. [[PR739488](#): This issue has been resolved.]
- Total ALG errors counter does NOT include SIP ALG errors. The total ALG errors counter is only 1080 whereas SIP ALG error counter is 71M. This issue is fixed now. [[PR739601](#): This issue has been resolved.]
- In Junos OS Releases 10.4 and later, the number of outstanding IPSec tunnels has changed to be 50 tunnels instead of 200 outstanding tunnels in previous releases. [[PR739683](#): This issue has been resolved.]
- When using Junos OS maintenance releases 11.2R6, 11.4R2 or 12.1R1 (or any Junos OS service releases based on the mentioned maintenance releases) SNMP traps are not generated when the service PIC cpu usage exceeds 85% and/or the memory usage changes from one zone to another. The SNMP traps are not generated due to an internal Junos OS mis-programming. According to the Juniper SP-MIB definitions the following traps should be generated:
 - a. **jnxSpSvcSetCpuExceeded** (OID 1.3.6.1.4.1.2636.4.10.0.3) - when service PIC CPU usage becomes bigger than 85%
 - b. **jnxSpSvcSetCpuOk** (OID 1.3.6.1.4.1.2636.4.10.0.4) - when service PIC CPU usage becomes smaller than 85%
 - c. **jnxSpSvcSetZoneEntered** (OID 1.3.6.1.4.1.2636.4.10.0.1) - when service PIC memory usage enters a specific zone (Yellow, Orange or Red)
 - d. **jnxSpSvcSetZoneExited** (OID 1.3.6.1.4.1.2636.4.10.0.1) - when service PIC memory usage leaves a specific zone (Yellow, Orange or Red)

In Junos OS Releases 11.2R7, 11.4R3, or 12.1R2 and all later versions this specific Junos OS feature has been modified to send the SNMP traps correctly. [[PR745190](#): This issue has been resolved.]

- This PR enables visibility of Address Pool Paired out of port errors via the cli command 'show services nat pool detail' user@router-re0> show services nat pool detail Interface: sp-7/0/0, Service set: nat44 NAT pool: public-pool, Translation type:

```
dynamic Address range: 100.100.0.1-100.100.0.254 Port range: 512-65535, Ports in use:
64512, Out of port errors: 0, Max ports used: 64512 AP-P out of port errors: 440601
<<-- errors are now shown here
```

[[PR746752](#): This issue has been resolved.]

- With large NAT rules configuration, installation of service-sets, on PIC level, is not happening. [[PR751858](#): This issue has been resolved.]

Software Installation and Upgrade

- The **request system snapshot partition** command does not create a swap partition that leads to various problems related to mounting memory-based file systems and usage of the swap itself. [[PR746678](#): This issue has been resolved.]

Subscriber Access Management

- The nas-port-extended-format now has an additional value ae-width, listed first. All doc references to nas-port-extended-format should include ae-width which can hold the ae number.

```
regress@grifton# help apropos ae-width
set access profile <profile-name> radius options nas-port-extended-format ae-width
<ae-width>
Number of bits for the aggregated ethernet identifier field [edit] regress@grifton#
...adius radius options nas-port-extended-format ?
Possible completions: adapter-width
Number of bits for the adapter field (0..32 bits) ae-width
Number of bits for the aggregated ethernet identifier field + apply-groups
Groups from which to inherit configuration data + apply-groups-except Don't inherit
configuration data from these groups port-width
Number of bits for the port field (0..32 bits) slot-width
Number of bits for the slot field (0..32 bits) stacked-vlan-width
Number of bits for the S-VLAN subinterface field (bits) vlan-width
Number of bits for the VLAN subinterface field (bits) [edit]
```

[[PR565353](#): This issue has been resolved.]

- PPPoE sessions stuck after invalid parameters from SRC. Fixed in 11.4R3. [[PR728969](#): This issue has been resolved.]

VPNs

- In MVPN environment, when P-tunnels are built using PIM supporting multiple families, after Routing Engine switchover, Kernel Routing Table (KRT) stuck might occur. [[PR717523](#): This issue has been resolved.]
- In a NG-MVPN setup a receiver PE will get the multicast traffic for a particular (C-S,C-G) even if it does not have any local C-Join for that particular group. This can happen if

the following conditions are met: - a selective (S-PMSI) tunnel is used for that particular (C-S,C-G) - the provider tunnel signaling protocol is mLDP - the receiver PE is running Junos OS Release 11.4R1.14 - another receiver PE (running any Junos OS Release) in the same MVPN instance has a C-Join for that (C-S,C-G) [[PR725027](#): This issue has been resolved.]

- In Rosen6 MVPN environment, with nonstop routing (NSR) enabled but PIM NSR disabled, after Routing Engine switchover, it may cause the mt- interface not to be added in the VRF of the new master Routing Engine. [[PR737819](#): This issue has been resolved.]
- The MTU value corresponding to the AC interface configured in L2VPN instances, is being advertised with an inappropriate format. This could prevent L2VPN VC's from being able to get established, should an MTU mismatch is detected between remote ends. [[PR740415](#): This issue has been resolved.]
- With BGP MVPNs when there are many interfaces in the vrf, it is possible that RPD may core. If a forwarding entry has a large number of outgoing interfaces, this memory error will occur. The exact number of oifs needed to trigger this issue is not known. [[PR749379](#): This issue has been resolved.]
- The issue happens when the ingress PE receives the type-4 leaf AD route before discovering the egress PE as a neighbor using a type-1 route. PE ignores the type-4 leaf AD route as there is no nbr. When the ingress PE receives the type-1 route, it only processes inclusive p-tnl and since it did not add the unicast IR tunnel as a leaf to the spmsi tunnel, the egress PE does not receive the traffic. [[PR755209](#): This issue has been resolved.]
- When the label for intra-AS AD route changes, it is not reflected in the intra-as AD route generated to the MVPN PE peers as a result the peers still use the old label information and results traffic drop. [[PR771059](#): This issue has been resolved.]

**Related
Documentation**

- [New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 223](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 284](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 404](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 432](#)

Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers

Errata

Class of Service

- The *Example: Configuring Scheduling Modes on Aggregated Interfaces* topic fails to mention the following additional information regarding the parameters that are scaled for aggregated interface member links when the scheduler parameters are configured using scheduler maps:

Apart from transmit rate and buffer size that are scaled when the parameters are configured using scheduler maps, shaping rate is also scaled if you configure it in bits per second (bps). Shaping rate is not scaled if you configure it as a percentage of the available interface bandwidth.

[*Class of Service, Schedulers on Aggregated Ethernet and SONET/SDH Interfaces*]

High Availability

- The MX Series Virtual Chassis documentation in the *Junos OS High Availability Configuration Guide* failed to include the following information about how slot numbering in the Virtual Chassis affects your use of SNMP.

Junos OS supports the use of SNMP to monitor the routers and other devices in your network. For example, the Juniper Networks jnxBoxAnatomy enterprise-specific Chassis MIB contains the jnxFruTable object, which shows the status of field-replaceable units (FRUs) in the chassis. Within the jnxFruTable object, the jnxFruSlot object displays the slot number where the FRU is installed.

If you are using the jnxFruSlot object in jnxFruTable to display the slot numbers of line cards installed in a member router of an MX Series Virtual Chassis, keep in mind that the offset used for slot numbering in an MX Series Virtual Chassis affects the value that appears for the jnxFruSlot object.

[Table 23 on page 405](#) lists the jnxFruSlot number that appears in the jnxFruTable of the jnxBoxAnatomy MIB, and the corresponding line card physical slot number in each member router of a two-member MX Series Virtual Chassis. For example, a jnxFruSlot value of 15 corresponds to physical slot 3 in member 0 of an MX Series Virtual Chassis. A jnxFruSlot value of 30 corresponds to physical slot 6 in member 1 of an MX Series Virtual Chassis.

Table 23: jnxFruSlot Numbers and Corresponding Slot Numbers in an MX Series Virtual Chassis

jnxFruSlot Number	Line Card Slot Number	MX Series Virtual Chassis Member ID
Line Cards in MX Series Virtual Chassis Member ID 0 (offset = 12):		
12	0	0
13	1	0
14	2	0
15	3	0
16	4	0
17	5	0
18	6	0
19	7	0
20	8	0
21	9	0
22	10	0
23	11	0

Table 23: jnxFruSlot Numbers and Corresponding Slot Numbers in an MX Series Virtual Chassis (*continued*)

jnxFruSlot Number	Line Card Slot Number	MX Series Virtual Chassis Member ID
Line Cards in MX Series Virtual Chassis Member ID 1 (offset = 24)		
24	0	1
25	1	1
26	2	1
27	3	1
28	4	1
29	5	1
30	6	1
31	7	1
32	8	1
33	9	1
34	10	1
35	11	1

[*Junos OS High Availability Configuration Guide, Junos OS SNMP MIBs and Traps Reference*]

Infrastructure

- The description of PR 675582 in the *Previous Releases* subsection under the *Issues in Junos OS Release for M Series, MX Series, and T Series Routers* main section, which lists the resolved issues, ambiguously mentions that unified ISSU is not supported when upgrading from Junos OS Release 10.x to 11.1 or 11.2 on M Series, T Series, and TX Series devices with multicast configuration. The correct statement regarding this limitation on unified ISSU is as follows:

Because of certain changes to the Junos OS multicast infrastructure in Junos OS Release 11.1 and later, unified ISSU is not supported when upgrading from Junos OS Release 10.x to 11.1 or 11.2 and later on M Series, T Series, and TX Series devices with multicast configuration.

[*Release Notes*]

- The following additional information regarding the configuration of peer IP addresses for ICCP peers and multichassis protection for MC-LAG applies to the *Configuring ICCP for MC-LAG* topic:

For Inter-Chassis Control Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer ip-address** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer ip-address** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0.0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

[*Network Interfaces, Ethernet Interfaces*]

- The following additional information regarding the behavior of the **accept-data** statement for MC-LAG in an active-active bridge domain applies to the *Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview* topic:

For a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain and with VRRP configured over an integrated routing and bridging (IRB) interface, you must include the **accept-data** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other layer 3 protocols and applications to work properly over multi-chassis aggregated Ethernet (mc-aeX) interfaces.

[*Network Interfaces, Ethernet Interfaces*]

Interfaces and Chassis

- With Junos OS Release 10.1 and later, you need not include the **tunnel** option or the **clear-dont-fragment-bit** statement when configuring **allow-fragmentation** on a tunnel.

[*Services Interfaces*]

- The new feature description that was entitled “**IPv6 support for inline flow monitoring**” has been replaced in its entirety by a new feature description entitled “**IPv6 support for inline sampling**”. The updated feature description includes a detailed discussion of the new **flow-table-size** configuration statement that enables you to configure the size of IPv4 and IPv6 hash tables in order to improve sampling performance. The **flow-table-size** statement was omitted from the full documentation set for Junos OS Release 12.1.

[*Services Interfaces*]

- **Single-core Routing Engine support for M7i and M10i routers**—The information about the single-core Routing Engine support on M7i and M10i routers is not updated as part of Junos OS Documentation, Release 12.1R2. We plan to add more information about this feature to the *Junos OS Interfaces Fundamentals Configuration Guide* and *Junos OS Installation and Upgrade Guide* in an upcoming release.
- The 20-port Gigabit Ethernet MIC (MIC-3D-20GE-SFP) does not have hardware counters for VLAN frames. Therefore, the **VLAN tagged frames** field displays 0 when the **show interfaces** command is executed on a 20-port Gigabit Ethernet MIC. In other words, the number of VLAN tagged frames cannot be determined for the 20-port Gigabit Ethernet MIC. This information is documented in Junos OS Documentation, Release 12.2 and later only.
- The **lmi-type** statement incorrectly states that Consortium LMI is supported only on M320 routers with Enhanced III FPCs and specific IQE PICs and on MX80, MX240, MX480, and MX960 routers with MICs specified in the *Configuring Tunable Keepalives for Frame Relay LMI* section. The following is the correct compatibility statement:

Consortium LMI is supported on all MPCs and I-chip based FPCs.

[*Frame Relay Interfaces*]

- The **interface-mode** configuration statement topic incorrectly states that you cannot use the **accept** option with the **interface-mode** statement at the [**edit interfaces interface-name unit logical-unit-number family bridge**] and [**edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family bridge**] hierarchy levels to configure a logical interface to accept untagged packets on MPCs or MICs. You can configure the **access** option with the **interface-mode** statement to accept untagged packets on MPCs or MICs on MX Series routers.

[*Network Interfaces, Ethernet Interfaces*]

- The output of **show system core-dumps** has been improved to display **total blocks** and **total files** instead of just **total**.
- **IPFIX sampling documentation did not reference the correct flow template** —The documentation for “Configuring Inline Sampling” and “Configuring Inline Sampling for MX80 Routers” referred to the topic “Configuring Flow Aggregation to Use Version 9

Flow Templates" for information about sampling output, leading customers to believe that the IPv4 BGP_NEXT_HOP was supported for inline sampling. Inline sampling does not use Version 9 templates; they are used only for sampling done on a services PIC.

To view the correct flow template topic, "Configuring Flow Aggregation to Use IPFIX flow Templates", see [PR788037](#).

Services Interfaces

- The **frame-error** configuration statement topic incorrectly states that the default window during which frame errors are counted until they reach the configured threshold is 100 milliseconds. The correct description of the default window is as follows:

The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.

[*Network Interfaces, Ethernet Interfaces*]

- In the *Chassis Conditions That Trigger Alarms* section, the introductory paragraph contains an incorrect link for Table 9. Instead of specifying Table 9 as a link, the "Chassis Component Alarm Conditions on M5 and M10 Routers" heading that points to Table 1 is presented. The correct description is as follows:

Table 1 through Table 9 list the alarms that the chassis components can generate.

Also, the following additional information regarding the generation of alarms when the management interface is down in routers with a single Routing Engine or the master Routing Engine applies to Table 1 through Table 8.

Table 24: Chassis Component Alarm Conditions

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	The Ethernet management interface (fxp0 or em0) on the Routing Engine is down.	- Check the interface cable connection. - Reboot the system. - If the alarm recurs, open a support case using the Case Manager link at http://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States)	Red

[*System Basics, Chassis-Level Features*]

- The **offline-on-fabric-bandwidth-reduction** configuration statement topic incorrectly states that this statement, which is available at the **[edit chassis fpc slot slot-number]**

hierarchy level, has been introduced in Junos OS Release 12.2. This statement was introduced in Junos OS Release 11.4R3 and Junos OS Release 12.1R1.

[*System Basics, Chassis-Level Features*]

- The **show chassis fabric reachability** and the **show chassis fabric unreachable-destinations** command topics fail to state that these commands are also supported on MX240, MX480, and MX960 routers from Junos OS Release 11.4R2 and Junos OS Release 12.1. The Supported Platforms section of this topic fails to mention MX240, MX480, and MX960 routers on which these commands are supported.

[*System Basics and Services Command Reference*]

- The *Interfaces and Chassis* subsection in the *New Features in Junos OS Release for M Series, MX Series, and T Series Routers* section of the Junos OS 12.1R1 and Junos OS 11.4R3 Release Notes fails to describe the following information regarding support for disabling FPCs with degraded fabric bandwidth. This feature is available in Junos OS Release 12.1R1 and later, and Junos OS Release 11.4R3 and later.

Support for disabling an FPC with degraded fabric bandwidth—An FPC working with degraded fabric bandwidth can affect the re-routing process and can cause partial traffic black holes. On an MX960, MX480, or MX240 router, you can now configure the option to bring down an FPC whose fabric bandwidth has degraded because of link errors or bad fabric planes. This configuration is particularly useful in partial black-hole scenarios where bringing the FPC offline results in faster re-routing.

To configure this option on an FPC, use the **offline-on-fabric-bandwidth-reduction** statement at the **[edit chassis fpc slot-number]** hierarchy level.

Configuring this feature does not affect the system. You can configure this feature without restarting the FPC or restarting the system.

[*Release Notes*]

- The *Interfaces and Chassis* subsection in the *New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers* section of the Junos OS 12.1R1 and Junos OS 11.4R2 Release Notes fails to describe the following information regarding support for limiting black hole-time by detecting unreachable destinations. This feature is available in Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later.

Limiting traffic black-hole time by detecting Packet Forwarding Engine destinations that are unreachable over the fabric (MX240, MX480, and MX960 routers)—Enables the MX240, MX480, and MX960 routers to limit traffic black-hole time by detecting unreachable destination Packet Forwarding Engines. The router signals neighboring routers when it cannot carry traffic because of the inability of some or all source Packet Forwarding Engines to forward traffic to some or all destination Packet Forwarding Engines on any fabric plane, after interfaces have been created. This inability to forward traffic results in a traffic black hole.

Packet Forwarding Engine destinations can become unreachable for the following reasons:

- The control boards go offline as a result of a CLI command or a pressed physical button.
- The fabric control boards are turned offline because of high temperature.

- Voltage or polled I/O errors in the SIBs detected by the SPMB.
- All Packet Forwarding Engines receive destination errors on all planes from remote Packet Forwarding Engines, even when the SIBs are online.
- Complete fabric loss caused by destination timeouts, even when the SIBs are online.

When the system detects any unreachable Packet Forwarding Engine destinations, healing from a traffic black hole is attempted. If the healing fails, the system turns off the interfaces, thereby stopping the traffic black hole.

The recovery process consists of the following phases:

1. Fabric plane restart phase: Healing is attempted by restarting the fabric planes one by one. This phase does not start if the fabric plane is functioning properly and a single Flexible PIC Concentrator (FPC) is bad. An error message is generated to specify that a black hole is the reason for the fabric plane being turned offline. This phase is performed for fabric plane errors only.
2. Fabric plane and FPC restart phase: The system waits for the first phase to be completed before examining the system state again. If the black-hole condition still persists after the first phase is performed or if the problem occurs again within a duration of 10 minutes, healing is attempted by restarting both the fabric planes and the FPCs. If you have configured the **action-fpc-restart-disable** statement at the **[edit chassis fabric degraded]** hierarchy level to disable restart of the FPCs when a recovery is attempted, an alarm is triggered to indicate that a traffic black hole has occurred. In this second phase, three steps are taken:
 1. All the FPCs that have destination errors on a PFE are turned offline
 2. The fabric planes are turned offline and brought back online, one by one, starting with the spare plane.
 3. The FPCs that were turned offline are brought back online.
3. FPC offline phase: The system waits for the second phase to be completed before examining the system state again. Traffic black hole is limited by turning the FPCs offline and by turning off interfaces because previous attempts at recovery have failed. If the problem is not resolved by restarting the FPCs or if the problem recurs within 10 minutes after restarting the FPCs, this phase is performed.

By default, the system limits black-hole time by detecting severely degraded fabric. You do not need to configure anything to enable this feature. However, you can limit recovery actions to fabric plane restart only. You need to fix the traffic black hole by performing steps 2 and 3 manually.

In Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later, new alarms are added to indicate which FPCs are creating a traffic black hole in the system and to provide information about FPCs that are turned offline to stop the black hole in the recovery process.

In Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later, new error messages are added to indicate whether the traffic black hole is detected by unreachable FPCs in the system, or it is due to all planes being offline. These messages

also indicate the actions taken on FPCs and planes to stop the black hole—for example, FPC online, FPC offline, FPC restart, FPC power off, plane online, and plane offline.

Two new CLI commands are introduced for this feature:

- The **show chassis fabric unreachable-destinations** command shows the list of destinations that have changed from reachable to unreachable.
- The **show chassis fabric reachability** command shows the current state of fabric destination reachability, based on periodic reachability checks.

[Release Notes]

- The **tunnel-services** configuration statement topic incorrectly states that you can use the **tunnel-services** statement to specify that the IQ2 or IQ2E PIC will work both as a regular PIC and as a tunnel PIC. The correct functionality of the **tunnel-services** statement is as follows:

You can specify the IQ2 and IQ2E PICs to work exclusively in tunnel mode or as a regular PIC. To configure exclusive tunnel mode, use the **tunnel-only** statement at the **[edit chassis fpc slot-number pic slot-number tunnel-services]** hierarchy level. The default setting uses IQ2 and IQ2E PICs as a regular PIC. If you do not configure the **tunnel-only** option, the IQ2 and IQ2E PICs operate as regular PICs.

[System Basics, Chassis-Level Features]

- The **forwarding-mode (100-Gigabit Ethernet)** configuration statement topic fails to mention that this statement is supported on MX Series routers from Junos OS Release 12.1. The Supported Platforms section of this topic fails to list MX Series routers on which this command is supported.

[Network Interfaces, Ethernet Interfaces]

- **New range for message-rate-limit** – The range for message-rate-limit under the syslog configuration for services has changed to 0 through 2147483647.
- **Changes to DDoS protocol groups (MX Series routers)**—The **ipv4-unclassified** and **ipv6-unclassified** DDoS protocol groups have been deprecated in the **protocols** statement at the **[edit system ddos-protection ddos]** hierarchy level. These two protocol groups have also been deprecated from the **show ddos-protection protocols** commands. These groups formerly were used to police all unclassified IPv4 and IPv6 host-bound traffic.

In their place, 10 new protocol groups have been added to the **protocols** statement and the **show ddos-protection protocols** commands:

- **control-layer2**—Unclassified layer 2 control packets.
- **control-v4**—Unclassified IPv4 control packets.
- **control-v6**—Unclassified IPv6 control packets.
- **filter-v4**—Unclassified IPv4 filter action packets; sent to the host because of reject terms in firewall filters.
- **filter-v6**—Unclassified IPv6 filter action packets; sent to the host because of reject terms in firewall filters.

- **host-route-v4**—Unclassified IPv4 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
- **host-route-v6**—Unclassified IPv6 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
- **other**—All unclassified packets that do not belong to another type.
- **resolve-v4**—Unclassified IPv4 resolve packets sent to the host because of a traffic request resolve action.
- **resolve-v6**—Unclassified IPv6 resolve packets sent to the host because of a traffic request resolve action.

[DDoS Configuration]

- “Starting with Junos OS Release 11.4R1, unified ISSU supports Type 2 FPC (MX-FPC2) and Type 3 FPC (MX-FPC3) on the MX Series routers.”

J-Web Interface

- To access the J-Web interface, your management device requires the following software:
 - Supported browsers—Microsoft Internet Explorer version 7.0 or Mozilla Firefox version 3.0
 - Language support—English-version browsers
 - Supported OS—Microsoft Windows XP Service Pack 3

Layer 2 Ethernet Services

- In the *Layer 2 Configuration Guide*, the examples provided in the sections, *Configuring Layer 2 Protocol Tunneling*, *Configuring BPDU Protection on Individual Interfaces*, and *Configuring BPDU Protection on All Edge Ports* are incorrect for configuring layer 2 tunneling with routing instances.
- **Support for WAN PHY mode on 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 routers)**—The information about the WAN PHY mode support on 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+, which is plugged into the Type 5 FPC of T4000 routers, is *not* updated as part of Junos OS Documentation, Release 12.1. We plan to add more information about this feature to the *Junos OS Ethernet Interfaces Configuration Guide* in an upcoming release.
- In the *MX Series Ethernet Services Routers Solutions Guide*, the configuration commands in the *Example: Configuring One VPLS Instance for Several VLANs* section incorrectly describe the usage of the **vlan-id all** statement for normalizing VLANs in VPLS instances. The following information replaces the configuration commands and supplements the description in that section for the sample topology:

Alternatively, instead of configuring a VPLS instance, you can define a virtual switch with a bridge domain and associate the logical interfaces as trunk ports with the bridge domain. This configuration is necessary if you want to configure a list or range of VLAN IDs on the logical interfaces and use the **vlan-id all** statement to normalize VLANs.

A Layer 2 virtual switch, which isolates a LAN segment with its spanning-tree protocol instance and separates its VLAN ID space, filters and forwards traffic only at the data link layer. Each bridge domain consists of a set of logical ports that participate in Layer 2 learning and forwarding.

You can configure VPLS ports in a virtual switch so that the logical interfaces of the Layer 2 bridge domains in the virtual switch can handle VPLS routing instance traffic. VPLS configuration no longer requires a dedicated routing instance of type vpls. Packets received on a Layer 2 trunk interface are forwarded within a bridge domain that has the same VLAN identifier. A trunk interface is implicitly associated with bridge domains based on VLAN membership.

You can use either of the following two mechanisms to normalize VLAN identifiers and process them in a bridge domain or a VPLS routing instance:

- By using the **input-vlan-map** and the **output-vlan-map** statements at the **[edit interfaces interface-name]** hierarchy level to configure VLAN mapping.
- By using either the **vlan-id** statement or the **vlan-tags** statement to configure a normalizing VLAN identifier.

The **vlan-id** and **vlan-tags** statements are used to specify the normalizing VLAN identifier under the bridge domain or VPLS routing instance. The normalizing VLAN identifier is used to perform the following functions:

- Translate, or normalize, the VLAN tags of received packets received into a learn VLAN identifier.
- Create multiple learning domains that each contain a learn VLAN identifier. A learning domain is a MAC address database to which MAC addresses are added based on the learn VLAN identifier.

If you configure the **vlan-id all** statement in a VPLS routing instance, we recommend using the **input-vlan-map pop** and **output-vlan-map push** statements on the logical interface to pop the service VLAN ID on input and push the service VLAN ID on output and in this way limit the impact of doubly-tagged frames on scaling. You cannot use the native **vlan-id** statement when the **vlan-id all** statement is included in the configuration.

For the same network topology illustrated in Figure 1, if VLANs 1 through 1000 for customer C1 span the same sites, you can normalize the VLANs by doing one of the following. Using either of these optimal methods, you can switch and normalize all of these VLANs in an effective, streamlined manner without configuring separately for each VLAN ID.

- By configuring a VPLS routing instance if the logical interfaces are specified with a range of consecutive VLANs or a list of non-contiguous VLAN IDs and using VLAN maps to rewrite the VLAN tags on all of the incoming and outgoing packets on the logical interfaces with a normalized VLAN ID
- By configuring a virtual-switch instance consisting of a set of bridge domains that are associated with one or more logical interfaces configured as a trunk port

You cannot use the **vlan-id** statement to enable VLAN normalization in VPLS instances, if the logical interfaces in the VPLS instance are configured with the **vlan-id-list** or

vlan-id-range statement. In such a scenario, you can use the **input-vlan-map** or the **output-vlan-map** option to achieve VLAN normalization.

The following example illustrates the use of the VLAN mapping functionality in VPLS routing instances to normalize VLANs. This method is beneficial in scenarios with flexible VLAN tagging (asymmetric tag depth). In such an environment, the VLAN configuration data that you specified applies the appropriate VLAN tags to the input and output VLAN maps for the ingress and egress logical interfaces respectively. For example, if certain packets are received as single-tagged packets and if the remaining packets are received as double-tagged packets, using VLAN mapping enables normalization.

Using the VLAN mapping capability is effective only if packets of unequal VLAN tags are received or transmitted from logical interfaces to achieve normalization. We recommend that you do not use VLAN mapping in environments in which the VLAN tags are of equal tag depths for optimal configuration. In such cases, you can use the **vlan-id all** statement to enable normalization of VLANs.

```
[edit]
interfaces ge-1/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id-range 1-1000;
    input-vlan-map {
      push; /* Push the service vlan on input */
      vlan-id 1200; # This VLAN ID is the normalized VLAN for incoming packets
    }
    output-vlan-map pop; /* Pop the service vlan on output */
  }
  unit 11 {
    encapsulation vlan-vpls;
    vlan-id 1500;
  }
}
interfaces ge-2/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id-range 1-1000; # Note the use of the VLAN id range statement.
    input-vlan-map {
      push; /* Push the service vlan on input */
      vlan-id 1300; # This VLAN ID is the normalized VLAN for incoming packets
    }
    output-vlan-map pop; /* Pop the service vlan on output */
  }
}
interfaces ge-3/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    vlan-id 1-1000;
```

```
input-vlan-map {
    push; /* Push the service vlan on input */
    vlan-id 1400; # This VLAN ID is the normalized VLAN for incoming packets
}
output-vlan-map pop; /* Pop the service vlan on output */
}
}
interfaces ge-6/0/0 {
    encapsulation flexible-ethernet-services;
    flexible-vlan-tagging;
    unit 11 {
        encapsulation vlan-vpls;
        vlan-id 1500;
    }
}
routing-instances {
    customer-c1-v1-to-v1000 {
        instance-type vpls;
        interface ge-1/0/0.1;
        interface ge-2/0/0.1;
        interface ge-3/0/0.1;
    } # End of customer-c1-v1-to-v1000
    customer-c1-v1500 {
        instance-type vpls;
        interface ge-1/0/0.11;
        interface ge-6/0/0.11;
    } # End of customer-c1-v1500
} # End of routing-instances
```

The following operations are performed when you use the VLAN mapping configuration:

- Packets received on logical interfaces **ge-1/0/0.1**, or **ge-2/0/0.1**, or **ge-3/0/0.1** with a single VLAN tag in the range from 1 through 1000 in the frame are accepted.
- The VLAN tags of a received packet are compared with the normalized VLAN tags specified with the **vlan-id** statement in the input VLAN map. If the VLAN tags of the received packet are different from the normalized VLAN tags, then the received VLAN tag is converted to the normalized VLAN tag of 1200 for packets received on the logical interface **ge-1/0/0.1**. Similarly, for logical interface **ge-2/0/0.1**, the normalized VLAN tag is 1300 for received packets and for logical interface **ge-3/0/0.1**, the normalized VLAN tag is 1400. Then, the source MAC address of a received packet is learned based on the normalized VLAN configuration.

For output packets, based on the **output vlan-map pop** statement configured on the logical interfaces **ge-1/0/0.1**, or **ge-2/0/0.1**, or **ge-3/0/0.1**, if the VLAN tags associated with an egress logical interface do not match the normalized VLAN tags within the packet, then the VLAN tags in the packets that are being transmitted from the egress logical interface are removed.

- Unknown source MAC addresses and unknown destination MAC addresses are learned based on their normalized VLAN values of 1 through 1000.

- All packets sent on the VPLS pseudowire have a normalized VLAN tag after the source MAC address field in the encapsulated Ethernet packet.
- The **input-vlan-map pop** and **output-vlan-map push** statements on the logical interface cause the service VLAN ID to be popped on input and the service VLAN ID to be pushed on output, and in this way, the impact of doubly-tagged frames on scaling is limited.

The following example illustrates the use of the **vlan-id all** statement in logical interfaces when a virtual switch instance with a bridge domain is associated with the logical interfaces. You can normalize VLANs and create learning domains for each VLAN. A routing instance uses a trunk bridge to connect different departments in an organization, each with their own VLANs, at two different sites. You must configure a bridge domain and VLAN identifier for each VLAN associated with the trunk interface.

```
[edit]
interfaces ge-1/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    family-bridge {
      interface-mode trunk;
      vlan-id-range 1-1000; # Note the use of the VLAN id range statement.
    }
  }
  unit 11 {
    encapsulation vlan-vpls;
    family-bridge {
      interface-mode trunk;
      vlan-id 1500;
    }
  }
}
interfaces ge-2/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    interface-mode trunk;
    vlan-id-range 1-1000; # Note the use of the VLAN id range statement.
  }
}
interfaces ge-3/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    interface-mode trunk;
    vlan-id-range 1-1000; # Note the use of the VLAN id range statement.
  }
}
interfaces ge-6/0/0 {
```

```
encapsulation flexible-ethernet-services;
flexible-vlan-tagging;
unit 11 {
    encapsulation vlan-vpls;
    interface-mode trunk;
    vlan-id 1500;
}
}
routing-instances {
    customer-c1-virtual-switch {
        instance-type virtual-switch;
        interface ge-1/0/0.1;
        interface ge-2/0/0.1;
        interface ge-3/0/0.1;
        bridge-domains {
            c1-vlan-v1-to-v1000 {
                vlan-id all; # Note the use of the VLAN id all statement
            }
        }
    } # End of customer-c1-v1-to-v1000
    customer-c2-virtual-switch {
        instance-type virtual-switch;
        interface ge-1/0/0.11;
        interface ge-6/0/0.11;
        bridge-domains {
            c1-vlan-v1500 {
                vlan-id all; # Note the use of the VLAN id all statement
            }
        }
    } # End of customer-c1-v1500
} # End of routing-instances
```

Note the use of the **vlan-id all** and **vlan-id-range** statements in the virtual-switch instance called **customer-c1-v1-to-v1000**. The **vlan-id all** statement implicitly creates multiple learning domains, each with its own normalized VLAN.

The following operations are performed when you use the **vlan-id all** configuration:

- The logical interfaces are configured as a trunk port that multiplexes traffic from multiple VLANs and usually interconnects switches.
- All the VLAN identifiers are associated with a Layer 2 trunk port. Each of the logical interfaces, **ge-1/0/0.1**, or **ge-2/0/0.1**, or **ge-3/0/0.1**, accepts packets tagged with any VLAN ID specified in the respective **vlan-id-range** statements.
- The association of the received packet to a logical interface is done by matching the VLAN tags of the received packet with the VLAN tags configured on one of the logical interfaces on that physical port. The **vlan-id all** configuration within the bridge domain **c1-vlan-v1-to-v1000** for **customer-c1-virtual-switch** sets the normalized VLAN value. For a logical interface with a single VLAN tag, a learning domain is implicitly created for each normalized VLAN of the interface.
- Bridge domain **c1-vlan-v1-to-v1000** for **customer-c1-virtual-switch** has three logical interfaces:

- Logical interface **ge-1/0/0.1** configured on physical port **ge-1/0/0**.
- Logical interface **ge-2/0/0.1** configured on physical port **ge-2/0/0**.
- Logical interface **ge-3/0/0.1** configured on physical port **ge-3/0/0**.
- Packets received on logical interfaces **ge-1/0/0.11** or **ge-6/0/0.11** with a single VLAN tag of 1500 in the frame are accepted.
- Unknown source MAC addresses and unknown destination MAC addresses are learned based on their normalized VLAN values of 1 through 1000.
- All packets sent on the VPLS pseudowire have a normalized VLAN tag after the source MAC address field in the encapsulated Ethernet packet.
- Although there are only three logical interfaces in the VPLS instance called **customer-c1-virtual-switch**, the same MAC address (for example, M1) can be learned on different logical interfaces for different VLANs. For example, MAC address M1 could be learned on logical interface **ge-1/0/0.1** for VLAN 500 and also on logical interface **ge-2/0/0.1** for VLAN 600.

[MX Series Ethernet Services Routers SolutionsGuide]

MPLS Applications

- Documentation support is missing for RFC 4875, *Extensions to RSVP-TE for Point-to-Multipoint TE LSPs*.

Multicast

- The listings for the following RFCs incorrectly state that Junos OS supports only SSM include mode. Both include mode and exclude mode are supported in Junos OS Release 9.3 and later.
 - RFC 3376, *Internet Group Management Protocol, Version 3*
 - RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

[Hierarchy and Standards Reference]

Network Address Translation (NAT)

- **Limitation on number of terms for NAT rules applied to inline services interfaces**—You are limited to a maximum of 200 for a NAT rule that is applied to an inline services (type si) interface. If you specify more than 200 terms, you will receive the following error when you commit the configuration:

```
[edit]
'service-set service-set-name'
  NAT rule rule-name with more than 200 terms is disallowed for si-x/y/z.n
error: configuration check-out failed
```

Network Management and Monitoring

- The *Supported Network Management Standards* topic fails to mention the following additional information:

On MX Series routers with MPC/MIC interfaces that use the ATM MIC with SFP, Junos OS substantially supports the following RFCs:

- RFC 5603, *PWE3 MIB*
- RFC 5601, *PW-FRAME-MIB*

[*Junos OS Supported Standards*]

- The documentation fails to clearly describe the characters that can be used for SNMPv3 authentication passwords. Besides numbers, uppercase letters, and lowercase letters, the following special characters are supported:

,./\<>;:'[]{}~!@#\$%^*_+=-`

In addition, the following special characters are also supported, but you must enclose them within quotation marks (""") if you enter them on the CLI; if you use a Network Management System to enter the password, the quotation marks are not required:

| & () ?

The documentation also fails to clearly state that characters entered by simultaneously pressing the Ctrl key and additional keys are not supported. [PR/883083: This issue has been resolved]

Routing Policy and Firewall Filters

- The conditions under which enhanced network services mode must be configured for use with firewall filters were documented incorrectly in the *Firewall Filters and Enhanced Network Services Mode Overview* topic. Use the following guidelines when configuring enhanced network services mode with firewall filters:
 - In configurations where interfaces are created either statically or dynamically and firewall filters are applied dynamically, you must configure the chassis network services to run in enhanced mode by including the **network-services** statement at the [edit chassis] hierarchy level.
 - In configurations where interfaces are created statically and firewall filters are applied statically, you must do both of the following:
 - Configure chassis network services to run in enhanced mode by including the **network-services** statement at the [edit chassis] hierarchy level,
 - Configure each firewall filter for enhanced mode by including the **enhanced-mode** statement at the [edit firewall filter *filter-name*], [edit firewall family *family-name* filter *filter-name*], [edit logical-system *logical-system-name* firewall filter *filter-name*], or [edit logical-system *logical-system-name* firewall family *family-name* filter *filter-name*] hierarchy level.
- In routing instances, when a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away.

[*Routing Protocols*]

Routing Protocols

- The following additional information regarding the behavior of MAC addresses in a VPLS dual-homed network with MSTP applies to the *Bridge Priority for Election of Root Bridge and Designated Bridge* topic:

Consider a sample scenario in which a dual-homed customer edge (CE) router is connected to two other provider edge (PE) routers, which function as the VPLS PE routers, with MSTP enabled on all these routers, and with the CE router operating as the root bridge. Integrated Routing and Bridging (IRB) interface is configured for the VPLS routing instances on the routers. In such a network, the MAC addresses that are learned in the VPLS domain continuously move between the LSI or virtual tunnel (VT) interfaces and the VPLS interfaces on both the PE routers. To avoid the continuous movement of the MAC addresses, you must configure root protection by including the **no-root-port** statement at the **[edit routing-instances routing-instance-name protocols mstp interface interface-name]** hierarchy level and configure the bridge priority as zero by including the **bridge priority 0** statement at the **[edit routing-instances routing-instance-name protocols mstp]** hierarchy level on the PE routers. This configuration on the PE routers is required to prevent the CE-side facing interfaces from becoming the root bridge.

[Layer 2 Configuration Guide]

- The *Supported MPLS Standards* topic fails to mention the following additional information:

On MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Junos OS substantially supports RFC 4385, *Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*.

[Junos OS Supported Standards]

- The *Supported Carrier-of-Carriers and Interprovider VPN Standards* topic fails to mention the following additional information:

On MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Junos OS substantially supports the following RFCs:

- RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture*
- RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

[Junos OS Supported Standards]

- The *Supported IPv4, TCP, and UDP Standards* topic fails to mention the following additional information:

Junos OS substantially supports RFC 950, *Internet Standard Subnetting Procedure*

[Junos OS Supported Standards]

Services Applications

- The **rate** statement for packet sampling is now configured at the **[edit forwarding options sampling input family *family*]** hierarchy level.

[Services Interfaces]

- The *Junos OS Release Notes* for Release 12.1R1 erroneously listed “Fragmentation support for GRE-encapsulated packets” as a supported feature for Junos OS Release 12.1R1. This feature is currently unavailable.
- The documentation for the configuration statement **deterministic-port-block-allocation** incorrectly lists the range of acceptable values for **blocksize** as 0 through 512. The correct range is 0 through 64,512.
- The **aes-128-cbc**, **aes-192-cbc**, and **aes-256-cb** options that you can configure with the **encryption-algorithm** statement at the **[edit services ipsec-vpn ipsec proposal proposal-name]** hierarchy level are incorrectly specified as **ase-128-cbc**, **ase-192-cbc**, and **ase-256-cb** options in the following topics in the *Security Services* section of the *System Basics Configuration Guide*:

- *Security Services Configuration Statements*
- *Configuring Minimum IKE Requirements for IPsec on an ES PIC*
- *Configuring an IKE Proposal for Dynamic SAs*
- *encryption-algorithm*

[System Basics, Security Services]

- The **show services stateful-firewall flow-analysis** command should be included in the System Basics and Services Command Reference Guide. This command displays stateful firewall flow statistics.
- The **show services stateful-firewall subscriber-analysis** command should be included in the System Basics and Services Command Reference Guide. This command displays information about the number of active subscribers on the service physical interface card (PIC).
- The following information should be added to the syntax of the “service-set (Services)” configuration statement topic in the *Services Interfaces Configuration Guide*. This information should appear under the **service-set service-set-name** level:

```
service-set-options {  
  bypass-traffic-on-exceeding-flow-limits;  
  bypass-traffic-on-pic-failure>;  
  enable-asymmetric-traffic-processing;  
  support-uni-directional-traffic;  
}
```

This issue was being tracked by PR888803.

- In the Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions Guide, the section “Configuring Address Pools for Network Address Port Translation” should be revised as follows: The following variables should be added
Nr_Addr_PR_Prefix – Number of usable pre-NAT IPv4 subscriber addresses in a “from”

clause match condition $Nr\_Addr\_PU\_Prefix$ – Number of usable post-NAT IPv4 addresses configured in the NAT pool $Rounded\_Port\_Range\_Per\_IP$ – $\lceil (Nr\_Addr\_PR\_Prefix / Nr\_Addr\_PU\_Prefix) \rceil * Block\_Size$ The Forward Translation formulas should be: 1. $Pr\_Offset = Pr\_Prefix - Base\_Pr\_Prefix$ 2. $Pr\_Port\_Offset = Pr\_Offset * Block\_Size$ 3. $Rounded\_Port\_Range\_Per\_IP = \lceil (Nr\_Addr\_PR\_Prefix / Nr\_Addr\_PU\_Prefix) \rceil * Block\_Size$ 4. $Pu\_Prefix = Base\_Public\_Prefix + \text{floor}(Pr\_Port\_Offset / Rounded\_Port\_Range\_Per\_IP)$ 5. $Pu\_Start\_Port = Pu\_Port\_Range\_Start + (Pr\_Port\_Offset \% Rounded\_Port\_Range\_Per\_IP)$ The Reverse Translation formulas should be: 1. $Pu\_Offset = Pu\_Prefix - Base\_Pu\_Prefix$ 2. $Pu\_Port\_Offset = (Pu\_Offset * Rounded\_Port\_Range\_Per\_IP) + (Pu\_Actual\_Port - Pu\_Port\_Range\_Start)$ 3. $Subscriber\_IP = Base\_Pr\_Prefix + \text{floor}(Pu\_Port\_Offset / Block\_Size)$

Subscriber Access Management

- In the *Junos OS Subscriber Access Configuration Guide*, the *Dedicated Queue Scaling for CoS Configurations on Trio MPC/MIC Interfaces Overview* topic incorrectly describes the number of subscribers that can be supported per MPC port on a 60-Gigabit Ethernet Enhanced Queuing MPC. Because this MPC is limited to 16,000 subscribers per PIC, you can accommodate a maximum of 16,000 subscribers per port whether you dedicate 4 or 8 queues per subscriber.

[Subscriber Access]

- In the *Junos OS Subscriber Access Configuration Guide*, the *DHCPv6 Local Server Overview* topic includes a note stating that the DHCPv6 local server does not support dynamic profiles or the local address-assignment pools feature. That note is incorrect—DHCPv6 local server supports address-assignment pool starting in Junos OS Release 10.0, and supports dynamic profiles starting in Junos OS Release 10.1.

[Subscriber Access]

- In the **show network-access aaa radius-servers** command topic, the description of the Status field is incorrect. The correct description is as follows:

RADIUS server status, **UP** or **DOWN**.

If status is **DOWN**, the Status field includes the number of seconds configured by the **revert-interval** statement. When a RADIUS server is unreachable and marked as **DOWN**, the router waits until the revert interval expires before attempting to reconnect to the RADIUS server.



NOTE: If only one RADIUS server is configured, the server status is never marked as **DOWN**. If the RADIUS server is disconnected, the router continues to attempt to contact the server and displays a status of **UP** for the server.

[Subscriber Access]

- The *Configuring Per-Subscriber Session Accounting* topic in the *Subscriber Access Configuration Guide*, incorrectly states that the **update-interval** statement rounds up an interval of 10 through 15 minutes to 15. The actual behavior is that all configured

values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

[Subscriber Access]

- The *DHCP in Broadband Networks* topic erroneously states that the Junos OS subscriber management solution currently supports only DHCP as a multiple-client configuration protocol. However, subscriber management solutions support DHCP and PPPoE as multiple-client configuration protocols.

[Broadband Subscriber Management Solutions]

- The *Configuring Service Packet Counting* topic in the *Junos OS Subscriber Access Configuration Guide* does not include the following configuration guideline. When you specify the **service-accounting** action for the term, you cannot additionally configure the **count** action in the same term.

[Subscriber Access]

- The table titled *Supported Juniper Networks VSAs in the Juniper Networks VSAs Supported by the AAA Service Framework* topic lists RADIUS VSA 26-157 (IPv6-NdRa-Pool-Name). This VSA is not supported and should not appear in the table.

[Subscriber Access]

- The *Configuring a Dynamic Profile for Client Access* topic erroneously uses the **\$junos-underlying-interface** variable when an IGMP interface is configured in the client access dynamic profile. The following example provides the appropriate use of the **\$junos-interface-name** variable:

```
[edit dynamic-profiles access-profile]
user@host# set protocols igmp interface $junos-interface-name
```

- The *Subscriber Access Configuration Guide* and the *System Basics Configuration Guide* contain information about the **override-nas-information** statement. This statement does not appear in the CLI and is not supported.

[Subscriber Access, System Basics]

- When you modify dynamic CoS parameters with a RADIUS change of authorization (CoA) message, Junos OS accepts invalid configurations. For example, if you specify a transmit rate that exceeds the allowed 100 percent, the system does not reject the configuration and returns unexpected shaping behavior.

[Subscriber Access]

- Juniper Networks does not support multicast RIF mapping and ANCP when configured simultaneously on the same logical interface. For example, configuring a multicast VLAN and ANCP on the same logical interface is not supported, and the subscriber VLANs are the same for both ANCP and multicast.

[Subscriber Access]

- The *Subscriber Access Configuration Guide* incorrectly describes the **authentication-order** statement as it is used for subscriber access management. When configuring the **authentication-order** statement for subscriber access management, you must always

specify the **radius** method. Subscriber access management does not support the **password** keyword (the default), and authentication fails when you do not specify an authentication method.

[*Subscriber Access*]

- In the *Subscriber Access Configuration Guide*, the *Juniper Networks VSAs Supported by the AAA Service Framework* table and the *RADIUS-Based Mirroring Attributes* table incorrectly describe VSA 26-59. The correct description is as follows:

Attribute Number	Attribute Name	Description
26-59	Med-Dev-Handle	Identifier that associates mirrored traffic to a specific subscriber.

[*Subscriber Access*]

- In the *Subscriber Access Configuration Guide*, the table titled "Supported Juniper Networks VSAs" in the "Juniper Networks VSAs Supported by the AAA Service Framework" topic lists RADIUS VSA 26-42 (Input-Gigapackets) and VSA 26-43 (Output-Gigapackets). These two VSAs are not supported.

[*Subscriber Access*]

- In the *Junos OS Subscriber Access Configuration Guide*, the "Qualifications for Change of Authorization" section in the topic titled "RADIUS-initiated Change of Authorization (CoA) Overview", has been rewritten as follows to clarify how CoA uses the RADIUS attributes and VSAs.

Qualifications for Change of Authorization

To complete the change of authorization for a user, you specify identification attributes and session attributes. The identification attributes identify the subscriber. Session attributes specify the operation (activation or deactivation) to perform on the subscriber's session and also include any client attributes for the session (for example, QoS attributes). The AAA Service Framework handles the actual request.

Table 25 on page 426 shows the identification attributes for CoA operations.



NOTE: Using the Acct-Session-ID attribute to identify the subscriber session is more explicit than using the User-Name attribute. When you use the Acct-Session-ID, the attribute identifies the specific subscriber and session. When you use the User-Name as the identifier, the CoA operation is applied to the first session that was logged in with the specified username. However, because a subscriber might have multiple sessions associated with the same username, the first session might not be the correct session for the CoA operation.

Table 25: Identification Attributes

Attribute	Description
User-Name [RADIUS attribute 1]	Subscriber username.
Acct-Session-ID [RADIUS attribute 44]	Specific subscriber and session.

Table 26 on page 426 shows the session attributes for CoA operations. Any additional client attributes that you include depend on your particular session requirements.

Table 26: Session Attributes

Attribute	Description
Activate-Service [Juniper Networks VSA 26–65]	Service to activate for the subscriber.
Deactivate-Service [Juniper Networks VSA 26–66]	Service to deactivate for the subscriber.

[Subscriber Access]

- The *Junos OS Subscriber Management Scaling Values (XLS)* spreadsheet erroneously states that the maximum number of PPPoE interfaces per MPC1 is 15,996. The correct value is 31,998.

[Subscriber Management Scaling]

- The *Example: HTTP Service Attached to a Static Interface* topic in the *Junos OS Subscriber Access Configuration Guide* provides an incorrect example for configuring a service filter as a walled garden. The correct example is as follows:

The following example uses a service filter as a walled garden by defining a rule named `redirect`, referencing the rule in a profile named `http-redirect`, configuring a service set named `http-redirect` that references the `http-redirect` captive portal content delivery profile, and attaching the `http-redirect` service set to static interface `ge-1/0/1.0`.

```
[edit services]
captive-portal-content-delivery {
  rule redirect {
    match-direction input;
    term t1 {
      from {
        destination-address {
          100.0.1.1/32;
        }
      }
      then {
        redirect http://www.google.com;
      }
    }
  }
  profile http-redirect {
    cpcd-rules redirect;
  }
}
service-set http-redirect {
  captive-portal-content-delivery-profile http-redirect;
  interface-service {
    service-interface ms-1/0/0;
  }
}

[edit interfaces ge-1/0/1]
unit 0 {
  family inet {
    service {
      input {
        service-set http-redirect service-filter walled;
      }
      output {
        service-set http-redirect;
      }
    }
    address 10.1.3.2/24;
  }
}
```

[Subscriber Access]

- In the *Junos OS Subscriber Access Configuration Guide*, the *Example: Configuring an L2TP LNS* topic fails to show the correct predefined variable, `$junos-ipv6-address`, that sets the IPv6 addresses for the dual-stack configuration when routing advertisement is used. Instead, the topic erroneously shows a non-existent predefined variable, `$junos-ipv6-interface`.

- Several statistics fields described in the *show interfaces demux0 (Demux Interfaces)* and *show interfaces (PPPoE)* topics in the *Junos OS Interfaces Command Reference* might be misleading.

The **IPv6 Transit statistics** and **Transit statistics** fields displayed by **show interfaces demux0** and **show interfaces oo0** are described as reporting statistics on a logical interface for traffic that transits the router. In fact, these statistics also count packets and octets that do not leave the router due to drop profiles and other filters. In contrast, accurate accounting statistics for the same logical interface are presented to RADIUS that count only packets and octets that leave the router.

- The **show subscribers** topic in the *Junos OS System Basics and Services Command Reference* omits the following information about using the **address** option for the **show subscribers** command.

When you issue the **show subscribers address** command, you must specify the IPv4 or IPv6 address prefix *without* a netmask, as shown in the following example:

```
user@host> show subscribers address 192.168.17.1 detail
```

If you specify the IP address as a prefix *with* a netmask, as shown in the following example, the router displays a message that the IP address is invalid, and rejects the command.

```
user@host> show subscribers address 192.168.17.1/32 detail
Invalid argument: invalid ip_address 192.168.17.1/32
```

[*System Basics and Services Command Reference*]

- The documentation for the subscriber management domain mapping feature in the *Subscriber Access Configuration Guide* describes using the **aaa-logical-system** and **target-logical-system** statements to configure mapping to a non-default logical system. Subscriber management is supported in the default logical system only. Configuring a non-default logical system is a future extension of subscriber management and is not supported in current Junos OS releases.

[*Subscriber Access*]

- The *Dedicated Queue Scaling for CoS Configurations on Trio MPC/MIC Interfaces Overview* topic in the *Junos OS Subscriber Access Configuration Guide* does not explain the queuing behavior on 30-Gigabit Ethernet Queuing MPCs with only one MIC. See [Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview](#) for a more complete explanation of dedicated queue scaling.
- In the *Subscriber Access Configuration Guide*, there is an error in the *Example: Configuring RADIUS-Based Subscriber Authentication and Accounting* topic. In the example, the **profile** stanza incorrectly includes the **authentication** statement. The correct statement is **authentication-order**, as shown in the following sample:

```
profile isp-bos-metro-fiber-basic {
  authentication-order radius;
}
```

[*Subscriber Access*]

- **L2TP support for SNMP statistics (MX Series routers)**—By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in the table have a default value of zero.

Table 27: SNMP Counters for L2TP Statistics

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	Tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 30,000 sessions, none of these statistics is more than 15 minutes old.



BEST PRACTICE: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

[Subscriber Access]

- The *MX Series 3D Universal Edge Router Interface Module Reference* does not state that VLAN demux configurations are not supported on MX Series routers that have any of the following line cards installed:
 - Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q)
 - Enhanced Queuing IP Services DPCs (DPCE-R-Q)

The nonsupport includes any configuration stacked on top of a VLAN demux. For example, although PPPoE is supported, PPPoE over aggregated Ethernet interfaces is not supported when one of these cards is installed, because this configuration requires PPPoE to be stacked on a VLAN demux.

Timing and Synchronization

- The *Supported Time Synchronization Standards* topic fails to mention the following additional information:

On MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Junos OS substantially supports RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

[*Junos OS Supported Standards*]

User Interface and Configuration

- The **show system statistics bridge** command displays system statistics on MX Series routers.

[*System Basics*]

- As of Junos OS Release 12.1 and later, using **request system rollback software** reverts to the last known good state before the most recent **request system software (add | delete)** command.
- There is an erroneous note in the *Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 Traffic* topic. This topic is present in the *Junos OS Firewall Filters Configuration Guide*.

The note mentions that the “Filter IPv4 traffic within an MPLS flow up to 5 labels is not supported on T4000 routers.” Instead this feature is not supported on Type 5 FPC of T4000 routers.

- Four new options for creating strong passwords have been added to the **[edit system login password]** hierarchy level. They include:
 - **minimum-nums** *number*
 - **minimum-upper-cases** *number*
 - **minimum-lower-cases** *number*
 - **minimum-punctuations** *number*

Using several password minimum requirement options will cause the **minimum-length** to be reset if the total sum of the required minimums exceeds the **minimum-length** setting.

- The note in the *Installing the J-Web Software* topic that mentions that M Series or T Series routers must be running Junos OS version 7.3 or later to support the J-Web interface is incomplete. The following note accurately describes the support of the J-Web interface on M Series and T series routers.

M Series routers or T320, T640, and TX Matrix routers must be running Junos OS version 7.3 or later to support the J-Web interface. Except the T320, T640, and TX Matrix routers, other T Series routers do not support the J-Web software.

[*J-Web Interface User Guide*]

VPNs

- In *Chapter 19, Configuring VPLS* of the *VPNs Configuration Guide*, an incorrect statement that caused contradictory information about which platforms support LDP BGP interworking has been removed. The M7i router was also omitted from the list of supported platforms. The M7i router does support LDP BGP interworking.

[VPNs]

- Documentation support is missing for RFC 2917, *A Core MPLS IP VPN Architecture*.
- The **l3vpn** statement documentation states that this statement is not supported on MX Series routers with both MS-DPCs and MPCs installed. However, it should state that the **l3vpn** statement is not supported on MX Series routers with both DPCs and MPCs installed.

[VPNs]

- The following guideline regarding the support of LSI traffic statistics on M Series routers is missing from the *General Limitations on IP-Based Filtering* section in the *Filtering Packets in Layer 3 VPNs Based on IP Headers* topic:

Label-switched interface (LSI) traffic statistics are not supported for Intelligent Queuing 2 (IQ2), Enhanced IQ (IQE), and Enhanced IQ2 (IQ2E) PICs on M Series routers.

[VPNs, Layer 3 VPNs]

- The following limitation regarding firewall filters configured in conjunction with the **vrf-table-label** statement is missing from the *General Limitations on IP-Based Filtering* in the *Filtering Packets in Layer 3 VPNs Based on IP Headers* topic:

Firewall filters cannot be applied to interfaces included in a routing instance on which you have configured the **vrf-table-label** statement.

This documentation is applicable to all J Series, M Series, T Series, and SRX Series routers.

[VPNs, Layer 3 VPNs]

Changes to the Junos OS Documentation Set

The following are the changes made to the Junos OS documentation set:

- Carrier-grade NAT and softwire documentation is no longer included in the *Junos OS Services Configuration Guide*. The documentation is now available at the following subject-based web page: Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions—http://www.juniper.net/techpubs/en_US/junos121/information-products/pathway-pages/ngna-solutions/next-generation-network-addressing-solutions.html
- ALG documentation for MX Series platforms has been updated. The topic has been reorganized and expanded, with particular emphasis on SIP and SIP-NAT interaction. An updated version of the documentation is available at the following PR link location: [PR817816](#)

- Stateless firewall filter and traffic policer documentation is no longer included in the *Junos OS Policy Framework Configuration Guide*. This material is now available in the *Junos OS Firewall Filter and Policer Configuration Guide* only.
- Routing policy, traffic sampling, forwarding, and monitoring documentation is no longer included in the *Junos OS Policy Routing Framework Configuration Guide*. Individual HTML pages have a **Print** link in the upper left corner of the text area on the page. This material is now available in the *Junos OS Policy Framework Configuration Guide*.
- A new topic, *CGN Implementation: Best Practices*, which provides experience-based recommendations for configuring carrier-grade NAT, has been added to the documentation set. The new topic is available at http://www.juniper.net/techpubs/en_US/junos12.1/topics/concept/nat-best-practices.html

Related Documentation

- [New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 223](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 284](#)
- [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 298](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 432](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers

This section discusses the following topics:

- [Basic Procedure for Upgrading to Release 12.1 on page 432](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 435](#)
- [Upgrading a Router with Redundant Routing Engines on page 435](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 436](#)
- [Upgrading the Software for a Routing Matrix on page 437](#)
- [Upgrading Using ISSU on page 438](#)
- [Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR on page 439](#)
- [Downgrading from Release 12.1 on page 440](#)

Basic Procedure for Upgrading to Release 12.1

In order to upgrade to Junos OS 10.0 or later, you must be running Junos OS 9.0S2, 9.1S1, 9.2R4, 9.3R3, 9.4R3, 9.5R1, or later minor versions, or you must specify the **no-validate** option on the **request system software install** command.

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 *Junos OS Installation and Upgrade Guide*.



NOTE: With Junos OS Release 9.0 and later, the compact flash disk memory requirement for Junos OS is 1 GB. For M7i and M10i routers with only 256 MB memory, see the Customer Support Center JTAC Technical Bulletin PSN-2007-10-001 at <https://www.juniper.net/alerts/viewalert.jsp?txtAlertNumber=PSN-2007-10-001&actionBtn=Search>.



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 *Junos OS System Basics Configuration Guide*.

The download and installation process for Junos OS Release 12.1 is different from earlier Junos OS releases.

Follow these steps for the download and installation process:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks web page:
<http://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 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 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.

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

```
user@host> request system software add validate reboot  
source/jinstall-12.1R61-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-12.1R61-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 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 can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 12.1 **jinstall** package, you cannot issue the **request system software rollback** command to return to the previously installed software. 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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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 <http://www.juniper.net/support/eol/junos.html>.

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 [Junos OS Installation and Upgrade Guide](#).

Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1

In releases prior to Junos OS Release 10.1, the draft-rosen multicast VPN feature implements the unicast **lo0.x** address configured within that instance as the source address used to establish PIM neighbors and create the multicast tunnel. In this mode, the multicast VPN loopback address is used for reverse path forwarding (RPF) route resolution to create the reverse path tree (RPT), or multicast tunnel. The multicast VPN loopback address is also used as the source address in outgoing PIM control messages.

In Junos OS Release 10.1 and later, you can use the router's main instance loopback (**lo0.0**) address (rather than the multicast VPN loopback address) to establish the PIM state for the multicast VPN. We strongly recommend that you perform the following procedure when upgrading to Junos OS Release 10.1 if your draft-rosen multicast VPN network includes both Juniper Network routers and other vendors' routers functioning as provider edge (PE) routers. Doing so preserves multicast VPN connectivity throughout the upgrade process.

Because Junos OS Release 10.1 supports using the router's main instance loopback (**lo0.0**) address, it is no longer necessary for the multicast VPN loopback address to match the main instance loopback address **lo0.0** to maintain interoperability.



NOTE: You might want to maintain a multicast VPN instance **lo0.x** address to use for protocol peering (such as IBGP sessions), or as a stable router identifier, or to support the PIM bootstrap server function within the VPN instance.

Complete the following steps when upgrading routers in your draft-rosen multicast VPN network to Junos OS Release 10.1 if you want to configure the routers's main instance loopback address for draft-rosen multicast VPN:

1. Upgrade all M7i and M10i routers to Junos OS Release 10.1 before you configure the loopback address for draft-rosen Multicast VPN.



NOTE: Do not configure the new feature until all the M7i and M10i routers in the network have been upgraded to Junos OS Release 10.1.

2. After you have upgraded all routers, configure each router's main instance loopback address as the source address for multicast interfaces. Include the **default-vpn-source interface-name loopback-interface-name** statement at the **[edit protocols pim]** hierarchy level.

3. After you have configured the router's main loopback address on each PE router, delete the multicast VPN loopback address (**lo0.x**) from all routers.

We recommend that you remove the multicast VPN loopback address from all PE routers from other vendors. In Junos OS releases prior to 10.1, to ensure interoperability with other vendors' routers in a draft-rosen multicast VPN network, you had to perform additional configuration. Remove that configuration from both the Juniper Networks routers and the other vendors' routers. This configuration should be on Juniper Networks routers and on the other vendors' routers where you configured the **lo0.mvpn** address in each VRF instance as the same address as the main loopback (**lo0.0**) address.

This configuration is not required when you upgrade to Junos OS Release 10.1 and use the main loopback address as the source address for multicast interfaces.



NOTE: To maintain a loopback address for a specific instance, configure a loopback address value that does not match the main instance address (**lo0.0**).

For more information about configuring the draft-rosen Multicast VPN feature, see the *Junos OS Multicast Configuration Guide*.

Upgrading the Software for a Routing Matrix

A routing matrix can use either a TX Matrix router as the switch-card chassis (SCC) or a TX Matrix Plus router as the switch-fabric chassis (SFC). By default, when you upgrade software for a TX Matrix router or a TX Matrix Plus router, the new image is loaded onto the TX Matrix or TX Matrix Plus router (specified in the Junos OS CLI by using the **scc** or **sfc** option) and distributed to all T640 routers or T1600 routers in the routing matrix (specified in the Junos OS CLI by using the **lcc** option). To avoid network disruption during the upgrade, ensure that the following conditions are met before beginning the upgrade process:

- A minimum of free disk space and DRAM on each Routing Engine. The software upgrade will fail on any Routing Engine without the required amount of free disk space and DRAM. To determine the amount of disk space currently available on all Routing Engines in the routing matrix, use the CLI **show system storage** command. To determine the amount of DRAM currently available on all the Routing Engines in the routing matrix, use the CLI **show chassis routing-engine** command.
- The master Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and T640 routers or T1600 routers (LCC) are all **re0** or are all **re1**.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and T640 routers or T1600 routers (LCC) are all **re1** or are all **re0**.
- All master Routing Engines in all routers run the same version of software. This is necessary for the routing matrix to operate.
- All master and backup Routing Engines run the same version of software before beginning the upgrade procedure. Different versions of Junos OS can have incompatible message formats especially if you turn on GRES. Because the steps in the process include changing mastership, running the same version of software is recommended.

- For a routing matrix with a TX Matrix router, the same Routing Engine model is used within a TX Matrix router (SCC) and within a T640 router (LCC) of a routing matrix. For example, a routing matrix with an SCC using two RE-A-2000s and an LCC using two RE-1600s is supported. However, an SCC or an LCC with two different Routing Engine models is not supported. We suggest that all Routing Engines be the same model throughout all routers in the routing matrix. To determine the Routing Engine type, use the CLI **show chassis hardware | match routing** command.
- For a routing matrix with a TX Matrix Plus router, the SFC contains two model RE-DUO-C2600-16G Routing Engines, and each LCC contains two model RE-DUO-C1800-8G Routing Engines.



NOTE: It is considered best practice to make sure that all master Routing Engines are re0 and all backup Routing Engines are re1 (or vice versa). For the purposes of this document, the master Routing Engine is re0 and the backup Routing Engine is re1.

To upgrade the software for a routing matrix:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine (re0), and save the configuration change to both Routing Engines.
2. Install the new Junos OS release on the backup Routing Engine (re1) while keeping the currently running software version on the master Routing Engine (re0).
3. Load the new Junos OS on the backup Routing Engine.
4. After making sure that the new software version is running correctly on the backup Routing Engine (re1), switch mastership back to the original master Routing Engine (re0) to activate the new software.
5. Install the new software on the new backup Routing Engine (re0).

For the detailed procedure, see the [Routing Matrix with a TX Matrix Feature Guide](#) or the [Routing Matrix with a TX Matrix Plus Feature Guide](#).

Upgrading Using ISSU

Unified 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. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the *Junos OS High Availability Configuration Guide*.

Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR

Junos OS Release 9.3 introduced NSR support for PIM for IPv4 traffic. However, the following PIM features are not currently supported with NSR. The commit operation fails if the configuration includes both NSR and one or more of these features:

- Anycast RP
- Draft-Rosen multicast VPNs (MVPNs)
- Local RP
- Next-generation MVPNs with PIM provider tunnels
- PIM join load balancing

Junos OS Release 9.3 introduced a new configuration statement that disables NSR for PIM only, so that you can activate incompatible PIM features and continue to use NSR for the other protocols on the router: the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. (Note that this statement disables NSR for all PIM features, not only incompatible features.)

If neither NSR nor PIM is enabled on the router to be upgraded or if one of the unsupported PIM features is enabled but NSR is not enabled, no additional steps are necessary and you can use the standard upgrade procedure described in other sections of these instructions. If NSR is enabled and no NSR-incompatible PIM features are enabled, use the standard reboot or ISSU procedures described in the other sections of these instructions.

Because the **nonstop-routing disable** statement was not available in Junos OS Release 9.2 and earlier, if both NSR and an incompatible PIM feature are enabled on a router to be upgraded from Junos OS Release 9.2 or earlier to a later release, you must disable PIM before the upgrade and reenable it after the router is running the upgraded Junos OS and you have entered the **nonstop-routing disable** statement. If your router is running Junos OS Release 9.3 or later, you can upgrade to a later release without disabling NSR or PIM—simply use the standard reboot or ISSU procedures described in the other sections of these instructions.

To disable and reenable PIM:

1. On the router running Junos OS Release 9.2 or earlier, enter configuration mode and disable PIM.

[edit]

user@host# **deactivate protocols pim**
user@host# **commit**
2. Upgrade to Junos OS Release 9.3 or later software using the instructions appropriate for the router type. You can either use the standard procedure with reboot or use ISSU.
3. After the router reboots and is running the upgraded Junos OS, enter configuration mode, disable PIM NSR with the **nonstop-routing disable** statement, and then reenable PIM.

[edit]

```
user@host# set protocols pim nonstop-routing disable
user@host# activate protocols pim
user@host# commit
```

Downgrading from Release 12.1

To downgrade from Release 12.1 to another supported release, follow the procedure for upgrading, but replace the 12.1 **jinstall** package with one that corresponds to the appropriate release.



NOTE: You cannot downgrade more than three releases. For example, if your routing platform is running Junos OS Release 11.4, you can downgrade the software to Release 10.4 directly, but not to Release 10.3 or earlier. As a workaround, you can first downgrade to Release 10.4 and then downgrade to Release 10.3.

For more information, see the *Junos OS Installation and Upgrade Guide*.

Related Documentation

- [New Features in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 223](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 284](#)
- [Issues in Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 298](#)
- [Errata and Changes in Documentation for Junos OS Release 12.1 for M Series, MX Series, and T Series Routers on page 404](#)

Junos OS Documentation and Release Notes

For a list of related Junos OS documentation, see <http://www.juniper.net/techpubs/software/junos/>.

If the information in the latest release notes differs from the information in the documentation, follow the *Junos OS Release Notes*.

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <http://www.juniper.net/books>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales 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 <http://www.juniper.net/customers/support/downloads/710059.pdf>.
- Product warranties—For product warranty information, visit <http://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: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/> .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at <http://www.juniper.net/support/requesting-support.html>.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net:pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to **support@juniper.net**. For documentation issues, fill out the bug report form located at <https://www.juniper.net/cgi-bin/docbugreport/>.

Revision History

06 August 2013—Revision 4, Junos OS 12.1 R7 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

30 July 2013—Revision 3, Junos OS 12.1 R7 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

24 July 2013—Revision 2, Junos OS 12.1 R7 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

23 July 2013—Revision 1, Junos OS 12.1 R7 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

11 July 2013—Revision 6, Junos OS 12.1 R6 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

21 June 2013—Revision 5, Junos OS 12.1 R6 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

14 June 2013—Revision 4, Junos OS 12.1 R6 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

07 May 2013—Revision 3, Junos OS 12.1 R6 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

30 April 2013—Revision 2, Junos OS 12.1 R6 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

23 April 2013—Revision 1, Junos OS 12.1 R6 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

28 March 2013—Revision 21, Junos OS 12.1 R5 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

21 February 2013—Revision 20, Junos OS 12.1 R5 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

31 January 2013—Revision 19, Junos OS 12.1 R5 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

30 January 2013—Revision 18, Junos OS 12.1 R5 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

23 January 2013—Revision 17, Junos OS 12.1 R5 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

22 January 2013—Revision 16, Junos OS 12.1 R5 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

29 October 2012—Revision 15, Junos OS 12.1 R4 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

24 September 2012—Revision 14, Junos OS 12.1 R3 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

30 August 2012—Revision 13, Junos OS 12.1 R3 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

21 August 2012—Revision 12, Junos OS 12.1 R3 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

14 August 2012—Revision 11, Junos OS 12.1 R3 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

19 July 2012—Revision 10, Junos OS 12.1 R2 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

28 June 2012—Revision 9, Junos OS 12.1 R2 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

18 June 2012—Revision 8, Junos OS 12.1 R2 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

07 June 2012—Revision 7, Junos OS 12.1 R2 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

05 June 2012—Revision 6, Junos OS 12.1 R2 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

15 May 2012—Revision 5, Junos OS 12.1 R1 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

12 April 2012—Revision 4, Junos OS 12.1 R1 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

10 April 2012—Revision 3, Junos OS 12.1 R1 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

03 April 2012—Revision 2, Junos OS 12.1 R1 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

28 March 2012—Revision 1, Junos OS 12.1 R1 – High End SRX Series, Branch SRX Series, J Series, EX Series, M Series, MX Series, and T Series.

Copyright © 2013, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.