

Release Notes: Junos[®] OS Release 14.1R7 for the EX Series, M Series, MX Series, PTX Series, and T Series

16 March 2016

Contents

Introduction	6
Junos OS Release Notes for EX Series Switches	6
New and Changed Features	6
Hardware	6
Flow Monitoring	7
Platform and Infrastructure	8
Changes in Behavior and Syntax	8
Dynamic Host Configuration Protocol	9
Interfaces	9
Network Management and Monitoring	9
Platform and Infrastructure	9
Known Behavior	10
Known Issues	10
Dynamic Host Configuration Protocol	10
Interfaces and Chassis	11
Network Management and Monitoring	11
Routing Protocols	11
Software Installation and Upgrade	11
Resolved Issues	12
Resolved Issues: Release 14.1R7	12
Resolved Issues: Release 14.1R6	13
Resolved Issues: Release 14.1R5	13
Resolved Issues: Release 14.1R4	15
Resolved Issues: Release 14.1R3	16
Resolved Issues: Release 14.1R2	17
Documentation Updates	18
Migration, Upgrade, and Downgrade Instructions	18
Upgrade and Downgrade Support Policy for Junos OS Releases	18

Product Compatibility	19
Hardware Compatibility	19
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers	21
New and Changed Features	21
Hardware	22
Authentication, Authorization and Accounting (AAA) (RADIUS)	26
Class of Service (CoS)	26
Dynamic Host Configuration Protocol (DHCP)	28
Forwarding and Sampling	28
General Routing	28
High Availability (HA) and Resiliency	29
Interfaces and Chassis	31
IPv6	35
Layer 2 Features	36
MPLS	37
Multicast	38
Network Management and Monitoring	38
Network Operations and Troubleshooting Automation	39
Platform and Infrastructure	40
Port Security	40
Routing Policy and Firewall Filters	41
Routing Protocols	42
Services Applications	43
Software Installation and Upgrade	46
Spanning-Tree Protocols	47
Subscriber Management and Services	48
User Interface and Configuration	53
VLAN Infrastructure	53
VPNs	54
Changes in Behavior and Syntax	57
Authentication, Authorization and Accounting (AAA)	57
Application Layer Gateways (ALGs)	57
Class of Service (CoS)	57
High Availability (HA) and Resiliency	58
Interfaces and Chassis	59
IPv6	61
MPLS	61
Network Management and Monitoring	61
Routing Policy and Firewall Filters	62
Routing Protocols	62
Security	63
Services Applications	63
Subscriber Management and Services	65
User Interface and Configuration	68
VPNs	68
Known Behavior	70
Hardware	70
High Availability (HA) and Resiliency	71

MPLS	72
Software Installation and Upgrade	72
Subscriber Management and Services	72
Known Issues	72
General Routing	73
Authentication and Access Control	80
Class of Service (CoS)	80
Forwarding and Sampling	80
General Routing	81
High Availability (HA) and Resiliency	81
Infrastructure	81
Interfaces and Chassis	82
Layer 2 Features	84
Layer 2 Ethernet Services	84
MPLS	84
Multicast	85
Network Management and Monitoring	85
Platform and Infrastructure	86
Routing Protocols	89
Services Applications	90
Subscriber Access Management	91
User Interface and Configuration	92
VPNs	92
Resolved Issues	93
Resolved Issues 14.1R7	93
Resolved Issues: 14.1R6	108
Resolved Issues: 14.1R5	121
Resolved Issues: 14.1R4	145
Resolved Issues: 14.1R3	154
Resolved Issues: 14.1R2	170
Documentation Updates	179
Adaptive Services Interfaces Feature Guide	180
Chassis-Level Feature Guide	181
Ethernet Interfaces Feature Guide	181
Firewall Filters Feature Guide for Routing Devices	182
High Availability Feature Guide for Routing Devices	182
Interchassis Redundancy Using Virtual Chassis Feature Guide for MX Series Routers	182
Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide	184
Junos OS Administration Library for Routing Devices	185
Layer 2 Configuration Guide, Bridging, Address Learning, and Forwarding	185
Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices	185
MPLS Applications Feature Guide for Routing Devices	185
Overview for Routing Devices	186
OSPF Feature Guide	186
Services Interfaces Configuration Guide	186
Standards Reference	189

Subscriber Management Network Access Feature Guide	189
Subscriber Management Provisioning Guide	189
System Log Messages Reference	189
Traffic Sampling, Forwarding, and Monitoring Feature Guide for Routing Devices	190
User Access and Authorization Feature Guide for Routing Devices	190
VPLS Feature Guide for Routing Devices	191
VPNs Library for Routing Devices	191
Migration, Upgrade, and Downgrade Instructions	192
Basic Procedure for Upgrading to Release 14.1	192
Upgrade and Downgrade Support Policy for Junos OS Releases	194
Upgrading a Router with Redundant Routing Engines	194
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1	195
Upgrading the Software for a Routing Matrix	196
Upgrading Using Unified ISSU	197
Downgrading from Release 14.1	197
Changes Planned for Future Releases	198
Product Compatibility	199
Hardware Compatibility	200
Junos OS Release Notes for PTX Series Packet Transport Routers	201
New and Changed Features	201
Hardware	201
Interfaces and Chassis	203
MPLS	205
Network Management and Monitoring	206
Routing Protocols	206
Changes in Behavior and Syntax	207
Interfaces and Chassis	208
Network Management and Monitoring	208
Routing Protocols	208
VPNs	208
Known Behavior	209
Known Issues	209
General Routing	209
Class of Service (CoS)	210
Interfaces and Chassis	210
MPLS	211
Routing Protocols	211
Resolved Issues	212
Resolved Issues: 14.1R7	212
Resolved Issues: 14.1R6	213
Resolved Issues: 14.1R5	215
Resolved Issues: 14.1R4	217
Resolved Issues: 14.1R3	218
Resolved Issues: 14.1R2	220
Documentation Updates	221

Migration, Upgrade, and Downgrade Instructions	222
Upgrading Using Unified ISSU	222
Upgrading a Router with Redundant Routing Engines	222
Basic Procedure for Upgrading to Release 14.1R6	222
Product Compatibility	225
Hardware Compatibility	226
Finding More Information	227
Documentation Feedback	227
Requesting Technical Support	227
Self-Help Online Tools and Resources	228
Opening a Case with JTAC	228
Revision History	229

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, J Series, M Series, MX Series, PTX Series, QFabric systems, QFX Series, SRX Series, and T Series.

These release notes accompany Junos OS Release 14.1R7 for the EX Series, M Series, MX Series, PTX Series, and T Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 14.1R7 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 14.1 for the EX Series.

- [Hardware](#)
- [Flow Monitoring](#)
- [Platform and Infrastructure](#)

Hardware

- **High-speed Switch Fabric module for EX9200 switches**—Starting with Junos OS Release 14.1, a high-speed Switch Fabric module, EX9200-SF2, is supported. Compared to the original SF module, EX9200-SF, EX9200-SF2 offers increased bandwidth, providing higher-capacity traffic support in settings that require greater interface density (slot and capacity scale).

SF modules are installed horizontally on the front panel of the switch chassis. You can install either one or two SF modules in an EX9204 or EX9208 switch and either two or three SF modules in an EX9214 switch.

The Switch Fabric serves as the central nonblocking matrix through which all network data passes. The key functions of the Switch Fabric are:

- Monitor and control system functions
- Interconnection of all line cards
- Clocking, system resets, and booting control
- Routing Engine carrier

EX9200-SF2 supports all EX9200 line cards.



NOTE: When you upgrade from an EX9200-SF module to an EX9200-SF2 module in an EX9200 switch, the SF module types (that is, EX9200-SF and EX9200-SF2) can co-exist in the switch *during* the upgrade. However, if you do replace just one module during an upgrade, you must replace any remaining EX9200-SF modules with EX9200-SF2 modules before you begin using the switch for normal operations—that is, for the switch to function properly, it must contain only one type of SF module. [See [Upgrading an EX9200-SF to an EX9200-SF2.](#)]

Flow Monitoring

- **EX9200 Virtual Chassis support for inline flow monitoring (EX9200 Virtual Chassis)**—Starting with Junos OS Release 14.1R3, you can configure inline flow monitoring for an EX9200 Virtual Chassis. Inline flow monitoring enables you to actively monitor the flow of traffic by means of a switch participating in the network. Inline flow monitoring for an EX9200 Virtual Chassis provides support for the following:
 - Active sampling and exporting of both IPv4 and IPv6 traffic flows
 - Sampling traffic flows in both ingress and egress directions
 - Configuring flow collection on either IPv4 or IPv6 devices
 - Using the IPFIX flow collection template for traffic sampling (both IPv4 and IPv6 export records)

Basic configuration of inline flow monitoring on the EX9200 Virtual Chassis comprises these steps:

1. Enable the Virtual Chassis configuration. See [Configuring an EX9200 Virtual Chassis](#).
2. Set up the export version, family to be sampled, and template to be used at the **[edit forwarding-options]** hierarchy level.
3. Configure the IPFIX template at the **[edit services flow-monitoring]** hierarchy level.
4. Configure the firewall term at the **[edit firewall]** hierarchy level.

5. Associate the sampling instance to the master or backup switch and the corresponding FPC slot at the `[edit chassis member member-number fpc slot slot-number]` hierarchy level.
6. Associate the firewall term with the interface on which you have enabled ingress or egress sampling.

Platform and Infrastructure

- **Allow DHCP clients to send packets without Option 255 (EX9200)**—On EX9200 switches, starting with Junos OS Release 14.1, you can configure DHCP relay to enable clients to send DHCP packets without Option 255 (end-of-options). The default behavior in Junos OS is to drop packets that do not include Option 255. To override this default behavior, you can configure the **allow-no-end-options** CLI statement under the `[edit forwarding-options dhcp-relay overrides]` hierarchy level.

Related Documentation

- [Changes in Behavior and Syntax on page 8](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 14.1R7 for the EX Series.

- [Dynamic Host Configuration Protocol](#)
- [Interfaces](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)

Dynamic Host Configuration Protocol

- **DHCP clients can send packets without Option 255 (EX9200)**—On EX Series switches, starting with Junos OS Release 14.1R4, you can override the default configuration of the DHCP local server and enable clients to send DHCP packets that do not include Option 255 (end-of-options). The default behavior in Junos OS is to drop packets that do not include Option 255. To override that default behavior, configure the **allow-no-end-options** CLI statement at the **[system services dhcp-local-server overrides]** hierarchy level.

Interfaces

- **Change to CLI options for interface speed**—On EX9200 switches, the CLI command **set interfaces *interface-name* speed auto-10m-100m** is not supported.

Network Management and Monitoring

- **New system log message indicating the difference in the Packet Forwarding Engine counter value (EX9200)**—Effective in Junos OS Release 14.1R3, if the counter value of a Packet Forwarding Engine is reported lesser than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the **MIB2D_COUNTER_DECREASING** system log message for that specific counter.

Platform and Infrastructure

- **Changes in show chassis hardware command output descriptions for EX9200 components**—Starting with Junos OS Release 14.1, the output of the **show chassis hardware** command includes descriptions for enhanced midplanes on EX9204 and EX9208 switches (enhanced midplanes are already on EX9214 switches and their descriptions included in the **show chassis hardware** command output) and the high-speed SF module EX9200-SF2, as highlighted in the following sample:

```
user@switch> show chassis hardware
Hardware inventory:
Item                Version  Part number  Serial number  Description
Chassis              JN1221A03RFC
Midplane             REV 01    750-053633   ACRA1451      EX9204-BP
FPM Board            REV 04    760-021392   ABCB4822      Front Panel Display
PEM 0                Rev 10    740-029970   QCS1251U020   PS 1.4-2.52kW; 90-264V
  AC in
PEM 1                Rev 10    740-029970   QCS1251U028   PS 1.4-2.52kW; 90-264V
  AC in
Routing Engine 0     REV 02    740-049603   9009153805    RE-S-EX9200-1800X4
Routing Engine 1     REV 02    740-049603   9009153993    RE-S-EX9200-1800X4
CB 0                 REV 08    750-048307   CAB6474       EX9200-SF2
CB 1                 REV 10    750-048307   CABH8948      EX9200-SF2
...
```

- Related Documentation**
- [New and Changed Features on page 6](#)
 - [Known Behavior on page 10](#)

- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

Known Behavior

There are no changes in known behavior in Junos OS Release 14.1R7 for the EX Series.

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

Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R7 for the EX Series.

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

- [Dynamic Host Configuration Protocol](#)
- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)
- [Routing Protocols](#)
- [Software Installation and Upgrade](#)

Dynamic Host Configuration Protocol

- On EX9200 switches, the DHCPv6 binding table as shown in the output of the **show dhcp-security ipv6 binding** might contain stale entries under the following conditions:
 1. There is a mismatch in the link local address between the link local binding and the dynamic binding.

2. There is no dynamic binding, and a SOLICIT message that matches the link local entry is received, causing the state of the IPv6 address to transition from BOUND to WAITING. This resets the lease timer and creates a stale entry.

The presence of stale entries in the DHCPv6 binding table might cause the jdncpd process to create a core file. [PR1012556](#)

Interfaces and Chassis

- On EX9200 switches, if an IRB logical interface is configured on an EX9200-6QS line card as part of a VLAN, any device connected through that interface will not be able to route outside of the subnet because EX9200-6QS drops all ARP requests. As a result, the EX9200-6QS line card drops all routed traffic, including both data plane and control plane traffic. Configuring static ARP on devices using the EX9200 as gateway is not a workaround, because the packets are still dropped if the Routing Engine of the EX9200 has the routes and ARP entry for the destination IP. The minimum supported release for the EX9200-6QS line card is Junos OS Release 14.2R1. However, if your switch configuration includes an IRB logical interface configured on an EX9200-6QS line card as part of a VLAN, as a workaround, upgrade your software to the release specified in [TSB16659](#). [PR1068396](#)
- On EX9200 switches, an IRB unicast next hop in a scenario with a Layer 2 LAG as the underlying interface might result in traffic blackholing. [PR1114540](#)

Network Management and Monitoring

- On EX9200 switches, the `ptopoConnLastVerifyTime` MIB might return an incorrect value. [PR1049860](#)

Routing Protocols

- On EX9200 switches, if a session is initiated with an unconfigured peer, and the peer AS is a member of a confederation, then an RPD core file is created. As a workaround, use an explicitly configured peer in the confederation ASs. [PR963565](#)

Software Installation and Upgrade

- On EX9200 switches, if you issue a unified ISSU from Junos OS Release 13.2 or earlier to Junos OS Release 14.1 or later, approximately 20 seconds of traffic drop occurs for IPv6 protocols (for example, OSPFv6, BGPv6, or RIPv6) that are enabled on integrated routing and bridging (IRB) interfaces. The problem occurs because different link local addresses have been generated in the two releases for the same IRB logical units. As a workaround, configure different local IPv6 interfaces for different IRB logical units. [PR1086775](#)

Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)
- [Known Behavior on page 10](#)
- [Resolved Issues on page 12](#)

- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

- [Resolved Issues: Release 14.1R7 on page 12](#)
- [Resolved Issues: Release 14.1R6 on page 13](#)
- [Resolved Issues: Release 14.1R5 on page 13](#)
- [Resolved Issues: Release 14.1R4 on page 15](#)
- [Resolved Issues: Release 14.1R3 on page 16](#)
- [Resolved Issues: Release 14.1R2 on page 17](#)

Resolved Issues: Release 14.1R7

Firewall Filters

- On EX9200 switches, starting with Junos OS Release 14.1R1, **32k** is the minimum value that you must configure for policer bandwidth limits. If you configure a policer bandwidth limit that is less than **32k**, an error message is displayed. [PR1109780](#)

Infrastructure

- On an EX Series switch acting as a DHCPv6 server, the server does not send a Reply packet after receiving a Confirm packet from the client; the behavior is not compliant with the RFC3315 standard. [PR1025019](#)
- Upon BFD flapping on aggregate interfaces, the Lookup chip (XL) might send illegal packets to the center chip (XMCHIP) and compromise packet forwarding and an FPC restart is needed to recover from this condition. If Fabric path side is affected, the fabric healing process will initiate this process automatically to recover from such conditions. MPC6E/MPC5E/NG-MPC are exposed to this problem. Corrupted parcels from Lookup chip LU/XL to Center Chip (XM) can also compromise packet forwarding and report DRD parcel timeout errors. An additional parcel verification check is added to prevent sending corrupted parcels to the center chip (XM). For a possible workaround, contact JTAC for an op-script to change internal registers on MPC6E/MPC5E/NG-MPC cards. Restoration: MPC restart is needed to recover. [PR1067234](#)
- Scheduler: Protect: Parity error for tick table single messages might appear on MPC3E/MPC4E/MPC5E/MPC6E/T4000-FPC5. [PR1083959](#)

Resolved Issues: Release 14.1R6

Dynamic Host Configuration Protocol (DHCP)

- On EX9200 switches, the Dynamic Host Configuration Protocol (DHCP) relay feature, which enables the client interface and the server interface to be in separate virtual routing and forwarding (VRF) instances, does not work when the client interface has been configured as an integrated routing and bridging (IRB) interface. As a restoration workaround, configure the client interface by using **flexible-vlan-tagging** and **vlan-id** statements. [PR1064889](#)
- On EX9200 switches, if DHCP relay is configured using the **forward-only** and **forward-only-replies** statements at the **[edit forwarding-options dhcp-relay]** hierarchy level, and the DHCP local server is also configured with the **forward-snooped-clients** statement at the **[edit system services dhcp-local-server]** hierarchy level, the configuration for **forward-snooped-clients** takes precedence over the configuration for **forward-only** and **forward-only-replies**. As a result, DHCP message exchange between VRFs might not work as expected. [PR1077016](#)
- On EX9200 switches, if you configure DHCP relay with the DHCP server and the DHCP client in separate routing instances, unicast DHCP reply packets (for example, a “DHCP ACK” in response to a “DHCP RENEW”) might be dropped. [PR1079980](#)

Infrastructure

- On EX Series switches, if you change the PIM mode from sparse to dense or dense to sparse, a pfem core file might be generated. [PR1087730](#)

Network Management and Monitoring

- On EX9200 switches, if you configure an invalid SNMP source address, SNMP traps might not be sent even after you change the SNMP source address to a valid interface address. As a restoration workaround, restart the snmpd process. [PR1099802](#)

Platform and Infrastructure

- On an EX9200-2C-8XS line card, when the flow-detection feature is enabled under the **[edit system ddos-protection]** hierarchy, if suspicious control flows are received, two issues might occur on the device:
 - The suspicious control flow might not be detected on the line card.
 - After suspicious control flows are detected, they might never time out, even if traffic flows no longer violate control parameters.

[PR1102997](#)

Resolved Issues: Release 14.1R5

Authentication and Access Control

- On EX9200 switches, the output for the ptopoConnRemotePort MIB might display an incorrect value for portIDMacAddr. [PR1061073](#)

- On EX9200 switches, when clients are authenticated with dynamic VLAN assignment on an interface enabled with 802.1X authentication, disabling 802.1X authentication on the interface might cause the Layer 2 Address Learning daemon (l2ald) to generate a core file. [PR1064491](#)

Dynamic Host Configuration Protocol

- On EX9200 switches, when the switch is configured as a DHCP relay agent with option 82, and the circuit ID is configured with the CLI statement **use-interface-description** with the **device** option, then the string of the option 82 field in the DHCP DISCOVER message that is forwarded to the DHCP server should include the switch name, physical interface name, and the VLAN name. Instead, the string contains integrated routing and bridging (IRB) information in place of the physical interface name. [PR1037687](#)

Firewall Filters

- On EX9200 switches, after upgrading Junos OS to Release 14.1R1 or later, the configuration of **ipv6-payload-protocol** as a firewall filter match condition might cause the related filters to stop working. [PR1066725](#)

Interfaces and Chassis

- On EX9200 switches, when the switch receives LACP control packets from an interface other than an aggregated Ethernet (ae) interface, it forwards the packets, causing LACP peer devices that receive the packets to reset LACP connections. This might cause continuous flaps on all aggregated Ethernet interfaces and multi-chassis aggregated Ethernet interfaces. [PR1034917](#)
- On EX9200 switches, when an MC-LAG is configured with two devices, and an MC-AE interface is deleted and then recreated, broadcast and multicast traffic that is flooded might loop for several milliseconds. [PR1082775](#)

Layer 2 Features

- On EX9200 switches, if MVRP is configured on an aggregated Ethernet (AE) interface, MVRP might become unstable when the CLI command **no-attribute-length-in-pdu** is configured. [PR1053664](#)

Platform and Infrastructure

- On EX9200 switches, the **show ethernet-switching table (vlan-name | display xml)** CLI command does not have the **vlan-name** attribute in the **l2ng-l2ald-rtb-macdb** XML tag. [PR955910](#)
- On EX9200 switches, if the **disable-logging** option is the only configured option under the **[edit system ddos-protection global]** hierarchy level, the kernel might generate a core file if this option is deleted. [PR1014219](#)
- On EX9200 switches, recurring LMEM data errors might cause a chip wedge. [PR1033660](#)
- On EX9200 switches, a process that fails multiple times in a short period of time might not generate a core file. [PR1058192](#)

Routing Protocols

- On EX9200 switches on which virtual private LAN service (VPLS) is enabled and the interfaces on the CE belong to multiple FPCs, when the links between the PE device and the CE device flap, or when the administrator clears the VPLS MAC table, traffic might keep flooding in the VPLS routing instance for more than 2 seconds during the MAC learning phase. [PR1031791](#)

Software Installation and Upgrade

- Due to a software defect in Junos OS Release 14.1R5.4, we strongly discourage the use of Release 14.1R5.4 on switches that contain EX9200-40T and EX9200-40F line cards. [PR108826](#)

Virtual Private LAN Service (VPLS)

- On EX9200 switches on which VPLS is configured, the label-switched interface (LSI) is not associated with the VLAN when the VPLS connection is in the **UP** state. As a result, the switch does not flood traffic to the LSI. [PR1083561](#)
- On EX9200 switches, when you add a VLAN to an existing virtual switch routing instance for VPLS, the label-switched interface (LSI) is not associated with the added VLAN. [PR1088541](#)

Resolved Issues: Release 14.1R4

Dynamic Host Configuration Protocol

- When the DHCP relay agent receives a DHCP DISCOVER packet from a client while the client already has a binding on the relay that is in BOUND state, the client state will change to TERMINATED a stale entry is created in the Session Database (SDB). As the number of such stale entries increases, the SDB memory size might be exhausted, preventing new DHCP clients from obtaining an IP address lease. [PR1031605](#)
- On EX9200 switches, Dynamic Host Configuration Protocol (DHCP) relay functionality might stop working and DHCP will not form new bindings when the number of subscribers exceeds 1000. [PR1033921](#)

Interfaces and Chassis

- On EX9200 switches, in an MC-LAG scenario, a MAC address might incorrectly point to an inter-chassis control link (ICL) after a MAC move from a single-home LAG to the MC-LAG. [PR1034347](#)

Platform and Infrastructure

- On EX9200 switches, the **restart chassis-control** CLI command might cause loss of unicast traffic. [PR1026125](#)
- On EX9200 switches, if the switch receives an ARP packet when the Forwarding Information Base (FIB) has exceeded the limit of 262144 routes, the kernel might generate a core file. [PR1028714](#)

Spanning Tree Protocols

- On EX9200 switches running the VLAN Spanning Tree Protocol (VSTP), incoming BPDUs might not be included in the output of the **show spanning-tree statistics interface** command. [PR847405](#)

Resolved Issues: Release 14.1R3

Layer 2 Features

- On EX Series switches, an Ethernet Switching daemon (eswd) memory leak might occur in the following two conditions:
 - If a VLAN acquires a VLAN index of 0 when the VLAN is deleted, but memory is not freed accordingly.
 - In a Multiple VLAN Registration Protocol (MVRP) scenario, when a VLAN map entry is deleted, but memory is not freed accordingly.

[PR956754](#)

Interfaces and Chassis

- On EX9200 switches, virtual private LAN service (VPLS) might not work as expected, causing traffic loss. [PR993029](#)
- On EX9200 switches, in a BOOTP relay agent scenario, DHCPACK messages responding to DHCPINFORM messages might not be forwarded to the DHCP client if these ACK messages are sent from a DHCP server that is different from the DHCP server in the DHCP relay agent's binding table. [PR994735](#)
- On an EX9200 switch, if the underlying Layer 2 interface of an IRB interface is changed from access mode to trunk mode and bidirectional traffic is sent from an interface on the same switch that has been changed from IRB over Layer 2 to Layer 3 mode, the Layer 3 traffic toward the IRB interface might be dropped and PPE thread timeout errors might be displayed. As a workaround, deactivate and then reactivate the Layer 2 trunk interface underlying the IRB interface where the traffic drop occurs. [PR995845](#)
- On EX9200 switches that are configured in a multicast scenario with PIM enabled, an (S,G) discard route might stop programming if the switch receives resolve requests from an incorrect reverse-path-forwarding (RPF) interface. Once the issue occurs, the (S,G) state might not be updated when the switch receives multicast traffic from the correct RPF interfaces, and multicast traffic might be dropped. [PR1011098](#)

Platform and Infrastructure

- On EX9200 switches, when **apply-groups** is used in the configuration, the expansion of **interfaces <*> apply-groups** is done against all interfaces during the configuration validation process, even if **apply-groups** is configured only under a specific interface stanza. This issue does not affect the configuration; if the configuration validation passes, **apply-groups** is expanded only on interfaces for which **apply-groups** is configured. [PR967233](#)

Routing Protocols

- On an EX9200 switch with an IGMP configuration in which two receivers are joined to the same (S,G) and IGMP immediate-leave is configured, when one of the receivers sends a leave message for the (S,G), the other receiver might not receive traffic for one through two minutes. [PR979936](#)

Resolved Issues: Release 14.1R2

- [Interfaces and Chassis](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)

Interfaces and Chassis

- On EX9200 switches, if you configure the interface alias feature, the feature might not work as expected and interfaces might go up and down after commit. [PR981249](#)
- On EX9200 switches, the configuration statement **mcae-mac-flush** is not available in the CLI; it is missing from the **[edit vlans]** hierarchy level. [PR984393](#)
- On EX9200 switches, when the native VLAN is configured on a LAG trunk interface, if the native VLAN is modified (for example, if the native VLAN ID is changed or if the native VLAN is disabled), a Packet Forwarding Engine thread timeout might occur, a chip error message such as **fpc0 LUCHIP(1) PPE_5 Errors thread timeout error** might be displayed, and traffic might be affected. [PR993080](#)
- On EX9200 switches with multichassis link aggregation group (MC-LAG) interfaces configured, the Layer 2 address learning process (l2ald) might crash and a core file might be generated if you configure an MC-LAG interface with the **mac-rewrite** statement. [PR997978](#)

Platform and Infrastructure

- On an EX9200 switch working as a DHCP server, when you delete an IRB interface or change the VLAN ID of a VLAN corresponding with an IRB interface, the DHCP process (jdhcpd) might create a core file after commit, because a stale interface entry in the jdhcpd database has been accessed. [PR979565](#)
- On EX9200 switches, if you configure the interface alias feature, the feature might not work as expected and interfaces might go up and down after commit. [PR981249](#)

Routing Protocols

- On EX9200 switches with IGMP snooping enabled on an IRB interface, some transit TCP packets might be incorrectly handled as IGMP packets, causing packets to be dropped. [PR979671](#)

Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

Documentation Updates

There are no errata or changes in Junos OS Release 14.1R7 for the EX Series switches documentation.

Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

Migration, Upgrade, and Downgrade Instructions

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

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

- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 18](#)

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>.

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

Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Documentation Updates on page 18](#)
- [Product Compatibility on page 19](#)

Product Compatibility

- [Hardware Compatibility on page 19](#)

Hardware Compatibility

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

To determine the features supported on EX Series switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at <http://pathfinder.juniper.net/feature-explorer/>.

Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 8](#)

- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 12](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)

Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers

These release notes accompany Junos OS Release 14.1R7 for the M Series, MX Series, and T Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 14.1R7 for the M Series, MX Series, and T Series.

- [Hardware on page 22](#)
- [Authentication, Authorization and Accounting \(AAA\) \(RADIUS\) on page 26](#)
- [Class of Service \(CoS\) on page 26](#)
- [Dynamic Host Configuration Protocol \(DHCP\) on page 28](#)
- [Forwarding and Sampling on page 28](#)
- [General Routing on page 28](#)
- [High Availability \(HA\) and Resiliency on page 29](#)
- [Interfaces and Chassis on page 31](#)
- [IPv6 on page 35](#)
- [Layer 2 Features on page 36](#)
- [MPLS on page 37](#)
- [Multicast on page 38](#)
- [Network Management and Monitoring on page 38](#)
- [Network Operations and Troubleshooting Automation on page 39](#)
- [Platform and Infrastructure on page 40](#)
- [Port Security on page 40](#)
- [Routing Policy and Firewall Filters on page 41](#)
- [Routing Protocols on page 42](#)

- [Services Applications on page 43](#)
- [Software Installation and Upgrade on page 46](#)
- [Spanning-Tree Protocols on page 47](#)
- [Subscriber Management and Services on page 48](#)
- [User Interface and Configuration on page 53](#)
- [VLAN Infrastructure on page 53](#)
- [VPNs on page 54](#)

Hardware

- **Support for guided cabling (TX Matrix Plus routers with 3D SIBs)**—Junos OS Release 14.1 and later support guided cabling in a routing matrix based on a TX Matrix Plus router with 3D SIBs. When you enable guided cabling, blinking LEDs on unconnected ports help you connect cables between the TXP-F13-3D and the TXP-LCC-3D SIBs.

Use the following commands to enable or disable guided cabling:

- To enable guided cabling, use the **request chassis fabric guided-cabling (all-lcc | lcc *lcc-number*) enable (plane-by-plane | port-by-port)** operational mode command.
- To disable guided cabling, use the **request chassis fabric guided-cabling (all-lcc | lcc *lcc-number*) disable** operational mode command.

[See [Guided Cabling Overview](#), [request chassis fabric guided-cabling enable](#), and [request chassis fabric guided-cabling disable](#)]

- **Support for simultaneous BITS/BITS redundancy on SCBE2 (MX240, MX480, and MX960)**—Starting with Junos OS Release 14.1, simultaneous BITS/BITS redundancy is supported on SCBE2 on MX240, MX480, and MX960 routers. You can configure both the external interfaces for BITS input. One of the BITS inputs is considered as a primary clock source and the other as a secondary clock source on the basis of the configured clock quality.

[See [Centralized Clocking Overview](#).]

- **Unified ISSU support (TX Matrix Plus router with 3D SIBs)**—Beginning with Junos OS Release 14.1, unified in-service software upgrade (ISSU) is supported on a TX Matrix Plus router with 3D SIBs. Unified ISSU enables you to upgrade from an earlier Junos OS release to a later one with no disruption on the control plane and with minimal disruption of traffic.

[See [Unified ISSU Concepts](#)]

- **Support for OTN MIC on MPC6E (MX2010 and MX2020)**—The 24-port 10-Gigabit Ethernet OTN MIC with SFPP (MIC6-10G-OTN) is supported on MPC6E on the MX2010 and MX2020 routers. The OTN MIC supports both LAN PHY and WAN PHY framing modes on a per-port basis.

The MIC supports the following features:

- Transparent transport of 24 10-Gigabit Ethernet signals with optical channel data unit 2 (ODU2) and ODU2e framing on a per port basis

- ITU-standard optical transport network (OTN) performance monitoring and alarm management
- Pre-forward error correction (pre-FEC)-based bit error rate (BER). Fast reroute (FRR) uses the pre-FEC BER as an indication of the condition of an OTN link

To configure the OTN options for this MIC, use the **set otn-options** statement at the **[edit interfaces interfaceType-fpc/pic/port]** hierarchy level.

- **OTN support for 10-Gigabit Ethernet and 100-Gigabit Ethernet interfaces on MPC5E and MPC6E (MX240, MX480, MX960, MX2010, and MX2020)**—Junos OS Release 14.1R3 extends optical transport network (OTN) support for 10-Gigabit Ethernet and 100-Gigabit Ethernet interfaces on MPC5E and MPC6E. MPC5E-40G10G and MPC5EQ-40G10G support OTN on 10-Gigabit Ethernet interfaces, and MPC5E-100G10G and MPC5EQ-100G10G support OTN on 10-Gigabit Ethernet interfaces and 100-Gigabit Ethernet interfaces. The OTN MICs MIC6-10G-OTN and MIC6-100G-CFP2 on MPC6E support OTN on 10-Gigabit Ethernet interfaces and 100-Gigabit Ethernet interfaces, respectively.

OTN support includes:

- Transparent transport of 10-Gigabit Ethernet signals with optical channel transport unit 2 (OTU2) framing
- Transparent transport of 100-Gigabit Ethernet signals with OTU4 framing
- ITU-T standard OTN performance monitoring and alarm management

Compared with SONET/SDH, OTN provides stronger forward error correction, transparent transport of client signals, and switching scalability. To configure the OTN options for the interfaces, use the **set otn-options** configuration statement at the **[edit interfaces interfaceType-fpc/pic/port]** hierarchy level.

- **Support for fixed-configuration MPC (MX240, MX480, MX960, MX2010, and MX2020)**—MX2020, MX2010, MX960, MX480, and MX240 routers support a new MPC, MPC5E (model number: MPC5E-40G10G). On the MX2010 and MX2020 routers, MPC5E is housed in an adapter card. MPC5E is a fixed-configuration MPC with four built-in PICs and does not contain separate slots for Modular Interface Cards (MICs). MPC5E supports two Packet Forwarding Engines, PFE0 and PFE1. PFE0 hosts PIC0 and PIC2 while PFE1 hosts PIC1 and PIC3. A maximum of two PICs can be kept powered on (PIC0 or PIC2 and PIC1 or PIC3). The other PICs are required to be kept powered off.

MPC5E supports:

- Flexible queuing option by using an add-on license
- Forwarding capability of up to 130 Gbps per Packet Forwarding Engine
- Intelligent oversubscription services
- Quad small form-factor pluggable plus transceivers (QSFP+) and small form-factor pluggable plus transceivers (SFP+) for connectivity
- Up to 240 Gbps of full-duplex traffic
- WAN-PHY mode on 10-Gigabit Ethernet Interfaces on a per-port basis



NOTE: On MX960 routers, all the MPC slots work with chassis temperature of up to 40°C (104°F). However, when the chassis temperature exceeds 40°C (104°F), slots 0 and 11 can only work with MPC1, MPC2, and the 16x10GE MPC.

[See [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5E](#).]

- **Support for new fixed-configuration queuing MPC (MX240, MX480, MX960, MX2010, and MX2020)**—MX2020, MX2010, MX960, MX480, and MX240 routers support a new queuing MPC, MPC5EQ (model number: MPC5EQ-40G10G). On the MX2010 and MX2020 routers, MPC5EQ is housed in an adapter card. MPC5EQ, like MPC5E is a fixed-configuration MPC with four built-in PICs and does not contain separate slots for Modular Interface Cards (MICs). MPC5EQ, like MPC5E supports two Packet Forwarding Engines, PFE0 and PFE1. PFE0 hosts PIC0 and PIC2 while PFE1 hosts PIC1 and PIC3. A maximum of two PICs can be kept powered on (PIC0 or PIC2 and PIC1 or PIC3). The other PICs are required to be kept powered off.

MPC5EQ supports 1 million queues per slot on all MX Series routers. All the other software features supported on MPC5E are also supported on MPC5EQ.



NOTE: On the MX960 router, FPC slot 0 and FPC slot 11 are not NEBS compliant beyond 104°F (40°C). This is a cooling restriction.

[See [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5E](#).]

- **Software feature support on the MPC5E**—MPC5E supports the following key features:
 - Basic Layer 2 features and virtual private LAN services (VPLS) functionality
 - Class of service (CoS)
 - Flexible Queuing option—By using an add-on license, MPC5E supports a limited number of queues (32,000 queues per slot including ingress and egress)
 - Hierarchical QoS
 - Intelligent oversubscription services
 - Interoperability with existing MPCs and DPCs
 - MPLS
 - MX Virtual Chassis

The following features are not supported on MPC5E:

- Active flow monitoring and services
- Subscriber management features

[See [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5E](#).]

- **Software feature support on the MPC5EQ**—MPC5EQ supports 1 million queues per slot on all MX Series routers. All the other software features supported on MPC5E are also supported on MPC5EQ.

[See [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5E.](#)]

- **Support for new 520-gigabit full duplex Modular Port Concentrator (MPC6E) with two Modular Interface Card (MIC) slots (MX2010 and MX2020)**—The MX2020 and MX2010 routers support a new MPC, MPC6E (model number: MX2K-MPC6E). MPC6E is a 100-Gigabit Ethernet MPC that provides increased density and performance to MX Series routers in broadband access networks for services such as Layer 3 peering, VPLS and Layer 3 aggregation, and video distribution.

MPC6E provides packet-forwarding services that deliver up to 520 Gbps of full-duplex traffic. It has two separate slots for MICs and supports four Packet Forwarding Engines with a throughput of 130 Gbps per Packet Forwarding Engine. It also supports two MIC slots as WAN ports that provide physical interface flexibility.

MPC6E supports:

- 100-Gigabit Ethernet interfaces
- Forwarding capability of up to 130 Gbps per Packet Forwarding Engine
- Intelligent oversubscription services
- Two separate slots for MICs (MIC6-10G and MIC6-100G-CXP)
- Two Packet Forwarding Engines for each MIC slot
- Up to 560 Gbps of full-duplex traffic for the two MIC slots
- WAN-PHY mode on 10-Gigabit Ethernet interfaces on a per port basis

[See [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5E.](#)]

- **Feature support on MPC6E**—MPC6E supports the following software features:
 - Basic Layer 2 features and virtual private LAN service (VPLS) functionality, except for Operation, Administration, and Maintenance (OAM)
 - Class of service (CoS)
 - Firewall filters and policers
 - Intelligent hierarchical policers
 - Internet Group Management Protocol (IGMP) snooping with bridging, integrated routing and bridging (IRB), or VPLS
 - Interoperability with existing DPCs and MPCs
 - Layer 2 trunk port
 - Layer 3 routing protocols
 - Multicast forwarding
 - MPLS

- MPLS-fast reroute (FRR) VPLS instance prioritization
- Precision Time Protocol (PTP) (IEEE 1588)
- Synchronous Ethernet
- Tunnel service

The following features are not supported on MPC6E:

- Active flow monitoring and services
- Fine-grained queuing and input queuing
- Unified in-service software upgrade (ISSU)
- Virtual Chassis support

[See [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5E.](#)]

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **RADIUS functionality over IPv6 for system AAA**—Starting in Release 14.1R2, Junos OS supports RADIUS functionality over IPv6 for system AAA (authentication, authorization, and accounting) in addition to the existing RADIUS functionality over IPv4 for system AAA. With this feature, Junos OS users can log in to the router authenticated through RADIUS over an IPv6 network. Thus, Junos OS users can now configure both IPv4 and IPv6 RADIUS servers for AAA. To accept the IPv6 source address, include the **source-address-inet6** statement at the **[edit system radius-server IPv6]** hierarchy level. (Note that if an IPv6 RADIUS server is configured without any **source-address**, default ::0 is considered as the source address.)

Class of Service (CoS)

- **Distributed periodic packet management support for aggregated Ethernet interfaces (T4000)**—Starting with Release 14.1, Junos OS extends support on T4000 routers for the Bidirectional Forwarding Detection (BFD) protocol to use the periodic packet management daemon (ppmd) to distribute IPv4 sessions over aggregated Ethernet interfaces. Only IPv4 BFD sessions over aggregated Ethernet interfaces are supported. The ppm process automatically runs on the Routing Engine and the Packet Forwarding Engine. To disable ppm on the Packet Forwarding Engine only, include the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level. The ppm process does not support IPv6 BFD sessions or MPLS BFD sessions over an aggregated Ethernet interface.

[See [ppm](#) and [no-delegate-processing](#).]

- **Support for limiting traffic black-hole time by detecting Packet Forwarding Engine destinations that are unreachable (T4000)**—Junos OS Release 14.1 and later releases extend support for T4000 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. By default, the system limits traffic black-hole time by detecting severely degraded fabric. No user interaction is necessary.

[See [Traffic Blackholing Caused by Fabric Degradation](#), [Disabling FPC Restart](#), [degraded, action-fpc-restart-disable](#), [show chassis fabric reachability](#), and [show chassis fabric unreachable-destinations](#).]

- **Setting IPv4 and IPv6 DSCP and MPLS EXP bits independently (T4000 and TXP-4000-3D)**—Junos OS Release 14.1 and later releases extend support to set the packet DSCP and MPLS EXP bits independently on IPv4 and IPv6 packets on T4000 Type 5 FPCs (model numbers: T4000-FBC5-3D and T4000-FPC5-LSR) in T4000 routers and the TXP-4000-3D chassis. To enable this feature for IPv4, include the **protocol mpls** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules dscp *rewrite-name*]** hierarchy level. To enable this feature for IPv6, include the **protocol mpls** statement at the **[edit class-of-service interfaces *interface-name* unit *logical-unit-number* rewrite-rules dscp-ipv6 *rewrite-name*]** hierarchy level. You can set DSCP IPv4 and IPv6 values only at the ingress MPLS node. The following rewrite combinations are supported:
 - DSCP or inet-precedence and EXP rewrite on IPv4 packets
 - DSCPv6 and EXP rewrite on IPv6 packets

[See [Applying Rewrite Rules to Output Logical Interfaces](#), [Setting IPv6 DSCP and MPLS EXP Values Independently](#), [Configuring DSCP Values for IPv6 Packets Entering the MPLS Tunnel](#), and [Configuring Rewrite Rules](#).]

- **Layer 2 CoS-based traffic metering (MX80 and MX Series routers with MPCs)**—Starting with Junos OS Release 14.1, Layer 2 accounting statistics are available on a per class-of-service basis. Both bytes and packet total are counted (flow rates are not).

A single, aggregate counter can be used with each forwarding class to count inet and inet6 flows. For ingress, only packets forwarded to the fabric are counted, and for egress, only packets forwarded to the WAN are counted. You can exclude overhead bytes from the count, as well as dropped packets and non-relevant network protocols such as ARP, BFD, and EOAM. You can configure counters with any or all of the following parameters:

- logical/physical interfaces
- IPv4/IPv6 traffic types
- unicast/multicast traffic
- ingress/egress flows

Configure the counters using the **enhanced** command under **forwarding-class-accounting** in the CLI.

- **Support for CoS hierarchical schedulers on MPC5E (MX240, MX480, MX960, MX2010, and MX2020)**—You can configure class-of-service (CoS) hierarchical schedulers on MPC5E interfaces. This feature is supported on egress only.

You can use hierarchical schedulers to define traffic control profiles, which set the following CoS parameters on a CoS interface:

- Delay buffer rate
- Excess bandwidth
- Guaranteed rate
- Overhead accounting
- Scheduler map
- Shaping rate

Dynamic Host Configuration Protocol (DHCP)

- **Recursive DNS server ICMPv6 router advertisement option support (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 14.1, you can configure a maximum of three recursive DNS server addresses and their respective lifetimes through static configuration at the interface level for IPv6 hosts. Previously, rpd supported only link-local address information, prefix information, and the link MTU. The router advertisement-based DNS configuration is useful in networks where an IPv6 host's address is auto-configured through an IPv6 stateless address and where there is no DHCPv6 infrastructure available.

To configure the recursive DNS server address, include the **dns-server-address** statement at the **[edit protocols router-advertisement interface *interface-name*]** hierarchy level.

[See [Example: Configuring Recursive DNS Address.](#)]

Forwarding and Sampling

- **Native analyzer support (MX240, MX480, and MX960)**—Starting in Junos OS Release 14.1, support is provided for native analyzers and remote port-mirroring capabilities on the MX240, MX480, and MX960. A native analyzer configuration contains both an input stanza and an output stanza in the analyzer hierarchy for mirroring packets. In remote port mirroring, the mirrored traffic is flooded into a remote mirroring VLAN that can be specifically created for the purpose of receiving mirrored traffic. The analyzer configuration is available at the **[edit-forwarding-options]** hierarchy level.

General Routing

- **Updated behavior in static link protection Mode (M Series, MX Series, and T Series)**—In static link protection mode you can designate a primary and backup physical link to support aggregated interfaces link protection. Starting with Junos OS Release 14.1, you can configure a backup link to accept ingress traffic, discard ingress traffic, or remain down until it becomes active and starts carrying traffic. By default, the backup link accepts ingress traffic. The following new attributes have been added to **link-protection** to control these settings:
 - **bkp-state-accept**: Default, accept ingress traffic on the backup link
 - **bkp-state-discard**: Discard ingress traffic on the backup link
 - **bkp-state-down**: Mark the backup link as Down while the primary link is active

- **Support for preserving prenormalized ToS value in an egress mirrored or sampled packet (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 14.1, on MPC-based interfaces, you can preserve the prenormalized type-of-service (ToS) value for egress mirrored or sampled packets. To retain the pre-rewrite ToS value in mirrored or sampled packets, configure the **pre-rewrite-tos** statement at the **[edit forwarding-options sampling]** hierarchy level. This preserves the pre-rewrite ToS value for all forms of sampling, such as Routing Engine-based sampling, port mirroring, flow monitoring, and so on. This statement is effective for egress sampling only.

High Availability (HA) and Resiliency

- **MX Series Virtual Chassis support for determining member router health (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, you can configure an IP-based packet connection, known as a *heartbeat connection*, between the master router and backup router in an MX Series Virtual Chassis. The heartbeat connection exchanges *heartbeat packets* that provide important information about the availability and health of each member router.

If a disruption or split occurs in the Virtual Chassis configuration, the heartbeat connection helps prevent the member routers from changing roles, which could cause undesirable results.

To configure a heartbeat connection, first create a secure and reliable route between the master router and backup router. You can then configure the connection by including the **heartbeat-address** and **heartbeat-timeout** statements at the **[edit virtual-chassis]** hierarchy level.

- **MX Series Virtual Chassis support for locality bias (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, you can configure *locality bias* for aggregated Ethernet and equal-cost multipath (ECMP) traffic in an MX Series Virtual Chassis. Locality bias directs unicast transit traffic for ECMP groups and aggregated Ethernet bundles to egress links in the same (local) member router in the Virtual Chassis rather than to egress links in the remote member router, provided that the local member router has an equal or larger number of available egress links than the remote member router.

Configuring locality bias enables you to conserve bandwidth on the Virtual Chassis port links by directing all ECMP and aggregated Ethernet data traffic to local egress links rather than across the Virtual Chassis port links between member routers.

To enable locality bias, configure the **locality-bias** statement at the **[edit virtual-chassis]** hierarchy level.



BEST PRACTICE: To avoid possible traffic loss and oversubscription on egress interfaces, make sure that you understand the utilization requirements for the local links in your network before changing the locality bias configuration.

- **MX Series Virtual Chassis support for unified ISSU (MX Series routers with MPCs/MICs)**—Starting in Junos OS Release 14.1, you can perform a unified in-service software upgrade (unified ISSU) on member routers in an MX Series Virtual Chassis

configuration. Unified ISSU enables you to upgrade the the system software on the Virtual Chassis member routers with minimal traffic disruption and no disruption on the control plane.

To start a unified ISSU in an MX Series Virtual Chassis, issue the **request system software in-service-upgrade package-name** command from the master Routing Engine in the Virtual Chassis master router (VC-Mm). This command always reboots each of the four Routing Engines in the Virtual Chassis.

[See [Unified ISSU in a Virtual Chassis](#) and [Unified ISSU System Requirements](#).]

- **MX Series Virtual Chassis support for Layer 2 spanning-tree protocols (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, an MX Series Virtual Chassis configuration supports the following Layer 2 Control Protocol (L2CP) features, known collectively as xSTP:
 - Spanning Tree Protocol (STP)
 - Rapid Spanning Tree Protocol (RSTP)
 - Multiple Spanning Tree Protocol (MSTP)
 - VLAN Spanning Tree Protocol (VSTP)

Spanning-tree protocols resolve the forwarding loops in a Layer 2 network, thereby creating a loop-free tree topology. Configuring spanning-tree protocols provides link redundancy in case of link failures, and prevents undesirable loops in the data path.

To configure and manage STP, RSTP, MSTP, or VSTP in a Virtual Chassis, you use the same procedures for a member router in an MX Series Virtual Chassis as you do for a standalone MX Series router.

[See [Spanning-Tree Protocols Supported](#) and [Virtual Chassis Components Overview](#).]

- **MX Series Virtual Chassis support for inline flow monitoring (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, you can configure inline flow monitoring for an MX Series Virtual Chassis. Inline flow monitoring enables you to actively monitor the flow of traffic by means of a router participating in the network.

Inline flow monitoring for an MX Series Virtual Chassis provides the following support:

- Active sampling and exporting of both IPv4 and IPv6 traffic flows
 - Sampling traffic flows in both the ingress and egress directions
 - Configuration of flow collection on either IPv4 or IPv6 devices
 - Use of the IPFIX flow collection template for traffic sampling (both IPv4 and IPv6 export records)
- **Support for LACP with fast hellos during unified ISSU (MX Series)**—Starting in Junos OS Release 14.1, MX Series routers support LACP with fast hellos during unified ISSU. This support is disabled by default. To enable it you need to enter the new CLI **set protocols lacp fast-hello-issu** on both the DUT and peer routers before starting unified ISSU. The peer router must also be an MX Series router for this functionality to work.

Interfaces and Chassis

- **Support for physical interface damping (T Series)**—Beginning with Junos OS Release 14.1, interface damping is supported on *physical interfaces* to address longer periodic flapping lasting 5 seconds or more, with an up and down duration of 1 second. This damping method limits the number of advertisements of longer interface up and down events to the upper-level protocols. For longer periodic interface flaps, configure interface damping with the **damping** statement at the **[edit interfaces *interface-name*]** hierarchy level. You use the **show interfaces extensive** command to view the interface damping values and link state.

[See [Damping Longer Physical Interface Transitions.](#)]

- **Support for MC-LAG on logical systems**—Starting with Junos OS Release 14.1, you can configure multichassis link aggregation (MC-LAG) interfaces on logical systems within a router. To configure ICCP for MC-LAG interfaces on logical systems, include the **iccp** statement at the **[edit logical-systems *logical-system-name* protocols]** hierarchy level. To view ICCP information for MC-LAG on logical systems, use the **show iccp logical-system *logical-system-name*** command. To view ARP statistics or remote MAC addresses for the multichassis aggregated Ethernet (MC-AE) nodes for all or specified redundancy groups on a logical system, use the **show l2-learning redundancy-groups *group-name* logical-system *logical-system-name* (arp-statistics | remote-macs)** command. To view neighbor discovery statistical details for MC-AE nodes on redundancy groups of a logical group, use the **show l2-learning redundancy-groups *group-name* logical-system *logical-system-name* nd-statistics** command.

[See [Multichassis Link Aggregation on Logical Systems Overview.](#)]

- **Inline Multilink PPP, Multilink Frame Relay, and Multilink Frame Relay End-to-End for time-division multiplexing WAN interfaces (MX Series)**— Starting in Junos OS Release 14.1, this feature allows support of Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC).

Traditionally, bundling services are used to bundle multiple low-speed links to create a higher bandwidth pipe. This combined bandwidth is available to traffic from all links and supports link fragmentation and interleaving (LFI) on the bundle, reducing high priority packet transmission delay.

This support includes multiple links on the same bundle as well as multiclass extension for MLPPP. Through this service you can enable bundling services without additional DPC slots to support Service DPC and free up the slots for other MICs.

For connecting many smaller sites in VPNs, bundling the TDM circuits with MLPPP/MLFR technology is the *only* way to offer higher bandwidth and link redundancy.

MLPPP enables you to bundle multiple PPP links into a single multilink bundle, and MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

[See [Inline MLPPP for WAN Interfaces Overview](#), [Example: Configuring Inline MLPPP and Multilink Frame Relay End-to-End \(FRF.15\) for WAN Interfaces](#), and [Example: Configuring Inline Multilink Frame Relay \(FRF.16\) for WAN Interfaces](#).]

- **SFPP-10G-CT50-ZR (MX Series)**—Starting in Junos OS Release 14.1, the SPFF-10G-CT50-ZR tunable transceiver provides a duplex LC connector and supports the 10GBASE-Z optical interface specification and monitoring. The transceiver is not specified as part of the 10-Gigabit Ethernet standard and is instead built according to Juniper Networks specifications. Only WAN-PHY and LAN-PHY modes are supported. To configure the wavelength on the transceiver, use the **wavelength** statement at the **[edit interfaces interface-name optics-options]** hierarchy level. The following interface modules support the SPFF-10G-CT50-ZR transceiver:

MX Series:

- 16-port 10-Gigabit Ethernet MPC (model number: MPC-3D-16XGE-SFPP)—Supported in Junos OS Release 12.3R6, 13.2R3, 13.3R2, 14.1, and later.

For more information about interface modules, see the “Cables and Connectors” section in the *Interface Module Reference* for your router.

[See [10-Gigabit Ethernet 10GBASE Optical Interface Specifications](#) and [wavelength](#).]

- **SFPP-10G-ZR-OTN-XT (MX Series, T1600, and T4000)**—Starting in Junos OS Release 14.1, the SFPP-10G-ZR-OTN-XT dual-rate extended temperature transceiver provides a duplex LC connector and supports the 10GBASE-Z optical interface specification and monitoring. The transceiver is not specified as part of the 10-Gigabit Ethernet standard and is instead built according to ITU-T and Juniper specifications. In addition, the transceiver supports LAN-PHY and WAN-PHY modes and OTN rates and provides a NEBS-compliant 10-Gigabit Ethernet ZR transceiver for the MX Series interface modules listed here. The following interface modules support the SFPP-10G-ZR-OTN-XT transceiver:

MX Series:

- 10-Gigabit Ethernet MIC with SFP+ (model number: MIC3-3D-10XGE-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later
- 16-port 10-Gigabit Ethernet (model number: MPC-3D-16XGE-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later
- 32-port 10-Gigabit Ethernet MPC4E (model number: MPC4E-3D-32XGE-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later
- 2-port 100-Gigabit Ethernet + 8-port 10-Gigabit Ethernet MPC4E (model number: MPC4E-3D-2CGE-8XGE)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later

T1600 and T4000 routers:

- 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (model numbers: PD-5-10XGE-SFPP and PF-24XGE-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later

- 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number: PF-12XGE-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later

For more information about interface modules, see the “Cables and Connectors” section in the *Interface Module Reference* for your router.

[See [10-Gigabit Ethernet 10GBASE Optical Interface Specifications](#).]

- **Support for mixed rates on an aggregated Ethernet bundle (MX Series)**—Starting with Junos OS Release 14.1R2, support for mixed rates on aggregated Ethernet bundles is extended to MX240, MX480, MX960, MX2010, and MX2020 routers, thereby enabling you to configure the member links with any combination of rates—10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet—on an aggregated Ethernet bundle.
- **Source class accounting (T4000)**—Starting with Junos OS Release 14.1R2, the source class accounting is performed at the ingress on a T4000 Type 5 FPC in T4000 routers.
- **Support for MPC5E on SCBE2 (MX Series)**—Starting with Junos OS Release 14.1R2, MPC5E is supported on SCBE2 on MX240, MX480, and MX960 routers.
- **New command to set the license mode for MPCs (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 14.1R2, you can set the license mode for enhanced MPCs such as MPC4E, MPC5E, and MPC6E by including the **ir-mode** configuration statement at the **[edit chassis fpc]** hierarchy level. Setting the license mode enables you to distinguish between an MPC with an IR license and an MPC with an R license after the MPC is installed on the router.



NOTE: You cannot set or alter the license of the MPC when you configure the mode. The license mode settings are used only to provide information.

The license mode settings are set per slot. If the MPC is installed on a different slot, or moved to another device, the license mode settings must be re-configured on the new slot or device. Also, the license mode settings configured on the previous slot must be removed. To view the current license mode settings, as well as the effect of the new settings, use the **show chassis fpc** and **show chassis hardware extensive** commands. To delete the license mode settings, use the **delete chassis fpc** command.

- **Loop prevention in VPLS network due to MAC moves (MX Series)**—Starting with Junos OS Release 14.1R2, you can use the base learning interface approach and the statistical approach to prevent a loop in a VPLS network by disabling the suspect customer-facing interface that is connected to the loop. Some virtual MACs can genuinely move between different interfaces and you can configure such MACs to disregard the moves. The cooloff time and statistical approach wait time are used internally to find out the looped interface. You can configure the interface recovery time to auto-enable the interface that gets disabled due to a loop in the network. To configure these parameters of VPLS MAC moves, include the **vpls-mac-move** statement at the **[edit protocols l2-learning]** hierarchy level. The **show vpls mac-move-action instance instance-name** command displays the learning interfaces that are disabled, in a VPLS instance due to a MAC move. The **clear vpls mac-move-action interface ifl-name** command enables an interface disabled due to a MAC move.

- **Entropy label support in mixed mode (MX Series)**—Beginning with Junos OS Release 14.1R2, the entropy label supported in mixed mode for chassis. MX Series 3D Universal Edge Router DPCs support the pop out entropy label but do not support the flow label. You can configure the entropy label without enhanced-ip configuration.
- **Support for Synchronous Ethernet on MPC5E and MPC6E (MX240, MX480, MX960, MX2010, and MX2020 routers)**—Junos OS extends Synchronous Ethernet support for MPC5E and MPC6E on the MX240, MX480, MX960, MX2010, and MX2020 routers. MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, MPC5EQ-100G10G, and MX2K-MPC6E support Ethernet Synchronization Message Channel (ESMC) and external clocking.

To configure Synchronous Ethernet, include the **synchronization** statement and its substatements at the **[edit chassis]** hierarchy level.

- **Support for ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on aggregated Ethernet interfaces (MX Series routers with MPCs)**—Starting with Junos OS Release 14.1R4, you can configure ITU-T Y.1731 standard-compliant Ethernet loss measurement (ETH-LM), Ethernet synthetic loss measurement (ETH-SLM), and Ethernet delay measurement (ETH-DM) capabilities on aggregated Ethernet interfaces. These performance monitoring functionalities are supported on MX Series routers with MPCs, where the same level of support for the Ethernet services OAM mechanisms as the level of support on non-aggregated Ethernet interfaces is available on aggregated Ethernet interfaces. ETH-DM is supported on MPC3E and MPC4E modules with only software timestamping. ETH-SLM is supported on MPC3E and MPC4E modules.
- **Logical interfaces summary (MX Series)**—Beginning with Junos OS Release 14.1R2, a new show command, **show interfaces summary**, is available to display the status and statistics on the logical interfaces configured on the device at the Flexible PIC Concentrator (FPC) level.

[See [show interfaces summary](#).]

- **Support for fabric black-hole detection and recovery (TX Matrix Plus) router**—Starting in Junos OS Release 14.1R6, TX Matrix Plus routers can detect and recover from fabric faults that are not caused by hardware failure but might be a result of a fabric black-hole condition.

To recover from a fabric black-hole condition, the routing matrix uses the following options:

- Destination reprogramming
- FPC reboot
- Related faults recovery
- SIB reboot

You can disable the automatic recovery feature by using the **auto-recovery-disable** statement at the **[edit chassis fabric degraded]** hierarchy level. You can configure the FPCs to go offline when a traffic black-hole condition is detected in the routing matrix by using the **fpc-offline-on-blackholing** statement at the **[edit chassis fabric degraded]** hierarchy level.

You can configure the FPCs to restart when a traffic black-hole condition is detected in the routing matrix by using the **fpc-restart** statement at the **[edit chassis fabric degraded]** hierarchy level.

[See [auto-recovery-disable](#) and [fpc-offline-on-blackholing](#).]

IPv6

- **Expanded ALG support with NAT64 (MX Series routers with MS-MPC or MS-MIC line cards)**—Starting with Junos OS Release 14.1, the FTP, TFTP, SIP, RTSP, and PPPT ALGs are supported. To configure the ALGs, include the **applications [applications-list]** statement at the **[edit services nat rule rule-name term termname from]** hierarchy level.

Include in the ALG list, *applications-list*, Junos OS identifiers for desired ALGs:

- **junos-ftp** for FTP
 - **junos-tftp** for TFTP
 - **junos-sip** for SIP
 - **junos-rtsp** for RTSP
 - **junos-pppt** for PPPT
- **Limit software flows per IPv6 prefix for DS-Lite (MX Series routers with MS-DPC interface cards)**—Junos OS provides a configurable option to limit the number of software flows from a subscriber's Basic Bridging Broadband (B4) device at a given point in time, thus limiting excessive use of addresses within the subnet available to a subscriber. This limitation reduces the risk of denial-of-service (DoS) attacks.

To specify the size of the subnet subject to limitation, include the **dslite-ipv6-prefix-length prefix-length** statement at the **[edit services service-set service-set-name software-options]** hierarchy level. Specify a prefix length of 56, 64, 98, or 128.

Starting in Junos OS Release 14.1, the **show services nat mappings address-pooling-paired** operational command output shows the mapping for the prefix. The mapping shows the address of the active B4.

The **show services software flows** output shows active and inactive software flows from the same prefix.

- **Support for hyper mode to increase packet processing rate on enhanced MPCs (MX240, MX480, MX960, MX2010, and MX2020)**—In Junos OS Releases 13.3R4 and 14.1R4, MPC3E, MPC4E, MPC5E, and MPC6E support the **hyper-mode** feature. Enabling the hyper mode feature increases the rate at which a data packet is processed, which results in the optimization of the lifetime of a data packet. Optimization of the data packet lifetime enables better performance and throughput.



NOTE: You can enable hyper mode only if the network-service mode on the router is configured as either **enhanced-ip** or **enhanced-ethernet**. Also, you cannot enable the hyper mode feature for a specific Packet Forwarding Engine on an MPC—that is, when you enable the feature, it is applicable for all Packet Forwarding Engines on the router.

When you enable the hyper mode feature, the following features are not supported:

- Creation of Virtual Chassis.
- Interoperability with legacy DPCs, including MS-DPCs. The MPC in hyper mode accepts and transmits data packets only from other existing MPCs.
- Interoperability with non-Ethernet MICs and non-Ethernet Interfaces such as channelized interfaces, multilink interfaces, and SONET interfaces.
- Padding of Ethernet Frames with VLAN.
- Sending Internet Control Message Protocol (ICMP) redirect messages.
- Termination or tunneling of all subscriber-based services.

To configure the hyper mode feature, use the **hyper-mode** statement at the **[edit forwarding-options]** hierarchy level. To view the changed configuration, use the **show forwarding-options hyper-mode** command.

Layer 2 Features

- **Support for configuring PPP NCP negotiation mode (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, both static and dynamic subscriber interfaces use passive PPP NCP negotiation by default. To enable active negotiation, use the new **initiate-ncp** configuration statement with the appropriate option:
 - For IPv4 (**inet** family) subscriber interfaces, use the **initiate-ncp ip** statement.
 - For IPv6 (**inet6** family) subscriber interfaces, use the **initiate-ncp ipv6** statement.

You can also configure the negotiation mode for the PPP server in an IPv4/IPv6 dual-stack configuration:

- For active negotiation, use the **initiate-ncp ip** statement for the IPv4 subscriber interface and the **initiate-ncp ipv6** statement for the IPv6 subscriber interface.
- For passive negotiation, use the **initiate-ncp dual-stack-passive** statement, which overrides the **initiate-ncp ip** and **initiate-ncp ipv6** statements if they are configured.

[See [PPP Network Control Protocol Negotiation Mode Overview](#).]

- **Global configuration for LAC interoperoperation using Cisco NAS Port Info AVP (MX Series)**—Starting in Junos OS Release 14.1, you can globally configure LAC interoperoperation with a Cisco Systems LNS by specifying the LAC's NAS port method as **cisco-avp** with the **nas-port** statement at the **[edit services l2tp tunnel]** hierarchy level. This causes the LAC to include the Cisco NAS Port Info AVP (100) in the ICRQ messages it sends to the LNS for all tunnels.

In earlier releases, you can configure interoperation only in a tunnel profile, so that it applies only to tunnels instantiated with that profile. The tunnel profile configuration now has precedence over the global configuration. You can override both by including the Tunnel-Nas-Port-Method VSA [26–30] in a RADIUS server configuration that modifies or creates a tunnel profile.

[See [Globally Configuring the LAC to Interoperate with Cisco LNS Devices.](#)]

- **Enhanced support for firewall filter match conditions based on IEEE 802.1p VLAN priority bits (M320 and MX Series)**—Starting in Junos OS Release 14.1, the M320 router supports firewall filter match conditions based on IEEE 802.1p VLAN priority bits. The M320 router also supports these match conditions with the presence of a control word in a VPLS instance. Also starting with Junos OS Release 14.1, MX Series routers support firewall filter match conditions based on IEEE 802.1p VLAN priority bits in both a VPLS instance and a Layer 2 VPN instance.

[See [Firewall Filter Match Conditions for VPLS Traffic](#) and [Firewall Filter Match Conditions for Layer 2 CCC Traffic.](#)]

MPLS

- **LSP selection for default forwarding class using CBF (M Series, MX Series, and T Series)**—When CoS-based forwarding (CBF) is configured on a VPLS PE router, VPLS BUM traffic (broadcast, unknown, and multicast traffic) uses the default forwarding class for label-switched path (LSP) selection. Starting in Junos OS Release 14.1, the LSP for the default forwarding class is configurable, enabling the association of VPLS BUM traffic with an LSP through CBF configuration.

[See [Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface.](#)]

- **Support for load balancing VPLS non-unicast traffic across member links of an aggregate interface (M Series, MX Series, and T Series)**—By default, VPLS non-unicast (or BUM — broadcast, unknown, and multicast) traffic sent across aggregate Ethernet interfaces is sent across only one member link of the aggregate interface. Starting in Junos OS Release 14.1, load balancing VPLS BUM traffic across all members of an aggregate interface can be enabled for each VPLS instance.

[See [Load Balancing VPLS Non-Unicast Traffic Across Member Links of an Aggregate Interface.](#)]

- **Entropy label and FAT label support (MX Series and T Series)**—Starting in Release 14.1, Junos OS supports entropy labels and Flow Aware Transport for Pseudowires (FAT) labels. Entropy labels and FAT labels when configured on the label-switching routers (LSRs) and label edge routers (LERs) perform load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload.

In Junos OS Release 14.1, entropy labels can be used for RSVP-signaled label-switched paths (LSPs) and point-to-point LDP-signaled LSPs. FAT flow labels can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for virtual private LAN service (VPLS) and virtual private wire service (VPWS) networks.

[See [Configuring the Entropy Label for LSPs](#) and [FAT Flow Labels Overview.](#)]

Multicast

- **Multicast-only fast reroute (MoFRR) (MX Series)**—Starting in Junos OS Release 14.1, MoFRR functionality is available, in which packet loss is minimized in PIM and multipoint LDP domains. This enhancement is available on the MX Series operating in enhanced IP mode and with MPC line cards. A new configuration statement, **stream-protection**, enables MoFRR. When establishing the primary and backup ECMPs, MoFRR attempts to select two separate upstream routers, if two such routers are available. If separate upstream routers are not available, but there are two links through the same upstream router, the protocol selects that router for both paths.



NOTE: MoFRR might select the same upstream router to establish the primary and the backup paths, even when two separate upstream routers are available.

[See [Example: Configuring Multicast-Only Fast Reroute in a PIM Domain](#) and [Example: Configuring Multicast-Only Fast Reroute in a Multipoint LDP Domain](#).]

Network Management and Monitoring

- **Configuring SNMP to match jnxNatObjects values for MS-DPC and MS-MIC (MX Series)**—In Junos OS Release 14.1R6, you can configure the **snmp-value-match-msmic** statement at the **[edit services service-set service-set-name nat-options]** hierarchy level.

In networks where both MS-DPC and MS-MIC are deployed, you can configure this statement to ensure that the values for MS-MIC-specific objects in the **jnxNatObjects** MIB table match the values for MS-DPC objects. By default, this feature is disabled. You can use the **deactivate services service-set service-set-name nat-options snmp-value-match-msmic** configuration mode command to disable this feature.

- **Forwarding Class extension to Interface MIB (MX Series)**—Beginning with Junos OS Release 14.1, a new enterprise-specific forwarding class MIB, **jnxIfAccountingStats**, is available to monitor the statistics for various accounting parameters configured on the interface with the available forwarding classes. This is an extension to the *Enterprise-Specific Interface MIB*. The Forwarding Class MIB is currently supported only on the MX Series.

[See [Interpreting the Enterprise-Specific Interface Accounting Forwarding MIB](#).]

- **SNMP notifying target for removed notify target configuration (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 14.1, when a trap target is deleted from Juniper Networks devices, either a syslog event or a syslog trap is generated as per the user configuration. The existing SNMP trap **jnxSyslogTrap** is sent to all target network management systems (NMSs) specified in the SNMP agent including the target NMS, which is being deleted. By default, in the event of target deletion, only a syslog event is generated. To trigger a trap on deletion of a trap target, configure a syslog event policy, which sends the syslog as a trap to the network management systems.

- **Alarm MIB support (MX Series)**—Beginning with Release 14.1, Junos OS supports RFC 3877, *Alarm MIB*, which provides the generic SNMP-based alarm management framework to address the problems occurring on a particular network resource. The `jnxAlarmMib` reports active alarms and the history of alarms through the SNMP MIB tables. A new daemon called alarm management daemon, **AlarmMgmtD**, reports notifications defined in the alarm model table. The Alarm MIB is currently supported only on the MX Series.

To configure alarm management, include the **alarm-management** statement at the **[edit snmp]** hierarchy level.

[See [Interpreting the Enterprise-Specific Alarm MIB](#).]

- **SNMP MIB support for Ethernet OAM (MX Series)**—Starting in Junos OS Release 14.1, SNMP MIB support is enabled for Ethernet OAM on MX Series routers. See *Standard SNMP MIBs Supported by Junos OS* to view the standard MIBs (in IEEE 802.1ag, Connectivity Fault Management and IEEE 802.1ap, Management Information Base (MIB) definitions for VLAN Bridges) that are supported for Ethernet OAM.
- **Subscriber accounting MIB support (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.1, a new enterprise-specific Subscriber MIB, `jnxSubscriberAccountingTable`, has been added to the `jnxSubscriberGeneral` MIB to monitor subscriber sessions that are configured for RADIUS accounting. The `jnxSubscriberAccountingTable` MIB is a subset of the `jnxSubscriberTable` MIB.
- **SNMP support to monitor subscriber count per port (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 14.1, a new enterprise-specific Subscriber MIB, `jnxSubscriberPortCountTable`, has been added to the `jnxSubscriberGeneral` MIB to provide the number of active subscribers per port for tunneled and terminated subscribers.
- **Enhancement for viewing the details of user authentication (M Series, MX Series, and T Series)**—Starting with Junos OS Release 14.1, you can configure the following statements to view the attribute values of a logged in user:
 - **enhanced-accounting**—This configuration statement displays the details such as access privileges, access modes, and remote port of a user logged in through the RADIUS server or the TACAC+ server or local database. To enable this feature, use the **set system radius-options enhanced-accounting** command for the RADIUS server or the **set system tacplus-options enhanced-accounting** command for the TACAC+ server.
 - **enhanced-avs-max**—This configuration statement helps to limit the number of attribute values to be displayed when **enhanced-accounting** is enabled. To enable this feature, use the **set system accounting enhanced-avs-max** command.

Network Operations and Troubleshooting Automation

- **Upgrade to automation libraries (M Series, MX Series, and T Series)**—SLAX is an alternative syntax for XSLT that is tailored for readability and familiarity, following the style of C and Perl. SLAX was originally developed as part of Junos OS. It is used for on-box scripting to allow users to customize and enhance the CLI. The Junos OS automation infrastructure uses the `libslax` and `libxslt` open source libraries. Beginning

in Junos OS Release 14.1, these libraries have been upgraded to libxslt-1.1.28 and libslax.0.14.1.

- **Script dampening (M Series, MX Series, and T Series)**—Beginning in Junos OS Release 14.1, the impact of processor-intensive scripts on the performance of the Routing Engine can be minimized by configuring Junos OS to dampen or slow down the execution of any commit, op, or event script. To slow down script execution, include the **dampen** statement at the **[edit event-options event-script]**, **[edit system scripts commit]**, or **[edit system scripts op]** hierarchy level.

[See [Dampening Script Execution](#).]

Platform and Infrastructure

- **Virtual route reflector (VRR)**—Starting in Junos OS Release 14.1R3, you can implement route reflector capability using a general-purpose virtual machine on a 64-bit Intel-based blade server or appliance. Benefits of the VRR are:
 - Improved scalability (depending on the server core hardware use)
 - Scalability of the BGP network with lower cost using VRR at multiple locations in the network
 - Fast and more flexible deployment using Intel servers rather than router hardware
 - Space savings through elimination of router hardware
- **Virtual MX Series router (vMX)**—Starting in Junos OS Release 14.1R5, you can deploy vMX routers on x86 servers. vMX supports most of the features available on MX Series routers and allows you to leverage Junos OS to provide a quicker and more flexible deployment. Some vMX benefits include:
 - Optimizing carrier-grade routing for the x86 environment
 - Leveraging Junos OS for MX Series routers
 - Simplifying operations through consistency with MX Series routers
 - Introducing new services without reconfiguring the current infrastructure

Port Security

- **Storm control support (MX240, MX480, and MX960)**—Starting in Junos OS Release 14.1, support exists for storm control that enables the router to monitor traffic levels and to drop broadcast, multicast, and unknown unicast packets when a specified traffic level – called the storm control level – is exceeded, thereby preventing packets from proliferating and degrading the LAN.

You can modify the storm-control configuration by configuring a storm control profile at the **[edit forwarding-options]** hierarchy level, and then binding the storm control profile to a specific logical interface or to a group of logical interfaces. The group can include a range of interfaces or all interfaces on the switch.

[See [Layer 2 Device Security Feature Guide for MX Series Routers](#).]

- **Access port security (MX240, MX480, and MX960)**—Starting in Junos OS Release 14.1, Layer 2 software access port security is supported on the MX240, MX480, and MX960:
 - DAI—DAI protects switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.
 - DHCP option 82—You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the router against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation.
 - DHCP snooping—DHCP snooping filters and blocks ingress DHCP server messages on untrusted ports, and builds and maintains an IP address to MAC address binding database. Most port security features depend on DHCP snooping.
 - IP source guard—You can use the IP source guard access port security feature to mitigate the effects of source IP address spoofing and source MAC address spoofing.
 - Static IP—You can add static (fixed) IP addresses and bind them to fixed MAC addresses in the DHCP snooping database.
 - Trusted DHCP server interface—You can configure any interface on a switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

[See [Layer 2 Port Security Feature Guide for MX Series Routers](#).]

Routing Policy and Firewall Filters

- **Firewall filter match condition support for IPv6 extension headers (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, IPv6 firewall filters support extension header types as match conditions. This feature enables you to control the transmission of IPv6 packets based on the presence of specified extension header types in the packet. In the first fragment of a packet, the filter searches for a match in any of the extension header types. When a packet with a fragment header is found (a subsequent fragment), the filter only searches for a match of the next extension header type.

[See [Standard Firewall Filter Match Conditions for IPv6 Traffic](#).]

- **Firewall filter match condition support for additional ICMPv6 types (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, IPv6 firewall filters support several additional ICMPv6 match conditions. This feature enables you to specify match conditions for the following ICMP message types:
 - certificate-path-advertisement (149)
 - certificate-path-solicitation (148)
 - home-agent-address-discovery-reply (145)
 - home-agent-address-discovery-request (144)
 - inverse-neighbor-discovery-advertisement (142)
 - inverse-neighbor-discovery-solicitation (141)

- mobile-prefix-advertisement-reply (147)
- mobile-prefix-solicitation (146)
- private-experimentation-100 (100)
- private-experimentation-101 (101)
- private-experimentation-200 (200)
- private-experimentation-201 (201)

[See [Standard Firewall Filter Match Conditions for IPv6 Traffic](#).]

- **IPv6 support for next-hop groups (MX Series)**—Starting in Junos OS Release 14.1R2, this feature allows support of next-hop groups of type inet6 (IPv6). The following features are also supported:
 - Configuration of interfaces of inet6 (IPv6) type at the **[edit forwarding-options port-mirroring family inet6 output]** hierarchy level or subgroups at the **[edit forwarding-options port-mirroring family inet6 output next-hop-group]** hierarchy level.
 - Configuration of next-hop groups as filter action.
 - Configuration of next-hop groups as port-mirror destination when specified at the **[edit forwarding-options port-mirroring family inet6 output]** hierarchy level.
- **Support for consistent load balancing for ECMP groups (MX Series routers with MPCs)**—Effective in Junos OS Release 14.1R2, on MX Series 3D Universal Edge Routers with Modular Port Concentrators (MPCs) only, you can prevent the reordering of flows to active paths in an ECMP group when one or more paths fail. Only flows that are inactive are redirected. This feature applies only to Layer 3 adjacencies learned through external BGP connections. It overrides the default behavior of disrupting all existing, including active, TCP connections when an active path fails. Include the **consistent-hash** statement at the **[edit policy-options policy-statement *policy-statement-name* then load-balance]** hierarchy level. You must also configure a global per-packet load-balancing policy.

[See [Actions in Routing Policy Terms](#).]

Routing Protocols

- **Nonstop active routing for BGP multicast VPNs (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.1, this feature enables nonstop active routing for the BGP multicast VPNs (MVPNs). This feature synchronizes the MVPN routes, cmcast, provider-tunnel and forwarding information between the master and the backup Routing Engines.

[See [advertise-from-main-vpn-tables](#).]

- **Advertising multiple paths in BGP (MX Series and T Series)**—Starting in Junos OS Release 14.1, this feature allows up to 20 BGP add-paths to be advertised for a subset of prefixes that match the **add-path prefix-policy**.

To enable this feature for a prefix, the **add-path prefix-policy** term matching the prefix should have a new **then** action to set **add-path send-count<2...20>**. This new action is

a not applicable if the policy-statement containing it is used anywhere other than **add-path prefix-policy**.

[See [Actions in Routing Policy Terms](#), [path-count](#), and [prefix-policy](#).]

- **Egress protection for BGP labeled unicast (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.1, fast protection for egress nodes is available to services in which BGP labeled unicast interconnects IGP areas, levels, or autonomous systems (ASs). If a provider router detects that an egress router (AS or area border router) is down, it immediately forwards the traffic destined to that router to a protector router that forwards the traffic downstream to the destination.

[See [Egress Protection for BGP Labeled Unicast](#).]

- **Selecting backup LFA for IS-IS routing protocol (M Series, MX Series, and T Series)**—Starting with Junos OS Release 14.1, the default loop-free alternate (LFA) selection algorithm or criteria can be overridden with an LFA policy. These policies are configured for each destination (IPv4 and IPv6) and a primary next-hop interface. These backup policies enforce LFA selection based on admin-group, srlg, neighbor, neighbor-tag, bandwidth, protection-type, and metric attributes of the backup path. During backup shortest-path-first (SPF) computation, each attribute (both node and link) of the backup path, stored per backup-next hop, is accumulated by IGP. For the routes created internally by IGP, the attribute set of every backup path is evaluated against the policy configured per destination per prefix primary next-hop interface. The first or the best backup path is selected and installed as the backup next-hop in the routing table. To configure the backup selection policy, include the **backup-selection** configuration statement at the **[edit routing-options]** hierarchy level. The **show backup-selection** command displays the configured policies for a given interface and destination. The display can be filtered against a particular destination, prefix, interface, or logical systems.

[See [Example: Configuring Backup Selection Policy for IS-IS Protocol](#).]

Services Applications

- **Support for inline video monitoring (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1, video monitoring using media delivery indexing (MDI) criteria is supported. MDI information enables you to identify devices that are causing excessive jitter or packet loss for streaming video applications. To configure inline video monitoring criteria, include the **templates** and **interfaces** statements at the **[edit services video-monitoring]** hierarchy level.

Inline video monitoring is available for the following MPC interface cards:

- MPCE1
- MPCE2
- MPC-16XGE

[See [Inline Video Monitoring Feature Guide](#).]

- **Enhancements to IPsec packet fragmentation (MX Series routers with MS-MICs and MS-MPCs)**—Starting with Junos OS Release 14.1, in packets that are transmitted

through static and dynamic endpoint IPsec tunnels, you can enable the value set in the Don't Fragment (DF) bit of the packet entering the tunnel to be copied only to the outer header of the IPsec packet and to not cause any modification to the DF bit in the inner header of the IPsec packet. To copy the DF bit value to only the outer header and not modify the inner header, use the **copy-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level for static tunnels and at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level for dynamic endpoints. To configure the DF bit in only the outer header of the IPsec packet and to leave the inner header unmodified, include the **set-dont-fragment-bit** statement at the **[edit services ipsec-vpn rule *rule-name* term *term-name* then]** hierarchy level for static tunnels and at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level for dynamic endpoints.

[See [copy-dont-fragment-bit \(Services IPsec VPN\)](#) , [set-dont-fragment-bit \(Services IPsec VPN\)](#), [copy-dont-fragment-bit \(Services Set\)](#), and [set-dont-fragment-bit \(Services Set\)](#).]

- **Support for configuring template ID, observation domain ID, and source ID for Version 9 and IPFIX flow templates**—Starting with Junos OS Release 14.1, you can define the template ID for version 9 and IPFIX templates for inline flow monitoring. To specify the template ID for version 9 flows, include the **template-id *id*** statement at the **[edit services flow-monitoring version9 template *template-name*]** hierarchy level. To specify the template ID for version IPFIX flows, include the **template-id** statement at the **[edit services flow-monitoring version-ipfix template *template-name*]** hierarchy level. To specify the options template ID for version 9 flows, include the **options-template-id** statement at the **[edit services flow-monitoring version9 template *template-name*]** hierarchy level. To specify the options template ID for version IPFIX flows, include the **options-template-id** statement at the **[edit services flow-monitoring version-ipfix template *template-name*]** hierarchy level. The template ID and options template ID can be a value in the range of 1024 through 65535.

Until Junos OS Release 13.3, the observation domain ID is predefined and is set to a fixed value, which is derived from the combination of FPC slot, sampling protocol, Packet Forwarding Engine Instance, and LU Instance fields. This derivation creates a unique observation domain per LU per family. Starting with Junos OS Release 14.1, you can configure the observation domain ID, which causes the first 8 bits of the field to be configured. For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID.

[See [Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows](#) and [Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows](#).]

- **Increased number of IPsec tunnels (MX80, MX240, MX480, and MX960)**—Starting with Junos OS Release 14.1, you can configure a maximum of up to 8000 IPsec tunnels using 6000 service sets on a router. In such a scenario, you can employ up to 8000 logical interfaces in your environment and configure IPv4, IPv6, and dead peer detection (DPD) protocols. Until Junos OS Release 13.3, the maximum number of IPsec tunnels supported with 6000 service sets was 6000 tunnels.
- **IPsec invalid SPI notification (MX Series and T Series)**—Starting in Junos OS Release 14.1R3, you can enable automatic recovery when peers in a security association (SA) become unsynchronized. When peers become unsynchronized, this can cause the

transmission of packets with invalid security parameter index (SPI) values and the dropping of those packets by the receiving peer. You can enable automatic recovery by using the new **respond-bad-spi max-responses** configuration statement, which appears under the hierarchy level **[edit services ipsec-vpn ike policy]**. This statement results in a resynchronization of the SAs.

The *max-responses* value has a default of 5 and a range of 1 through 30.

- **Data plane inline support added for 6rd and 6to4 tunnels connecting IPv6 clients to IPv4 networks (MX Series routers with MPC line cards)**—Starting with Release 14.1R3, Junos OS supports inline 6rd and 6to4 on Modular Port Concentrator (MPC) line cards with Trio chipsets, saving customers the cost of using MS-DPCs for the required tunneling, encapsulation, and decapsulation processes. Anycast is supported for 6to4 (next-hop service interfaces only). Hairpinning is also supported for traffic between 6rd domains.

There are no CLI changes for 6rd and 6to4 configurations. To implement the inline functionality, configure service interfaces on the MPC card as inline services interfaces (**si-**) rather than as MultiServices (**ms-**) interfaces.

Two new operational commands have been added: **show services inline software statistics** and **clear services inline software statistics**.

- **Support for RPM probes with IPv6 sources and destinations (MX Series routers with MPCs)**—Starting with Junos OS Release 14.1R4, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address. To specify the destination IPv6 address used for the probes, include the **target (url ipv6-url | address ipv6-address)** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. You can also define the RPM client or the source that sends RPM probes to contain an IPv6 address. To specify the IPv6 protocol-related settings and the source IPv6 address of the client from which the RPM probes are sent, include the **inet6-options source-address ipv6-address** statement at the **[edit services rpm probe owner test test-name]** hierarchy level.
- **Support for PCP version 2 (MX Series)**—Starting in Junos OS Release 14.1R4, Junos OS supports Port Control Protocol (PCP) version 2, defined by IETF RFC 6887. PCP version 2 uses the client once for authentication. Junos OS is able to decode and process version 2 and version 1 messages. There are no CLI changes for PCP version 2 support.

Software Installation and Upgrade

- **Unified ISSU support for LFM (M Series and MX Series)**—Starting in Junos OS Release 14.1, the LFM protocol supports unified ISSU on M Series and MX Series with some restrictions. Connectivity failures that occur during the unified ISSU period are not detected until after unified ISSU has completed. If unified ISSU is initiated while discovery is in progress, the discovery completes only after unified ISSU has finished. LFM features that require Routing Engine involvement do not work during the unified ISSU period. Unified ISSU cannot run on the local and remote ends at the same time. The peer router must also be a Juniper Networks router running Junos OS that supports LFM ISSU for this feature to work on the local end.
- **Unified ISSU support (MX104)**—Starting with Junos OS Release 14.1, unified ISSU is supported on the MX104.

Unified ISSU is supported on the following MICs on MX104 routers:

- Gigabit Ethernet MIC with SFP (MIC-3D-20GE-SFP)
- Gigabit Ethernet MIC with SFP (E) (MIC-3D-20GE-SFP-E)
- Gigabit Ethernet MIC with SFP (EH) (MIC-3D-20GE-SFP-EH)
- 10-Gigabit Ethernet MICs with XFP (MIC-3D-2XGE-XFP)
- Tri-Rate Copper Ethernet MIC (MIC-3D-40GE-TX)

When unified ISSU is not supported on a MIC, at the beginning of the upgrade, Junos OS issues a warning that the MIC will be taken offline. After the MIC is taken offline and unified ISSU is complete, the MIC is brought back online.

Unified ISSU is not supported on the following MICs on MX104 routers:

- ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM)
- Channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE)
- Channelized E1/T1 Circuit Emulation MIC (H) (MIC-3D-16CHE1-T1-CE-H)
- Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE)
- Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (H) (MIC-4COC3-1COC12-CE-H)
- Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP (MIC-3D-4CHOC3-2CHOC12)
- Channelized SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP (MIC-3D-8CHOC3-4CHOC12)
- DS3/E3 MIC (MIC-3D-8DS3-E3)
- SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP (MIC-3D-4OC3OC12-1OC48)
- SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP (MIC-3D-8OC3OC12-4OC48)
- SONET/SDH OC192/STM64 MIC with XFP (MIC-3D-1OC192-XFP)

During unified ISSU, the protocols and applications that are not supported on MX104 routers are the same as those that are not supported on other MX Series routers undergoing unified ISSU.

[See [Unified ISSU System Requirements](#).]

- **Support for LACP with fast hellos during unified ISSU (MX Series)**—Starting in Junos OS Release 14.1, MX Series routers support LACP with fast hellos during unified ISSU. This support is disabled by default. To enable it you need to enter the new CLI knob **set protocols lacp fast-hello-issu** on both the DUT and peer routers before starting unified ISSU. The peer router must also be an MX Series router for this functionality to work.
- **Unified ISSU support on L2TP LNS (M Series, MX Series, and T Series)**—Junos OS Release 14.1 and later releases support unified ISSU on the L2TP network server (LNS). When an upgrade is initiated, the LNS completes any L2TP negotiations that are in progress but rejects any new negotiations until the upgrade has completed. No new tunnels or sessions are established during the upgrade.

[See [L2TP for Subscriber Access Overview](#).]

- **Unified ISSU support (TX Matrix Plus router with 3D SIBs)**—Starting in Junos OS Release 14.1, unified ISSU is supported on TX Matrix Plus routers with 3D SIBs. Unified ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.

[See [Unified ISSU System Requirements](#).]

- **Validate system software against running configuration on remote host**—Beginning with Junos OS Release 14.1R6, you can use the **on (host *host* <username *username*> | routing-engine *routing-engine*)** option with the **request system software validate *package-name*** command to verify candidate system software against the running configuration on the specified remote host or Routing Engine.
- **Validate system software add against running configuration on remote host or routing engine**—Beginning with Junos OS Release 14.1R6, you can use the **validate-on-host *hostname*** and **validate-on-routing-engine *routing-engine*** options with the **request system software add *package-name*** command to verify a candidate software bundle against the running configuration on the specified remote host or Routing Engine.

Spanning-Tree Protocols

- **Enhancements to STP logs (MX Series)**—Beginning with Release 14.1R1, Junos OS supports:
 - Logging of information in the internal ring buffer about events like Spanning Tree (such as STP, MSTP, RSTP, or VSTP) interface role or state change without having to configure STP traceoptions.
 - Capturing information as to what triggered the spanning-tree role or state change.

You can use the operational mode commands [show spanning-tree statistics message-queues](#), [show spanning-tree stp-buffer see-all](#), [show spanning-tree statistics bridge](#), and [show spanning-tree statistics interface](#) to get the information from ring-buffer,

bridge, and port statistics. `clear spanning-tree stp-buffer` clears the stp-buffer, and `clear spanning-tree statistics bridge` clears the statistics of the bridge.



NOTE: `show spanning-tree statistics interface` is not supported in Release 14.1R1 but is supported from Release 14.1R2.

Subscriber Management and Services



NOTE: Although present in the code, the subscriber management features are not supported in Junos OS Release 14.1R7. Documentation for subscriber management features is included in the Junos OS Release 14.1 documentation set.

-
- **RADIUS VSAs in output of test aaa command when authentication is unsuccessful (MX Series)**—Starting in Junos OS Releases 13.2R3 and 14.1R1, when you run the `test aaa` command, the command output includes all subscriber attributes when authentication is unsuccessful. In previous releases, the `test aaa` command returned a partial list of attributes when authentication was unsuccessful.

[See [Testing a Subscriber AAA Configuration](#).]

- **Using DHCP relay agent optional information to enhance security (MX Series)**—Starting in Junos OS Release 14.1, you can provide additional security by configuring DHCP relay agent to include optional information in client requests that the relay forwards to the DHCP server. The optional information helps minimize potential security shortcomings that might exist when a DHCP server on a central LAN allows connections from central access devices.

For DHCPv4, DHCP relay agent inserts Relay Agent Information Option (option 82) Agent Remote ID (suboption 2) into the relayed client requests. For DHCPv6, DHCPv6 relay agent inserts Relay Agent Remote-ID (option 37) into the relayed (RELAY-FORW) DHCPv6 messages.

[See [Using DHCP Relay Agent Option 82 Information](#) and [DHCPv6 Relay Agent Options](#).]

- **Support for Agent-Remote-Id when testing subscriber authentication (MX Series)**—Starting in Junos OS Release 14.1, you can use the `agent-remote-id ari` option with the `test aaa dhcp user` and `test aaa ppp user` commands to verify DHCP and PPP subscriber authentication in those networks that use the DSL Forum Agent-Remote-Id (VSA 26-2). If the ARI value that you specify includes special characters, such as a phone number that includes parentheses and a hyphen, you must enclose the value in quotation marks (""), as in the following example:

```
test aaa ppp user agent-remote-id "(202)555-1212"
```

[See [Testing a Subscriber AAA Configuration](#).]

- **RADIUS-based usage thresholds for subscriber services (MX Series)**—Starting in Junos OS Release 14.1, you can set usage thresholds for subscriber services that are dynamically activated or modified.

Subscriber management supports two types of usage thresholds—traffic volume and time. You use Juniper Networks VSAs to set the usage thresholds. The VSAs are transmitted in RADIUS Access-Accept messages for dynamically activated services, or in RADIUS-initiated CoA-Request messages for existing services. The traffic volume threshold sets the maximum amount of traffic that can use the service before the service is deactivated. The time threshold sets the maximum length of time that the service can be active.

[See [Usage Thresholds for Subscriber Services](#).]

- **Overriding short DHCP leases offered by third-party DHCP servers (MX Series)**—Starting in Junos OS Release 14.1, you can specify the minimum DHCP lease time allowed by the DHCP local server or DHCP relay agent. This feature enables you to avoid potential issues when a third party owns or manages the DHCP server or address-assignment pool that provides the client lease. In some cases, the third party might provide address leases that are unsuitable for the subscriber access environment. For example, extremely short lease times can create unnecessary traffic that results in reduced performance in the network.

In addition to specifying a minimum lease time, you can also specify the action the router takes when receiving a DHCP lease time that is less than the minimum acceptable value.

[See [DHCP Lease Time Violation](#).]

- **Support for L2TP AVPs that report access line information to the LNS (MX Series)**—Starting in Junos OS Release 14.1, you can configure the LAC to include L2TP AVPs in ICRQ messages to convey attributes such as line identification and traffic rates. The LAC receives the information from the DSLAM (ANCP access node) associated with the subscriber line; the values can be sourced from the ANCP agent or PPPoE intermediate agent tags carried in PADI and PADR discovery packets. You can also configure the LAC to send Connect-Speed-Update-Notification messages to the LNS to report updates to the subscriber connection speeds compared to the initial values conveyed by L2TP AVP 24 and AVP 38.

[See [Forwarding of Subscriber Access Line Information by the LAC](#) and [Configuring the LAC to Report Access Line Information to the LNS](#).]

- **Support for RADIUS accounting message retry and timeout (MX Series)**—Starting in Junos OS Release 14.1, you can include the new **accounting-retry** and **accounting-timeout** statements to specify retry and timeout values for RADIUS accounting messages separately from authentication messages. When you do so, the existing **retry** and **timeout** statements apply only to authentication messages; otherwise, they also apply to accounting messages.

Separate settings are useful for the following reasons:

- Authentication is time critical. Consequently, dropped packets need to be retransmitted quickly and short timeouts are desirable. Fewer retransmissions are sufficient because an unsuccessful subscriber is likely to attempt another login quickly.
- Accounting is less time critical, but it is important not to lose the accounting messages. Long timeouts and more retransmissions reduce packet loss.

[See [accounting-retry](#) and [accounting-timeout](#).]

- **Conserving IPv4 addresses for dual-stack PPP subscribers (MX Series routers with MPCs or MICs)**—Beginning in Junos OS Release 14.1, the IPv4 address saving feature for dual-stack PPP subscribers when they are not using the IPv4 service is expanded. During IPv4 address negotiation, if the broadband network gateway (BNG) receives an Access-Reject response from the RADIUS server that includes the Unisphere-Ipv4-release-control VSA and Reply Message attribute #18, the BNG sends an IPCP terminate request to the CPE. The CPE is then allowed to renegotiate IP NCP. However, if Unisphere-Ipv4-release-control VSA and Reply Message attribute #18 are not included in the Access-Reject response, the CPE must renegotiate the LCP link before being allowed to renegotiate IP NCP.
- **Dynamic Domain Name System (DNS) Resolver for IPv6 (MX Series)**—Beginning in Junos OS Release 14.1, in a network that uses Neighbor Discovery Router Advertisement (NDRA) to provide IPv6 addressing, the DNS server address can be provided in Router Advertisements sent to IPv6 hosts. The address is included in a field called Recursive DNS Server (RDNSS). This feature is useful in networks that are not running DHCPv6.

To configure (the default lifetime is 1800 seconds):

```
[edit dynamic-profiles profile-name protocols router-advertisement interface
 $junos-interface-name]
user@host# set dns-server-address $junos-ipv6-dns-server-address lifetime
#-of-seconds
```

[See [DNS Resolver for IPv6 DNS Overview](#).]

- **Subscriber interfaces over point-to-point MPLS pseudowires (MX Series routers with MPCs or MICs)**—Beginning in Junos OS Release 14.1, pseudowire subscriber interfaces support the following features:
 - Access Node Control Protocol (ANCP), which is used to monitor subscriber access lines and to report and modify subscriber traffic on the access lines between the subscribers and the access nodes.
 - Agent circuit identifier (ACI) interface sets, which are dynamic VLAN subscriber interfaces that are created based on ACI information and that originate at the same household or on the same access-loop port.
 - CoS shaping-rate and overhead-accounting attributes for dynamic ACI interface sets.
- **Minimum retransmission interval for L2TP control packets (MX Series)**—Starting in Junos OS Release 14.1, you can give a remote L2TP peer more or less time to respond to a control message sent by the local peer by including the **minimum-retransmission-interval** statement to configure the minimum interval that the local peer waits for a response. You can configure a minimum value of 1, 2, 4, 8, or 16 seconds; previously, the minimum interval was fixed at 1 second. The peer retransmits the message if a response is not received before the timeout expires, but waits for double the previous interval. The interval doubles with each retransmission until the maximum of 16 seconds is reached.

[See [Retransmission of L2TP Control Messages](#).]

- **Support for dynamic VLAN authentication based on subscriber packet type (MX Series)**—Starting in Junos OS Release 14.1, you can limit the packet types that trigger RADIUS authentication for dynamic, auto-sensed VLANs. In earlier releases, authentication is triggered by packet types configured with the **accept** statement in VLAN dynamic profiles.

Now you can specify that a subset of accepted packet types triggers authentication by including the **packet-types** statement at the **[edit interfaces *interface-name* auto-configure vlan-ranges authentication]** or **[edit interfaces *interface-name* auto-configure stacked-vlan-ranges authentication]** hierarchy level.

Because PPPoE subscribers are authenticated by PPP, you can conserve resources in a mixed PPPoE and IP environment by limiting VLAN authentication to the IP packets. You can also use this statement with the Client-Profile-Name VSA [26-174] to override a dynamic profile for certain subscriber types in a mixed access environment.

- **Clear DS-Lite mappings and flows (MX Series routers with MS-DPC interface cards)**—In Junos OS Release 14.1 and later releases, you can clear DS-Lite mapping statistics and flows for a specific subscriber, Basic Bridging Broadband Device (B4), or host behind a B4 using the following new operational commands.
 - **clear services nat mappings app**—Clear address-pooling paired mappings.
 - **clear services nat mappings eim**—Clear endpoint independent mappings.
 - **clear services nat mappings pcp**—Clear port control protocol (PCP) mappings.
 - **clear services nat mappings service-set**—Clear all NAT mappings for a service-set.
 - **clear services nat flows**—Clear all NAT flows. This command has the following scope options:
 - **b4address**—Clear all flows for a subscriber B4 address.
 - **service-set**—Clear all flows for a service set.
 - **subscriber**—The subscriber address.
- **Support for ATM virtual path shaping on ATM MICs with SFP (MX Series)**—Starting in Junos OS Release 14.1, class-of-service (CoS) hierarchical shaping for ATM virtual paths (VPs) is supported on MIC-3D-8OC3-2OC12-ATM.

The following configuration requirements apply to ATM VP shaping:

- All ATM interfaces that are members of an interface set must share the same virtual path identifier (VPI) and have a unique virtual circuit identifier (VCI).
- The ATM interface set can include only ATM interfaces. It cannot include Ethernet interfaces.
- The ATM interface set cannot include PPPoE over ATM interfaces, but it can include the underlying ATM interface over which PPPoE over ATM is carried.

To configure an ATM interface set and its members, use the **interface-set** stanza at the **[edit interfaces]** or **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level, specifying the ATM physical interface (**at-slot/mic/port**) and logical unit numbers.

After you configure the ATM interface set, you must create a CoS traffic control profile that includes the **peak-rate** (peak cell rate, or PCR), **sustained-rate** (sustained cell rate, or SCR), and **max-burst-size** (maximum burst size, or MBS) statements to shape the ATM cells transmitted on the ATM MIC. You then associate the traffic control profile to the ATM interface set.

- **Modifications to output fields of test aaa command (MX Series)**—Starting in Junos OS Release 14.1, the output of the **test aaa [dhcp | ppp] user** command is modified to improve readability. The modifications include the following:
 - The output now includes the corresponding tag for service-related attributes. For example, the following output includes the tag number (1) for the filter-service.
Service Name (1) - filter-service(100,200)
 - The output now includes the service activation type. For example:
Service Activation Type (1) - 1
 - The **junos-adf-rule-v4** output field is now titled **IPv4 ADF Rule**.
 - The **junos-adf-rule-v6** output field is now titled **IPv6 ADF Rule**.
- **DHCPv6 local server and relay agent username and option 37 (MX Series)**—Starting in Junos OS Releases 12.3R7, 13.2R4, 13.3R2, and 14.1R1, the MX Series router supports the generation of an ASCII version of the authentication username. When you configure a DHCPv6 local server or relay agent to concatenate the authentication username with the Agent Remote-ID option 37, the router uses only the remote-id portion of option 37 and ignores the enterprise number.

The router no longer supports the **enterprise-id** and **remote-id** options for the **relay-agent-remote-id** statement.

- **Realm name parsing (MX Series)**—Starting in Junos OS Release 14.1, the router supports realm name delimiters and parsing, when determining domain names that are used for the domain mapping feature. The realm name support is similar to the existing domain name support, and is used when subscriber usernames are presented in the realm name format (such as, **abc.com\marilyn**) rather than in the typical domain name format (such as, **joseph@abc.com**). You use the **parse-order** statement to specify the order in which the router searches for the domain name—you can specify that the router searches first for either the domain name or the realm name in the subscriber username. You can also specify the unique character that is the realm name delimiter, and the parsing direction the router uses to identify the resulting domain name that is used for domain mapping operations.
- **Specifying a domain map for usernames without a domain or realm name (MX Series)**—Starting in Junos OS Release 14.1, you can specify a domain map name of **none** for the **map domain-map-name** statement at the **[edit access domain]** hierarchy level. The router uses the domain map named **none** to perform domain map operations for subscriber usernames that do not include a domain or realm name.
- **MLPPP support for LNS and PPPoE subscribers (MX Series)**—Multilink PPP (MLPPP) support is provided for static and dynamic LNS (L2TP network server) and PPPoE (Point-to-Point Protocol over Ethernet) terminated and tunneled subscribers running on the MX Series with access-facing MPC2 slots. The following features are supported:

- Mixed mode for customers with both MLPPP and single link PPP subscribers
- Fragmentation-maps for both static and dynamic inline service si interfaces
- Co-existence support for member link IFL and the bundle IFL on different lookup engines
- Link fragmentation and interleaving (LFI) for a single-link bundle
- Minimization of fragment reordering
- **Subscriber management and services feature and scaling parity (MX104)**—Starting in Junos OS Release 14.1, the MX104 router supports all subscriber management and services features that are supported by the MX80 router. In addition, the scaling and performance values for the MX104 router match those of the MX80 router.
[See [Protocols and Applications Supported by MX5, MX10, MX40, and MX80 Routers](#).]
- **Subscriber management and services feature and scaling parity (MX2010 and MX2020)**—Starting in Junos OS Release 14.1, the MX2010 and the MX2020 routers support all subscriber management and services features that are supported by the MX240, MX480, and MX960 routers. In addition, the scaling and performance values for the MX2010 and the MX2020 match those of MX960 routers.
- **Support for up to 256 L2TP tunnel groups (MX Series)**—Starting in Junos OS Release 14.1R6, you can configure and commit up to 256 tunnel groups. In earlier releases, the CLI prevents you from committing the configuration when you create more than 32 groups.

User Interface and Configuration

- **New commit check for static label uniqueness**—Previously, applications, such as MPLS LSPs and Layer 2 circuits that use static labels, did not check to ensure that an incoming label name was not being used by another application. This caused the routing protocol process (RPD) to generate a core file. Starting in Junos OS Release 14.1, a commit check has been implemented to ensure the uniqueness of static labels across applications.

VLAN Infrastructure

- **VXLAN gateway support (MX80, MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 14.1R2, the MX80, MX240, MX480, MX960, MX2010, and MX2020 support Virtual Extensible Local Area Network (VXLAN) Gateways. Each VXLAN Gateway supports the following functionalities:
 - 32,000 VXLANs with one VXLAN per bridge domain
 - 8000 VXLAN Tunnel End Points (VTEPs)
 - 32,000 multicast groups
 - Switching functionality with traditional Layer 2 networks and VPLS networks
 - Inter VXLAN routing and VXLAN-only bridging domain with IRB
 - Virtual switches

- VXLAN with VRF functionality
- Configurable load balancing
- Statistics for remote VTEP
- **OVSDB support (MX80, MX240, MX480, MX960)**—Starting in Release 14.1, the Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which VMware NSX controllers and MX Series routers that support OVSDB can communicate. In an NSX multi-hypervisor environment, NSX controllers and MX Series routers can exchange control and statistical information through the OVSDB schema, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to entities in a physical network and the reverse.

You can set up a connection between the MX management interface (fxp0) and an NSX controller by using the **edit protocols ovssdb controller ip-address** statement, and OVSDB-managed Virtual Extensible LANs (VXLANs) by using the **edit bridge-domains bridge-domain-name vxlan ovssdb-managed** or **edit routing-instances routing-instance-name bridge-domain bridge-domain-name vxlan ovssdb-managed** statement.

VPNs

- **Control word for BGP VPLS (M320 and MX Series)**—For hash calculation, transit routers must determine the payload. While parsing an MPLS encapsulated packet for hashing, a transit router can incorrectly calculate an Ethernet payload as an IPv4 or IPv6 payload if the first nibble of the DA MAC is 0x4 or 0x6, respectively. This false positive can cause out-of-order packet delivery over a pseudowire. Starting in Junos OS Release 14.1, this issue can be avoided by configuring a BGP VPLS PE router to request that other BGP VPLS PE routers insert a control word between the label stack and the MPLS payload.

[See [Control Word for BGP VPLS Overview](#).]

- **Group VPN member support (MX240, MX480, and MX960)**—Starting with Junos OS Release 14.1, MX Series routers with MS-MPC-PIC and MS-MIC-16G line cards provide the group VPN member functionality support with one or more Cisco group controller or key servers (GC/KS). The group members can connect to a maximum of four Cisco GC/KSs with minimum interoperability with the cooperative servers.

This feature also provides system logging support for the group VPN functionality and routing instance support for both control and data traffic.

[See [Example: Configuring Group VPN on Routing Devices](#).]

- **IRB interface on EVPNs (MX Series routers with MPCs and MICs only)**—In an Ethernet VPN (EVPN) solution, multiple bridge domains can be defined within a particular EVPN instance, and one or more EVPN instances can be associated with a single Layer 3 VPN VRF. In general, each data center tenant is assigned a unique Layer 3 VPN VRF, although the tenant can consist of one or more EVPN instances or bridge domains per EVPN instance.

To support this flexibility and scalability factor, beginning with Junos OS Release 14.1, the EVPN solution provides support for the integrated routing and bridging (IRB)

interface on MX Series routers containing MPC interfaces to facilitate optimal Layer 2 and Layer 3 forwarding along with virtual machine mobility. The IRB interfaces are configured on each configured bridge domain including the default bridge domain for an EVPN instance.

[See [Example: Configuring EVPN with IRB Solution.](#)]

- **Virtual switch support for EVPNs (MX Series routers with MPCs and MICs only)**—Starting with Junos OS Release 14.1, the Ethernet VPN (EVPN) solution on MX Series routers with MPC interfaces is extended to provide virtual switch support that enables multiple tenants with independent VLAN and subnet space within an EVPN instance. Virtual switch provides the ability to extend Ethernet VLANs over a WAN using a single EVPN instance while maintaining data-plane separation between the various VLANs associated with that instance. A single EVPN instance can stretch up to 4094 bridge domains defined in a virtual switch to remote sites.

[See [Example: Configuring EVPN with Support for Virtual Switch.](#)]

- **Multihoming support for EVPNs (MX Series routers with MPCs and MICs only)**—Starting with Junos OS Release 14.1, the Ethernet VPN (EVPN) solution on MX Series routers with MPC interfaces is extended to provide multihoming functionality in the active-standby redundancy mode of operation.

To enable EVPN active-standby multihoming, include the **single-active** statement at the **[edit interfaces esi]** hierarchy level.

[See [Example: Configuring EVPN Multihoming.](#)]

- **VRF localization (MX Series routers with MPC line cards)**—Starting with Junos OS Release 14.1R3, VRF localization provides a mechanism for localizing routes of VRF to specific line cards to help maximize the number of routes that a router can handle. CE-facing interfaces localize all the routes of instance type VRF to specific line cards. If CE-facing interfaces are logical interfaces like AE or RLSQ or IRB, then the line card number has to be configured to localize routes. Core-facing line cards store all the VRF routes. These cards have to be configured as VPN core-facing only or VPN core-facing default. To configure VRF localization, configure the **localized-fib** configuration statement at the **[edit routing-instances instance-name routing-options]** hierarchy level and configure **vpn-localization** at the **[edit chassis fpc fpc-slot]** hierarchy level. The **show route vpn-localization** command displays the localization information of all the VRFs in the system.
- **Integrating EVPN with VXLAN for Layer 2 data center interconnect (MX Series routers with MPCs and MICs only)**—Virtual Extensible Local Area Network (VXLAN) is a technology that provides intra data center connectivity using a tunneling scheme to stretch Layer 2 connections over an intervening Layer 3 network.

The Ethernet VPN (EVPN) technology, on the other hand, provides a solution for multipoint Layer 2 VPN services with advanced multihoming capabilities, using BGP control plane over MPLS/IP network.

Although several solutions are available for data center connectivity, the integration of EVPN with VXLAN starting in Junos OS Release 14.1R4 provides an added advantage over the existing MPLS data center interconnect (DCI) technologies.

EVPN provides mechanisms for next generation DCI by adding extended control plane procedures to exchange Layer 2 MAC address and Layer 3 IP address information among the participating Data Center Border Routers (DCBRs). EVPN with its advanced features like active-active redundancy, aliasing, and mass MAC withdrawal helps in addressing the DCI challenges, such as seamless VM mobility and optimal IP routing, thus making it essential to provide VXLAN solutions over EVPN.

- **Leveraging DPCs for EVPN deployment (MX Series routers with DPCs)**—Starting with Junos OS Release 14.1R4, DPCs can be leveraged to provide support for Ethernet VPN (EVPN) functionality. Earlier, the EVPN functionality was provided by MX Series routers with MPC and MIC interfaces only.

The DPC support for EVPN is provided with the following considerations:

- DPCs provide support for EVPN in the active-standby mode of operation including support for the following:
 - EVPN instance (EVI)
 - Virtual switch (VS)
 - Integrated routing and bridging (IRB) interfaces
- DPCs intended for providing the EVPN active-standby support should be the customer edge (CE) device-facing line card. The provider edge (PE) device interfaces in the EVPN domain should use only MPC and MIC interfaces.
- **Active-active multihoming support for EVPNs (MX Series routers with MPCs and MICs only)**—Starting with Junos OS Release 14.1R4, the Ethernet VPN (EVPN) solution on MX Series routers with MPC and MIC interfaces is extended to provide multihoming functionality in the active-active redundancy mode of operation. This feature enables load balancing of Layer 2 unicast traffic across all the multihomed links on and toward a customer edge device.

The EVPN active-active multihoming feature provides link-level and node-level redundancy along with effective utilization of resources.

To enable EVPN active-active multihoming, include the **all-active** statement at the **[edit interfaces esi]** hierarchy level.

Related Documentation

- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 14.1R7 for the M Series, MX Series, and T Series.

- [Authentication, Authorization and Accounting \(AAA\) on page 57](#)
- [Application Layer Gateways \(ALGs\) on page 57](#)
- [Class of Service \(CoS\) on page 57](#)
- [High Availability \(HA\) and Resiliency on page 58](#)
- [Interfaces and Chassis on page 59](#)
- [IPv6 on page 61](#)
- [MPLS on page 61](#)
- [Network Management and Monitoring on page 61](#)
- [Routing Policy and Firewall Filters on page 62](#)
- [Routing Protocols on page 62](#)
- [Security on page 63](#)
- [Services Applications on page 63](#)
- [Subscriber Management and Services on page 65](#)
- [User Interface and Configuration on page 68](#)
- [VPNs on page 68](#)

Authentication, Authorization and Accounting (AAA)

- **Statement introduced to enforce strict authorization**—Starting in Junos OS Release 14.1R6, customers can use the **set system tacplus-options strict-authorization** statement to enforce strict authorization to the users. When a user is logging in, Junos OS issues two TACACS+ requests—first is the authentication request and then the authorization request. By default, when the authorization request is rejected by the TACACS+ server, Junos OS disregards this rejection and provides full access to the user. When the **set system tacplus-options strict-authorization** statement is set, Junos OS denies access to the user on failure of the authorization request.

Application Layer Gateways (ALGs)

- **Handling noncompliant IPv6 address in RTSP ALG (MX Series)**—Starting in Junos OS Release 14.1, Real-Time Streaming Protocol (RTSP) application-level gateway (ALG) cannot convert a noncompliant IPv6 address in its payload to an IPv4 address. The packet is not dropped, but it is forwarded to the receiving end of RTSP, which decides further processing of the packet.

Class of Service (CoS)

- **Change to TWAMP connection/session**—Beginning with Junos OS Release 14.1, a TWAMP connection/session comes up only if the session padding length is greater

than or equal to 27 bytes on the TWAMP Client. The valid range of padding length supported by the TWAMP Server is 27 bytes through 1400 bytes.

If IXIA is used as the TWAMP Client, packet length is supported from 41 bytes through 1024 bytes.

- **Change to interpolated WRED drop probability**—In Junos OS Releases 13.2R4, 13.3R2, and 14.1 and later, the interpolated fill level of 0 percent has a drop probability of 0 percent for weighted random early detection (WRED). In earlier Junos OS releases, interpolated WRED can have a nonzero drop probability for a fill level of 0 percent, which can cause packets to be dropped even when the queue is not congested or the port is not oversubscribed.

High Availability (HA) and Resiliency

- **Unified ISSU support for ATM MIC with SFP (MX Series)**—Starting in Junos OS Release 14.1, the ATM MIC with SFP (MIC-3D-8OC3-2OC12-ATM) supports unified ISSU with the following guidelines:
 - The PPP keepalive interval must be 10 seconds or greater. PPP requires three keepalives to fail before it brings down the session. Thirty seconds (10 seconds x 3) provides a safe margin to maintain PPP sessions across the unified ISSU in case of any traffic loss during the operation. Configure the interval with the **keepalives** statement at the **[edit interfaces at-interface-name]** or **[edit interfaces at-interface-name unit logical-unit-number]** hierarchy level.
 - The OAM F5 loopback cell period must be 20 seconds or greater to maintain ATM connectivity across the unified ISSU. Configure the interval with the **oam-period** statement at the **[edit interfaces at-interface-name unit logical-unit-number]** hierarchy level.
- **Enhanced show virtual-chassis heartbeat command (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1R3, a new state, **Detected**, has been added to the **show virtual-chassis heartbeat** command display output. When you configure a heartbeat connection in an MX Series Virtual Chassis, the **Detected** state indicates that the master Routing Engine in the specified member router has successfully exchanged a heartbeat connection message with the other member router when an adjacency disruption or split occurs in the Virtual Chassis. The **Detected** state persists until the heartbeat connection is reset, or until the Virtual Chassis forms again and a master router (protocol master) and backup router (protocol backup) are elected.

In previous releases, the **show virtual-chassis heartbeat** command displayed the **Alive** state for both split and merged Virtual Chassis conditions when a heartbeat message was successfully exchanged between the member routers. As a result, the only way to detect whether a heartbeat connection was in use during an adjacency split or disruption was to check for the **Heartbt** status in the **show virtual-chassis status** command. The new **Detected** state in the **show virtual-chassis heartbeat** command enables you to use a single command to determine whether or not the heartbeat message was successfully exchanged during an adjacency split.
- **Improved command output for determining GRES readiness in an MX Series Virtual Chassis (MX Series routers with MPCs)**—Starting in Junos OS Release 14.1R4, the

request virtual-chassis routing-engine master switch check command displays the following output when the member routers in a Virtual Chassis are ready to perform a graceful Routing Engine switchover (GRES):

```
{master:member0-re0}
```

```
user@host> request virtual-chassis routing-engine master switch check
Switchover Ready
```

In earlier releases, the **request virtual-chassis routing-engine master switch check** command displays no output to confirm that the member routers are ready for GRES.

The output of the **request virtual-chassis routing-engine master switch check** command has not changed when the member routers are not yet ready for GRES.

Interfaces and Chassis

- **Display revision number of Routing Engines (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 14.1, you can use the **show system commit revision** command to display the revision number of the Routing Engines in a dual Routing Engines-based router.

A commit error message is issued when overlapping subnets are configured within a logical interface.

- **Changes to DDoS protection policers for PIM and PIMv6 (MX Series routers with MPCs, T4000 with FPC5)**—Starting in Junos OS Release 14.1, the default values for bandwidth and burst limits have been reduced for PIM and PIMv6 aggregate policers to prevent starvation of OSPF and other protocols in the presence of high-rate PIM activity.

Policer Limit	New Value	Old Value
Bandwidth (pps)	8000	20,000
Burst (pps)	16,000	20,000

To see the default and modified values for DDoS protection packet-type policers, issue one of the following commands:

- **show ddos-protection protocols parameters brief**—Displays all packet-type policers.
- **show ddos-protection protocols *protocol-group* parameters brief**—Displays only packet-type policers with the specified protocol group.

An asterisk (*) indicates that a value has been modified from the default.

- **Changes to distributed denial of service statement and command syntax**—Starting in Junos OS Release 14.1, the protocol group and packet type syntax has changed for the **protocols** statement at the **[edit system ddos-protection]** hierarchy level and for the various **show ddos-protection protocols** commands.

The **filter-v4** and **filter-v6** packet types have been moved from the **unclassified** protocol group to the new **filter-action** protocol group.

The **resolve-v4** and **resolve-v6** packet types have been removed from the **unclassified** protocol group. They are replaced by the new **mcast-v4**, **mcast-v6**, **ucast-v4**, and **ucast-v6** packet types in the new **resolve** protocol group.

Both protocol groups also include an **aggregate** option for all unclassified packets in the group and an **other** option for unclassified packets that are not IPv4 or IPv6.

[See [protocols \(DDoS\)](#) and [show ddos-protection protocols](#).]

- **Deleting PTP clock client (MX104)**—Starting with Junos OS Release 14.1, on MX104 routers, when you toggle from a secure slave to an automatic slave or vice versa in the configuration of a Precision Timing Protocol (PTP) boundary clock, you must first delete the existing PTP clock client or slave clock settings and then *commit* the configuration. You can delete the existing PTP clock client or slave clock settings by using the **delete clock-client *ip-address* local-ip-address *local-ip-address*** statement at the **[edit protocols ptp master interface *interface-name* unicast-mode]** hierarchy level. You can then add a new clock client configuration by using the **set clock-client *ip-address* local-ip-address *local-ip-address*** statement at the **[edit protocols ptp master interface *interface-name* unicast-mode]** hierarchy level and committing the configuration. However, if you attempt to delete the existing PTP clock client and add the new clock client before committing the configuration, the PTP slave clock remains in the free-run state and does not operate in the auto-select state (to select the best clock source). This behavior is expected when PTP client or slave settings are modified.
- **Disabling distribution of connectivity fault management sessions on aggregated Ethernet interfaces (MX Series)**—Starting with Junos OS Release 14.1, connectivity fault management (CFM) sessions operate in distributed mode and are processed on the Flexible PIC Concentrator (FPC) on aggregated Ethernet interfaces by default. Starting with Junos OS Release 14.1, to disable the distribution of CFM sessions on aggregated Ethernet interfaces and to operate in centralized mode, include the **no-aggregate-delegate-processing** statement at the **[edit protocols oam ethernet connectivity-fault-management]** hierarchy level.

[See [IEEE 802.1ag OAM Connectivity Fault Management Overview](#).]

- **Preventing the filtering of packets by ARP policers (MX Series routers with MPCs)**—Beginning with Junos OS Release 14.1, you can configure the router to disable the processing of the specified ARP policers on the received ARP packets. Disabling ARP policers can cause denial-of-service (DoS) attacks on the system. Due to this possibility, we recommend that you exercise caution while disabling ARP policers. To prevent the processing of ARP policers on the arriving ARP packets, include the **disable-arp-policer** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet policer]** or the **[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet policer]** hierarchy level. You can configure this statement only for interfaces with inet address families and on MX Series routers with MPCs. When you disable ARP policers per interface, the packets are continued to be policed by the distributed DoS (DDoS) ARP policer. The maximum rate of is 10000 pps per FPC.

[See [Network Interfaces, Protocol Family and Interface Address Properties](#).]

- **Disabling the control word with active CFM sessions**—Starting in Junos OS Release 14.1, if you attempt to disable the control word by configuring the **no-control-word**

statement at the `[edit routing-instances routing-instance-name protocols l2vpn]` or `[edit protocols l2circuit neighbor neighbor-id interface interface-name]` hierarchy level for all Layer 2 VPNs and Layer 2 circuits over which you are running CFM MEPs, the existing CFM sessions are dropped. To prevent this problem, you must first deactivate the Layer 2 circuit, disable the control word, and reactivate the Layer 2 circuit on both the MEPs of a CFM session.

[See [Network Interfaces, Ethernet OAM.](#)]

IPv6

- **IPv6 support for SNMP traps (M Series, MX Series, and T Series)**—In Release 14.1R3 and later, Junos OS supports IPv6 source addresses of the SNMP traps.

MPLS

- **Enhanced support for GRE interfaces for GMPLS (MX Series)**—Starting in Junos OS Release 14.1, on GRE interfaces for Generalized MPLS control channels, you can enable the inner IP header's ToS bits to be copied to the outer IP packet header. Include the `copy-tos-to-outer-ip-header` statement at the `[edit interfaces gre unit logical-unit-number]` hierarchy level. Previously, the `copy-tos-to-outer-ip-header` statement was supported for GRE tunnel interfaces only.

[See [copy-tos-to-outer-ip-header.](#)]

- **Changes to MPLS protection options**—In Junos OS releases earlier than Release 14.1, you can configure both fast reroute and node and link protection on the same LSP. Starting in Junos OS Release 14.1, you can still configure both fast reroute and node and link protection on the same LSP; however, when you attempt to commit a configuration where both features are enabled, a syslog warning message is displayed that states: **The ability to configure both fast-reroute and link/node-link protection on the same LSP is deprecated and will be removed in a future release.**
- **Enhanced transit LSP statistics collection**—Starting in Junos OS Release 14.1R3, RSVP no longer periodically polls for transit LSP statistics. This change does not affect the `show mpls lsp statistics` command or automatic bandwidth operations for ingress LSPs. To enable the polling and display of transit LSP statistics, include the `transit-statistics-polling` statement at the `[edit protocols mpls statistics]` hierarchy level. You cannot enable transit LSP statistics collection if MPLS statistics collection is disabled with the `no-transit-statistics` statement at the `[edit protocols mpls statistics]` hierarchy level.

Network Management and Monitoring

- **New system log message indicating the difference in the Packet Forwarding Engine counter value (M Series, MX Series, and T Series)**—Effective in Junos OS Release 14.1R3, if the counter value of a Packet Forwarding Engine is reported less than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the `MIB2D_COUNTER_DECREASING` system log message for that specific counter.

[See [MIB2D_COUNTER_DECREASING.](#)]

- **SNMP proxy feature (M Series, MX Series, and T Series)**—Starting with Junos OS Release 14.1R6, you must configure the **interface <interface-name>** statement at the **[edit snmp]** hierarchy level for the proxy SNMP agent. In previous releases, configuring the interface for the proxy SNMP agent was not mandatory.
- **Enhancement for SONET interval counter (M Series, MX Series, and T Series)**—Starting with Junos OS Release 14.1R6, only the **Current Day Interval Total** output field in the **show interfaces interval** command for SONET interfaces are reset after 24 hours. In addition, the **Previous Day Interval Total** output field displays the last updated time in **hh:mm**.

[See [show interfaces interval](#).]

Routing Policy and Firewall Filters

- **New firewall filter match condition supported on MPCs**—Starting in Release 13.3R2, Junos OS supports the **gre-key** firewall filter match condition on MPC line cards on MX Series 3D Universal Edge Routers. To configure the **gre-key** firewall filter match condition, include the **gre-key** statement at the **[edit firewall family inet filter filter term term from]** hierarchy level.

Routing Protocols

- **Modification to the default BGP extended community value (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.1, the default BGP extended community value used for MVPN IPv4 VRF route import (RT-import) has been modified to the IANA-standardized value. Thus, the default behavior has changed such that the behavior of the **mvpn-iana-rt-import** statement has become the default. The **mvpn-iana-rt-import** statement is deprecated; we recommend that you remove it from configurations.
- **Removal of support for provider backbone bridging (MX Series)**—Starting with Junos OS Release 14.1, the provider backbone bridging (PBB) capability is disabled and not supported on MX Series routers. The **pbb-options** statement and its substatements at the **[edit routing-instances routing-instance-name]** hierarchy level, and the **pbb-service-options** statement and its substatements at the **[edit routing-instances routing-instance-name service-groups service-group-name]** hierarchy level are no longer available for configuring customer and provider routing instances for PBB.
- **BGP route advertisement**—Starting in Junos OS Release 14.1, if you include the **advertise-peer-as** statement in a BGP configuration, BGP advertises routes learned from one external BGP (EBGP) peer back to another EBGP peer in the same autonomous system (AS) but not back to the originating peer. In earlier Junos OS releases, if you include the **advertise-peer-as** statement in the configuration, BGP advertises routes learned from one EBGP peer back to another EBGP peer in the same AS and also to the originating peer.
- **BGP hides a route received with a label block size greater than 256 (M Series, MX Series, and T Series)**—Starting with Junos OS Release 14.1R5, when a BGP peer (running Junos OS) sends a route with a label block size greater than 256, the local speaker hides the route and does not re-advertise this route. The output of the **show route detail/extensive hidden/all** commands displays the hidden route and states the reason as **label block size exceeds max supported value**. In earlier Junos OS releases, when a

peer sent a route with a label block size greater than 256, the routing protocol process (rpd) terminated abnormally.

- **Configuring and establish targeting sessions with third-party controllers using LDP targeted neighbor (M Series and MX Series)**— Starting with Junos OS Release 14.1R5, you can configure LDP targeted neighbor to third-party controllers for applications such as route recorder that wants to learn label-FEC bindings of an LSR. LDP targeted neighbor helps to establish a targeted session with controllers for a variety of applications.

Security

- **Packet types added for DDoS protection L2TP policers (MX Series with MPCs, T4000 with FPC5)**—The following eight packet types have been added to the DDoS protection L2TP protocol group to provide flexibility in controlling L2TP packets:

cdn	scccn
hello	sccrq
iccn	stopccn
icrq	unclassified

Previously, no individual packet types were available for this protocol group and all L2TP packets were policed the same based on the aggregate policer value. The default values for the bandwidth and burst policers for all packet types is 20,000 pps. The default **recover-time** is 300 seconds for each of the L2TP packet types.

Services Applications

- **Restrictions for maximum blocksize for NAT port block allocation**—Beginning with Junos OS Release 14.1, the maximum blocksize for NAT port block allocation (PBA) is 32,000.
- **Support for display of NAT type for EIF flows (MX Series routers with MS-MICs and MS-MPCs)**—In the output of the **show services sessions extensive** command, the Translation Type field displays the value as NAPT-44 for Endpoint Independent Filtering (EIF) flows. Also, the label, EIF, is displayed beside the translation type parameter to enable easy identification of EIF flows.
- **Increased maximum number of logical interfaces for services (MX Series routers with MS-MPCs and MS-MICs)**—Until Junos OS Release 13.3, for every media logical interface on which services were configured (interface style services), a logical interface alias was internally created. This interface alias stores the topology chains for features that are performed on the logical interface after an input service was processed to avoid packet loops in the system. With interface aliases, the maximum number of logical interfaces supported with services was reduced to half the supported maximum number because each logical interface consumed two entries, namely, one for the interface itself and the other for the interface alias.

Starting with Junos OS Release 14.1R4, input interface aliases are not created for MS-MPCs and MS-MICs. As a result, the maximum number of logical interfaces that are supported with services PICs is equal to the maximum number supported on the system. After input service processing by MS-MPCs and MS-MICs, the services PIC sends the packet to the Packet Forwarding Engine on the multiservices (ms-) logical interface where the corresponding service is configured. Post-services are not supported on MS-MPCs and MS-MICs in Junos OS Release 13.2 and later.

- **Interoperation of ingress sampling and PIC-based flow monitoring (MX Series)**—If PIC-based flow monitoring is enabled on an ms- logical interface, a commit check error occurs when you attempt to configure ingress traffic sampling on that particular ms- logical interface. This error occurs because a combination of ingress sampling and PIC-based flow monitoring operations on an ms- logical interface causes undesired flow monitoring behavior and might result in repeated sampling of a single packet. You must not configure ingress traffic sampling on ms- logical interfaces on which PIC-based flow monitoring is enabled.
- **Changed range for maximum lifetime for PCP mapping**—Starting in Junos OS Release 14.1R5, the range for the maximum lifetime, in seconds, for PCP mapping that you can configure by using the **mapping-lifetime-max** statement at the **[edit services pcp]** hierarchy level is modified from 0 through 4294667, instead of the previous range that existed from 0 through 2147483647.
- **Change in support for service options configuration on service PICs at the MS and AMS interface levels (MX Series)**—Starting in Junos OS Release 14.1R4, when a multiservices PIC (ms- interface) is a member interface of an AMS bundle, you can configure the service options to be applied on the interface only at the ms- interface level or the AMS bundle level by including the **services-options** statement at the **[edit interfaces interface-name]** hierarchy level at a point in time. You cannot define service options for a service PIC at both the AMS bundle level and at the ms- interface level simultaneously. When you define the service options at the MS level or the AMS bundle level, the service options are applied to all the service sets on the ms- interface or AMS interface defined at **ms-fpc/pic/port.logical-unit** or **amsN**, respectively.
- **Generation of mspmand core file for flow control (MX Series with MS-MICs and MS-MPCs)**—Starting with Junos OS Release 14.1R5, instead of an eJunos kernel core file, the multiservices PIC management daemon core file is generated when a prolonged flow control occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the **dump-on-flow-control** option). The watchdog functionality continues to generate a kernel core file in such scenarios.
- **Optional inclusion of Flags in DTCP LIST Messages (MX Series)**—Starting in Junos OS Release 14.1R4, the Flags field is not a required parameter in the DTCP LIST message. The LIST request is not rejected if the LIST message does not contain the Flags field. If the DTCP LIST message contains the Flags field, the value of that field is processed. If the LIST message does not contain the Flags field, the CRITERIA field parameter is used for the Flags field.
- **Support for RPM probes for IPv4 and IPv6 sources and targets (TX Matrix Plus)**—Starting with Junos OS Release 14.1R5, you can configure the TXP-T1600, TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D router as the real-time performance monitoring (RPM) client router (the router or switch that originates the

RPM probes) which send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv4 or IPv6 address. RPM enables you to configure active probes to track and monitor traffic. The support for configuring RPM probes and RPM clients on TX Matrix Plus routers is in addition to the support for RPM that existed on M Series, MX Series, T1600, and T4000 routers in previous releases.

- **Changes in the format of session open and close system log messages (MX Series router with MS-MICs and MS-MPCs)**—Starting with Junos OS Release 14.1R5, with the Junos OS Extension-Provider packages installed and configured on the device for MS-MPCs and MS-MICs, the formats of the MSVCS_LOG_SESSION_OPEN and MSVCS_LOG_SESSION_CLOSE system log messages are modified to toggle the order of the destination IPv4 address and destination port address displayed in the log messages to be consistent and uniform with the formats of the session open and close logs of MS-DPCs.

The following shows the modified format of the MSVCS_LOG_SESSION_OPEN and MSVCS_LOG_SESSION_CLOSE system log messages:

```
month date hh:mm:ss syslog-server-ip-address yyyy-mm-dd hh:mm:ss
{NAT-type}<MSVCS_LOG_SESSION_CLOSE or MSVCS_LOG_SESSION_OPEN>:App:
application, source-interface-name fpc/pic/port\address in hexadecimal format
source-address:source-port source-nat-information ->
destination-address:destination-port destination-nat-information (protocol-name)
```

The following shows an example of the session closure message generated for MS-MPCs and MS-MICs:

```
Nov 26 13:00:07 10.137.159.1 2014-11-26 07:22:44:
{Dynamic-NAT-64-SS-NHS-1}MSVCS_LOG_SESSION_CLOSE: application:none, ae4.454
2402:8100:1:160:1:2:d384:463c:36822 [49.14.64.37:12261] -> [141.101.120.14]
64:ff9b::8d65:780e:80 (TCP)
```

- **Change in the test-interval range for RPM tests (MX Series)**—Starting in Junos OS Release 14.1R6, the minimum period for which the RPM client waits between two tests (configured by using the `test-interval interval` statement at the `[edit services rpm probe owner test test-name]` hierarchy level is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 0 seconds. A test interval of 0 seconds causes the RPM test to stop after one iteration.

Subscriber Management and Services



NOTE: Although present in the code, the subscriber management features are not supported in Junos OS Release 14.1R7. Documentation for subscriber management features is included in the Junos OS Release 14.1 documentation set.

- **CLI prompt to confirm clearing of all current PPPoE subscriber sessions (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.1, when you enter the **clear pppoe sessions** command and fail to include the name of an interface associated with the subscriber session that you want to gracefully terminate, the CLI prompts you to confirm that you want to clear all current PPPoE subscriber sessions. In earlier releases, the CLI does not prompt you and instead immediately terminates all the sessions.
- **Change to unicast reverse path forwarding (RPF) check and filter-based forwarding (FBF) compatibility (MX Series)**—Starting in Junos OS Release 14.1, the unicast RPF check is compatible with FBF actions. uRPF check is processed for source address checking before any FBF actions are enabled for static and dynamic interfaces. This applies to both IPv4 and IPv6 families.
- **Support for processing Cisco VSAs in RADIUS messages for service provisioning**—Starting with Junos OS Release 14.1, Cisco VSAs are supported for provisioning and management of services in RADIUS messages, in addition to the supported Juniper Networks VSAs for administration of subscriber sessions. In a deployment in which customer premises equipment (CPE) is connected over an access network to a broadband remote access gateway, the Steel-Belted Radius Carrier (SBRC) application might be used as the authentication and accounting server using RADIUS as the protocol, and the Cisco BroadHop application might be used as the Policy Control and Charging Rules Function (PCRF) server for provisioning services using RADIUS change of authorization (CoA) messages. Both the SBRC and the Cisco BroadHop servers are considered to be connected with the broadband gateway in such a topology.

By default, service accounting is disabled. If you configure service accounting using both RADIUS attributes and the CLI interface, the RADIUS setting takes precedence over the CLI setting. To enable service accounting using the CLI, include the **accounting** statement at the **[edit access profile *profile-name* service]** hierarchy level. To enable interim service accounting updates and configure the amount of time that the router waits before sending a new service accounting update, include the **update-interval *minutes*** statement at the **[edit access profile *profile-name* service accounting]** hierarchy level.

You can configure the router to collect time statistics, or both volume and time statistics, for the service accounting sessions being managed by AAA. To configure the collection of statistical details that are time-based only, include the **statistics time** statement at the **[edit access profile *profile-name* service accounting]** hierarchy level. To configure the collection of statistical details that are both volume-time-based only, include the **statistics volume-time** statement at the **[edit access profile *profile-name* service accounting]** hierarchy level.

- **Specifying the UDP port for RADIUS dynamic-request servers**—Beginning with Junos OS Release 14.1, you can define the UDP port number to configure the port on which the router that functions as the RADIUS dynamic-request server must receive requests from RADIUS servers. By default, the router listens on UDP port 3799 for dynamic requests from remote RADIUS servers. You can configure the UDP port number to be used for dynamic requests for a specific access profile or for all of the access profiles on the router. To define the UDP port number, include the **dynamic-request-port**

port-number statement at the **[edit access profile *profile-name* radius-server *server-address*]** or **[edit access radius-server *server-address*]** hierarchy level.

- **Support for applying access profiles to DHCP local server and DHCP relay agent**—Access profiles enable you to specify subscriber access authentication and accounting parameters. After access profiles are created, you can attach them at the **[edit system services dhcp-local-server]** hierarchy level on a DHCP local server for DHCP or DHCPv6 subscribers and at the **[edit forwarding-options dhcp-relay]** hierarchy level on a DHCP relay agent for DHCP or DHCPv6 subscribers, group of subscribers, or group of interfaces.

If you configured a global access profile at the **[edit access profile *profile-name*]** hierarchy level for all DHCP or DHCPv6 clients on a router that functions as a DHCP local server or a DHCP relay agent, the access profile configured at the **[edit system services dhcp-local-server]** or **[edit system services dhcpv-local-server dhcpv6]** hierarchy level on a DHCP local server for DHCP or DHCPv6 subscribers and at the **[edit forwarding-options dhcp-relay]** or **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level on a DHCP relay agent for DHCP or DHCPv6 subscribers take precedence over the global access profile.

Configuring an access profile for DHCP subscribers at the DHCP relay agent level or the DHCP local server level provides you with the flexibility and effectiveness of enabling DHCP authentication and accounting for specific subscribers instead of enabling them at a global level. If no access profile is configured at the DHCP relay agent level or the DHCP local server level, the global access profile becomes effective.

- **Support for specifying preauthentication port and password**—Starting in Junos OS Release 14.1, you can configure a router that operates as the RADIUS client to contact a RADIUS server for authentication and preauthentication requests on two different UDP ports and using different secret passwords. Similar to configuring the port numbers for authentication and accounting requests, you can define a unique port number that the router uses to contact the RADIUS server for logical line identification (LLID) preauthentication requests. You can also define a unique password for preauthentication requests. If you do not configure a separate UDP port or secret for preauthentication purposes, the same UDP port and secret that you configure for authentication messages is used.

To configure a unique UDP port number to be used to contact the RADIUS server for preauthentication requests, include the **preauthentication-port *port-number*** statement at the **[edit access radius-server *server-address*]** or **[edit access profile *profile-name* radius-server *server-address*]** hierarchy level.

To configure the password to be used to contact the RADIUS preauthentication server, include the **preauthentication-secret *password*** statement at the **[edit access radius-server *server-address*]** or **[edit access profile *profile-name* radius-server *server-address*]** hierarchy level.

The output of the **show network-access aaa radius-servers** command has been enhanced to display the preauthentication port number. The output of the **show network-access aaa radius-servers detail** command has been enhanced to display statistical information on the RADIUS messages exchanged during the preauthentication phase and the port number used for preauthentication.

- **On-demand IPv4 address re-allocation for dual-stack PPP subscribers**—Beginning in Junos OS Release 14.1R4, the behavior of the on-demand IPv4 address re-allocation process when there are no IPv4 addresses available is changed. During IPv4 address negotiation, if the RADIUS server sends an Access-Reject response to the broadband network gateway (BNG) that includes the Unisphere-ipv4-release-control VSA, the BNG sends an IPCP terminate request to the CPE. The CPE is then allowed to renegotiate NCP and request another IP address without the need to renegotiate the link.
- **LAC configuration no longer required for L2TP tunnel switching with RADIUS attributes (MX Series)**—Starting in Junos OS Release 14.1R5, when you use Juniper Networks VSA 26-91 to provide tunnel profile information for L2TP tunnel switching, you no longer have to configure a tunnel profile on the LAC. In earlier releases, tunnel switching failed when you did not also configure the LAC, even when the RADIUS attributes were present.

User Interface and Configuration

- **Configuring regular expressions (M Series, MX Series, and T Series)**—In all supported Junos OS releases, you can no longer configure regular expressions if they require more than 64 MB of memory or more than 256 recursions for parsing.

This change in the behavior of Junos OS is in line with the FreeBSD limit. The change was made in response to a known consumption vulnerability that allows an attacker to cause a denial-of-service (resource exhaustion) attack by using regular expressions containing adjacent repetition operators or adjacent bounded repetitions. Junos OS uses regular expressions in several places within the CLI. Exploitation of this vulnerability can cause the Routing Engine to crash, leading to a partial denial of service. Repeated exploitation can result in an extended partial outage of services provided by the routing protocol process (rpd).

- **Change in show route protocol evpn output**—In all supported Junos OS releases prior to Release 14.1, the output of the command **show route protocol evpn** does not provide any information for correlating the routes installed in the forwarding plane with routes exchanged in the signaling plane.

Starting with Junos OS Release 14.1, the command **show route protocol evpn** output provides additional correlation detail between forwarding plane and signaling plane routes.

[See [show route protocol](#).]

VPNs

- **Group VPN ike proposal commit check (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.1, the **proposals** option for the **policy** statement under the following hierarchies is mandatory and is checked on a commit:

```
[edit security group-vpn member ike policy policy-name]  
[edit security group-vpn server ike policy policy-name]  
[edit security ike policy policy-name]
```

Prior to Junos OS Release 14.1, the **proposals** option was not checked on a commit.

- **New output field added to the show route forwarding-table family vpls**

command—Starting in Junos OS Release 14.1, the **show route forwarding-table family vpls** command output contains an extra field to show “Enabled Protocols” for a routing table instance. The following sample output of the **show route forwarding-table family vpls** command shows the **Enabled Protocols** field when broadcast, unknown unicast, and multicast (BUM) hashing is enabled by configuring the **bum-hashing** statement at the **[edit routing-instances green protocols vpls]** hierarchy level:

```
user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0          4.4.3.2      dscd   519      1
lsi.1048832      intf  0          Push 262145    621    2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0          ucst   590      5 ge-2/3/9.0
0x30003/51       user  0          comp   627      2
ge-2/3/9.0       intf  0          ucst   590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0          ucst   619      4 ge-3/1/3.0
0x30002/51       user  0          comp   600      2
0x30001/51       user  0          comp   597      2
```

The following sample output of the **show route forwarding-table family vpls** command shows the **Enabled Protocols** field when broadcast, unknown unicast, and multicast (BUM) hashing is enabled by configuring the **bum-hashing** statement at the **[edit routing-instances green protocols vpls]** hierarchy level and MAC Statistics is enabled by configuring the **mac-statistics** statement at the **set routing-instances green protocols vpls** hierarchy level:

```
user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination      Type RtRef Next hop      Type Index  NhRef Netif
default          perm  0          4.4.3.2      dscd   519      1
lsi.1048834      intf  0          Push 262145    592    2

ge-3/0/0.0
00:19:e2:25:d0:01/48 user  0          ucst   590      5 ge-2/3/9.0
0x30003/51       user  0          comp   630      2
ge-2/3/9.0       intf  0          ucst   590      5 ge-2/3/9.0
ge-3/1/3.0       intf  0          ucst   591      4 ge-3/1/3.0
0x30002/51       user  0          comp   627      2
0x30001/51       user  0          comp   624      2
```

- **EVPN interface status commit check**—Starting in Junos OS Release 14.1, there is a commit check enforced for disabled interfaces in EVPN-type routing instances and for bridge domains that have EVPN configured.

Prior to Junos OS Release 14.1, there was a warning displayed when using the **show routing-instance** or **show routing-instance instance-name** configuration command at

the **[edit]** hierarchy level, which stated: **interface not defined**, but later commits did still succeed.

- **XML validate output for show evpn mac-table command (MX Series)**—Starting in Junos OS Release 14.1R7, the XML output has changed when entering the **display xml validate** option for the following **show** commands:
 - **show bridge mac-table**
 - **show evpn mac-table**
 - **show vpls mac-table**

In earlier releases, these show commands generated an XML invalid error in the output. Additionally, two new XML tags have been added to the XML output: **l2-mac-entry** and **l2-mac-entry-pvlan**.



NOTE: You must update any existing scripts that are capturing XML tags.

Related Documentation

- [New and Changed Features on page 21](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

Known Behavior

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 14.1R7 for the M Series, MX Series, and T Series.

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

- [Hardware on page 70](#)
- [High Availability \(HA\) and Resiliency on page 71](#)
- [MPLS on page 72](#)
- [Software Installation and Upgrade on page 72](#)
- [Subscriber Management and Services on page 72](#)

Hardware

- **Support for MIC-3D-8OC3-2OC12-ATM Revision 22 and later**—MIC-3D-8OC3-2OC12-ATM Revision 22 is supported only by the following Junos OS releases:

- Junos OS Release 12.3—12.3R9 and later
- Junos OS Release 13.3—13.3R6 and later
- Junos OS Release 14.1—14.1R4 and later
- Junos OS Release 14.2—14.2R3 and later
- Junos OS Release 15.1 and later

You must upgrade to a supported Junos OS release to use MIC-3D-8OC3-2OC12-ATM Revision 22 and later.

High Availability (HA) and Resiliency

- The MPC5E, MPC5EQ, and MP6E cards do not support unified ISSU on an MX Series Virtual Chassis.
- In an MX Series Virtual Chassis configuration, a unified in-service software upgrade (ISSU) from Junos OS Release 14.1 or 14.1R2 to Junos OS Release 14.2 fails with traffic loss. As a workaround, download the latest build of Junos OS Release 14.1R3, which contains a fix for this issue, and perform a unified ISSU to this build from Junos OS Release 14.1R1 or 14.1R2. You can then successfully perform a unified ISSU from the latest build of Junos OS Release 14.1R3 to Junos OS Release 14.2 in an MX Series Virtual Chassis.

MPLS

- **Removal of SRLG details from the SRLG table only on the next reoptimization of the LSP**—If an SRLG is associated with a link used by an ingress LSP in the router, then on deleting the SRLG configuration from that router, the SRLG gets removed from the SRLG table only on the next reoptimization of the LSP. Until then the output displays Unknown-XXX instead of the SRLG name and a non zero srlg-cost of that SRLG for the `run show mpls srlg` command.

Software Installation and Upgrade

- **Downgrading to Junos OS Release 12.3 when the configuration includes the `targeted-distribution` statement**—In Junos OS Release 12.3, the `targetted-distribution` statement at the `[edit interfaces demux0 unit logical-unit-number]` hierarchy level is misspelled. Starting in Junos OS Release 14.1R5, the spelling for this statement is corrected to `targeted-distribution`. If you use the misspelled `targetted-distribution` statement in Junos OS Release 14.1R5 or later, the CLI corrects the spelling to `targeted-distribution` in your configuration, so existing scripts still work. The correct spelling is not backward compatible; Junos OS Release 12.3 supports only the `targetted-distribution` spelling. If you downgrade from Release 14.1R5 or later to Release 12.3, all correctly spelled `targeted-distribution` statements are removed from the configuration, and configuration scripts with the correct spelling fail.

Subscriber Management and Services

- **Commands display different values for the LCP state of a tunneled subscriber on the LAC**—The `show ppp interface interface-name extensive` and `show interfaces pp0` commands display different values for the LCP state of a tunneled subscriber on the LAC. The `show ppp interface interface-name extensive` command displays STOPPED whereas the `show interfaces pp0` command displays OPENED (which reflects the LCP state before tunneling). As a workaround, use the `show ppp interface interface-name extensive` command to determine the correct LCP state for the subscriber.

Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R7 for the M Series, MX Series, and T Series.

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

- [General Routing on page 73](#)
- [Authentication and Access Control on page 80](#)
- [Class of Service \(CoS\) on page 80](#)
- [Forwarding and Sampling on page 80](#)
- [General Routing on page 81](#)
- [High Availability \(HA\) and Resiliency on page 81](#)
- [Infrastructure on page 81](#)
- [Interfaces and Chassis on page 82](#)
- [Layer 2 Features on page 84](#)
- [Layer 2 Ethernet Services on page 84](#)
- [MPLS on page 84](#)
- [Multicast on page 85](#)
- [Network Management and Monitoring on page 85](#)
- [Platform and Infrastructure on page 86](#)
- [Routing Protocols on page 89](#)
- [Services Applications on page 90](#)
- [Subscriber Access Management on page 91](#)
- [User Interface and Configuration on page 92](#)
- [VPNs on page 92](#)

General Routing

- During the Multicast Tunnel (mt) interfaces flapping, if some errors occur (in this case, there are some errors in tunnel being added in IFL tables), the tunnel data will be freed. But now, there is a bug, there is no checking for the IFL type, so some data are freed wrongly. So the kernel might crash. [PR823304](#)
- When both Routing Engines in a dual-Routing Engine system reboot too quickly with GRES enabled, 'ipsec-key-management' process would require a manual restart. [PR854794](#)
- After the local gateway address is removed from the Dynamic End Point (DEP) tunnel side, the IP Security (IPSec) security associations (SAs) still remain up and traffic continues to flow through the tunnel. [PR856068](#)
- In multicast scenario, when a interface which used for multicast Next-hopes (NHs) is deactivated/deleted, the operation might cause Layer 2 (L2) descriptor memory leaking, and then Flexible PIC Concentrator (FPC) might crash. [PR858643](#)
- 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 dump core due to memory corruption or memory double free. [PR870661](#)

- On M Series routers, packets are dropped upon setting AE LP discard-data knob, however there is no cli command to display the drop count. [PR876190](#)
- The ifl count is incorrect and will not be repaired until a pic restart. [PR882406](#)
- Under fabric blackholing condition, where fabric ASICs are not pulling traffic out from queuing ASIC and ISSU is performed, FPCs may crash while doing ASIC re-init during ISSU reboot because of residue traffic in queuing ASIC. If queuing ASIC does not provide a way to flush traffic, FPC may crash, as no traffic is primary requirement for queuing ASIC re-init during ISSU. [PR889475](#)
- Observed a traffic-drd daemon might hang once after logging into service PIC and restarting the net-monitor daemon. [PR889982](#)
- rpd crash at krt_iflchange() when VRF instance is deleted and added [PR911547](#)
- 6k service-sets are configured, when the service-set are activated and deactivated in very quick succession, the mspmand process might crash, the service running on the service PIC might be affected. [PR915784](#)
- PIC removal without PIC offline can cause FPC crash in case of PD-5-10XGE-SFPP PIC. [PR922655](#)
- In the large scaled dual stack PPPoE subscribers environment (in this case, 16k dual stack PPPoE subscribers), when IPv6 router-advertisement is configured for common edge IPv6 subscribers, if flapping dual stack PPPoE subscribers multiple times, in rare condition the rpd process might crash. The routing protocols are impacted and traffic disruption will be seen due to loss of routing information. [PR934081](#)
- There is 50Kpps drop in performance due to addition of new functionality over previous release. [PR935393](#)
- On an MX Virtual Chassis platform, when we restart one or both of the standby Routing Engines (REs), the log message "ksyncd_select_control_plane_proto: rhost_sysctlbyname_get: No such file or directory" might be observed as the ksyncd daemon attempts to select a communication protocol (UDP/TCP). After several tries, it will fall back to TCP and proceed as normal. [PR945925](#)
- Flow-control asserted whenever the "show services sessions" command is executed with large of sessions present in ms-mic resulting in traffic drop. [PR947674](#)
- By setting the tunnel MTU appropriately we can avoid getting into the situation where solution need to handle fragmented ESP packets. [PR948034](#)
- In MX Virtual Chassis (MX-VC) environment, the private local nexthops and routes pointing to private local nexthops are sent to PFE from master RE and not sent to slave RE, then an RE switchover happens. Now as the new master RE does not know about such nexthops and routes, they are not cleaned up. When a nexthop with same index is added on new master RE and sent to PFE, the PFE might crash due to a stale nexthop exist. [PR951420](#)
- These debug message which are fixed in the releases mentioned below. [PR961558](#)
- In packet-triggered subscribers and policy control (PTSP) scenario, reboot the MS-DPC may result in new subscribers not coming up on MX Series Routers. [PR967070](#)

- BFD session flap is expected in scaled environment by restarting "chassisd" or by restarting "FPC". [PR969023](#)
- Trigger for seeing fabric errors: ----- Have traffic between 2 FPCs. Offline the destination FPC Result of the trigger: ----- Fabric errors like "Fo: Request timeout errors" will be seen on the source FPC. As the source fpc has outstanding requests to the offlined destination FPC, the requests timeout and hence these errors are seen. These messages are expected. [PR971956](#)
- When configured inline MLPPP, packets are indiscriminately dropped on the member links when the traffic is oversubscribed. [PR982175](#)
- On MX series routers with MS-MPC, when MS-MPC is booting a PIC may fail to boot successfully the first time. This results in increased boot time for this particular PIC. [PR986166](#)
- When a MAC moves from one VTEP to another VTEP, it is not learnt behind the new VTEP until the old VTEP ages out this MAC. This will cause traffic for this MAC to black hole till it ages out on the old VTEP. [PR988270](#)
- IPsec endpoint fails to decrypt packets on some of the tunnels with NAT between IPsec endpoints. [PR989054](#)
- Issue the CLI command "restart packet-triggered-subscribers" might cause sessions out-of-sync between the MX SAE and SRC (external policy manager), which results in new subscribers unable to be created. [PR990788](#)
- An inconsistency between JUNIPER-VPN-MIB and MPLS-L3VPN-STD-MIB with the number of interfaces for an routing-instance has been identified. For example with the following configuration: user@router-re0> show configuration routing-instances ri1 instance-type vrf; interface ge-2/0/8.10; interface lo0.10; route-distinguisher 65000:1; vrf-target target:65000:1; vrf-table-label; According to the MPLS-L3VPN-STD-MIB there are two interfaces in this routing-instance: MPLS-L3VPN-STD-MIB :: mplsL3VpnVrfAssociatedInterfaces: OID: 1.3.6.1.2.1.10.166.11.1.2.2.1.8 Description: Total number of interfaces connected to this VRF (independent of ifOperStatus type). {master} user@router-re0> show snmp mib walk 1.3.6.1.2.1.10.166.11.1.2.2.1.8 mplsL3VpnVrfAssociatedInterfaces.3.114.105.49 = 2 However according to JUNIPER-VPN-MIB there are three interfaces in this VRF: JUNIPER-VPN-MIB :: jnxVpnIfStatus OID: 1.3.6.1.4.1.2636.3.26.1.3.1.10 Description: Status of a monitored VPN interface. user@router-re0> show snmp mib walk 1.3.6.1.4.1.2636.3.26.1.3.1.10 jnxVpnIfStatus.2.3.114.105.49.733 = 5 jnxVpnIfStatus.2.3.114.105.49.754 = 5 jnxVpnIfStatus.2.3.114.105.49.774 = 5 The interfaces in the example are: {master} user@router-re0> show snmp mib walk 1.3.6.1.2.1.2.2.1.2 ifDescr.733 = ge-2/0/8.10 ifDescr.754 = lo0.10 ifDescr.774 = lsi.0 The fix for this issue adjusts this by removing the dynamic interface (in this case, lsi.0) from the interface list of JUNIPER-VPN-MIB. [PR1011763](#)
- On MX series router, when a instance type is changed from VPLS to EVPN, and in the same commit an interface is added to the EVPN instance, the newly added EVPN interface might not be able to come up. [PR1016797](#)

- When recovering from a split master Virtual Chassis (VC), the line-cards in the new VC-Bm chassis may contain provisioning data out of synchronization versus the master RE. [PR1036795](#)
- On Trio-based platform, when the feature flow-control is disabled (enabled by default) by using CLI command "no-flow-control" knob (for example, under "gigether-options" hierarchy), after bringing up or rebooting the MPC, due to the fact that status of the hardware may not be updated correctly, the flow control on that MAC may remain enabled. [PR1045052](#)
- If the flow routes (flow route is an aggregation of match conditions for IP packets) are active in the kernel, the rpd process might crash after executing command "show route table <X>.inetflow.0 extensive". [PR1047271](#)
- The following message is generated every 5 second in MX104 on 14.2R1~R3 and 15.1R1. xxx chassisd[1362]: Cannot read hw.chassis.startup_time: No such file or directory .. [PR1049015](#)
- In Layer2 port mirroring scenario with maximum-packet-length configured, packets are not getting mirrored to any of next-hop-subgroup interfaces when a new interface is added to the next-hop-subgroup. As a workaround, we can remove the configuration of "maximum-packet-length". [PR1052559](#)
- When the JAM package is to be un-installed, the NG-MPC cards must be offlined before doing so. When there is large number of RLSQ interfaces configured, it takes few minutes for all the interfaces to be deleted from the system. Attempting to un-install the JAM package before all the interfaces are deleted, can result in crashing of the router's RE. To avoid such a situation sufficient time gap must be given between offlining the NG-MPC and un-installing the JAM package so that all the interfaces are deleted from the system. [PR1052571](#)
- In IPsec environment, trying to send traffic over site-to-site IP Security (IPsec) VPN, the mspmand process, which manages the Multiservice PIC, might crash due to invalid memory access on the Multiservice PICs when initiating IPsec VPN session. [PR1057214](#)
- On MX platform with "subscriber-management" enabled, after clients are established/bound, in a corner case, the Trio-based line cards might crash after the bbe-smgd process restart. [PR1057483](#)
- LACP is not supported with virtIO. [PR1059231](#)
- With inline L2TP IP reassembly feature configured, the MX Series routers with MPCs/MICs might crash due to a memory allocation issue. [PR1061929](#)
- rpd core may happen during install nexthop to kernel with EEXIST error. This issue might be encountered on new master RE after performing multiple RE switchover. [PR1062405](#)
- When using the 'mpls-ipv4-template' sampling template for non-IP traffic encapsulated in MPLS, log messages such as this one can be seen frequently (depending upon the rate of traffic - it could in the range of few messages to 2000-3000 messages per minute) Feb 18 09:28:47 Router-re0 : %DAEMON-3: (FPC Slot 2, PIC Slot 0) ms20 mspmand[171]: jflow_process_session_close: Could not get session extension: 0x939d53448 sc_pid: 5 Depending upon the frequency of the messages per sec, eventd

PR1065788

- traffic may get dropped on the device. [PR1081579](#)

- If a router has Service PIC equipped but without any Service PIC specific configurations, the CPU usage on this PIC/FPC might be high. Have some configurations under below knob could prevent from this issue: [system processes process-monitor traceoptions] OR [chassis fpc <fpc-slot> pic <pic-slot> adaptive-services service-package extension-provider] OR [services] [PR1081736](#)
- The Trio-based line card might reboot immediately after restarting l2tp daemon at L2TP network server (LNS) during login/logout of scaled (e.g. 10k) L2TP clients. [PR1082321](#)
- If an RX OTN signal is received with non-zero values in the OPU4 PSI[2:81] overhead bytes, software may report the following alarm: "ODU-MSIM MOD-1-2 MAJOR SA 2015-04-16 ODU PM MSIM Multiplex Structure Identifier Mismatch" This is a false alarm, is not service-affecting and can be safely ignored. The datapath is not affected by MSI values when OPU4 payload type is 100GBaseR (PT=0x7). [PR1083141](#)
- Sparks-IPSec Performance drop while testing with 1 Next-hop style service-set and 10 sessions for PPS 600000, Framesize 1518 [PR1084376](#)
- Invalid Ethernet Synchronization (ESMC) frames may be transmitted by MX router when activating LAG and tag-protocol-id under interfaces. [PR1084606](#)
- TCP messages do not have their MSS adjusted by the Multiservices MIC and MPC if they do not belong to an established session. [PR1084653](#)
- Memory leak likely given VC-Bm Routing Engine memory utilization sustained at 97% usage with high active and inactive memory utilization compared with VC-Mm Routing Engine. [PR1088132](#)
- When BGP multipath is enabled in a Virtual Routing and Forwarding (VRF), if "auto-export" and "rib-group" are configured to leak BGP routes from this VRF table to another, for example, the default routing table, then traffic coming from the default routing instance might not be properly load balanced due to the multipath-route leaked into the default routing table is not the active route. This is a random issue. As a workaround, only use "auto-export" to exchange the routes among the routing tables. [PR1099496](#)
- Some of the new revisions (for example, REV 30, REV 31) of the MICs cannot come up with NG-MPC2 or NG-MPC3 line card. We can check the MIC version by CLI command "show chassis hardware detail | no-more". [PR1100073](#)
- On MX platforms, in subscriber management environment, when carrying scaling subscribers, as the PFE process (pfed) memory usage will grow along with the number of subscribers, the pfed memory usage limit may get reached (that is, 512M) because of the subscriber scale and number of service attached to the subscribers (for example, when carrying more than 140k single stack PPPoE subscribers per chassis, 4 services per subscriber), in this situation, the pfed crash may occur due to memory exhaustion. [PR1102522](#)
- cpcdd core observed in scaled scenario [PR1103675](#)
- On MX platform, the output of CLI command "show system subscriber-management route" may be shown as empty. [PR1104808](#)

- Dynamic vlan ifl is not removed with 'remove when-no-subscriber' configuration [PR1106776](#)
- On MX platform with "subscriber-management" enabled, while high scaled subscribers (for example, 126K dual-stack DHCP v4/v6 subscribers over VLAN demux) login/logout at high rate, Trio-based line cards which hold subscribers might crash after the bbe-smgd process restart. [PR1109280](#)
- Resolved problem with Syslog messages generated like "krt_decode_resolve for 239.255.255.250, 101.11.67.33: no logical interface for index 1073741825" when Multicast packets are received on Subscriber interfaces. [PR1110967](#)
- CLI core dump is due to repeated mismatched XML open/close directives in the "show pppoe lockout" output. This issue is most likely to occur when there is a ratio of 8 PPPoE clients in lockout per VLAN. [PR1112326](#)
- When using the SIP-ALG with MS-MIC and MS-MPC cards in the 13.3R8 and earlier builds the MSPMAND process can core. [PR1120100](#)
- On MX240/MX480/MX960/MX2010/MX2020 products with MPC2E-3D-NG/MPC2E-3D-NG-Q/MPC3E-3D-NG/MPC3E-3D-NG-Q with MIC-3D-4XGE-XFP, IFD flap detection is much slower. It might lead to high FRR time, some traffic might be lost. [PR1122589](#)
- In multihoming EVPN scenario and the customer facing interface is an AE interface, after moving an interface from the EVPN instance into a VPLS instance, traffic loss might be seen on CE facing FPC. [PR1126155](#)
- An incorrect destination MAC address is applied to the packet when a DHCPv6 Offer/Advertise packet is sent back to the subscriber from a non-default routing instance across a pseudowire. [PR1127364](#)
- On MS-MIC, TCP session Up/Down causes JSERVICES_NAT_* and JSERVICES_SESSION_* messages though severity level "none" is configured for services [PR1137596](#)

Authentication and Access Control

- A reflected cross site scripting (XSS) vulnerability in SRX Web Authentication (webauth) may allow the stealing of sensitive information or session credentials from firewall users. This issue affects the device only when Web Authentication is used for firewall user authentication. Refer to JSA10640 for more information. [PR907664](#)

Class of Service (CoS)

- In subscriber management environment with ANCP COS adjustment, after enable and disable SRL service, ANCP Bandwidth correction might not get properly adjusted, the COS "overhead adjustment mode" and "byte adjust" fail to restore. [PR1114519](#)

Forwarding and Sampling

- MX480 pfed: downward spike recived from pfe for ibytes_reply:1692427
ibytes_record:1798030 for ifl:1600 These downward spike messages are informational and do not indicate that an error exists in the router. The message originates from the router's PFED (packet forwarding engine daemon) and indicates that information is being downloaded from a particular interface. It is a transient condition and gets cleared. [PR833091](#)
- Due to incorrect handling of an invalid IPC (inter-process communication) message PFED might get stuck in a loop trying to process the invalid message over and over again. This will cause PFED to consume a high percentage of CPU. Also it can cause a memory leak inside PFED. [PR893428](#)
- This defect is seen only when an existing child link from an AE is moved to a newly created AE, simultaneously from both-ends. The new AE is listed as child link in the existing AE in 'show interface ae<>.0 extensive' CLI. [PR965872](#)
- This issue affects a system with two routing engines (REs) with "graceful-switchover" configured. When performs upgrade to Junos OS Release 13.3 from previous releases, without deactivating "graceful-switchover", master and backup routing engines are likely to become unresponsive due to running out of memory. The Routing Engines need a power reset to restore service. [PR1033926](#)
- When VRRP is configured on Trio based MX interfaces, Static mac entries are installed on PFE in the MAC-DB as part of the mac-filter installations. RPF mib-walk triggering a walk over the MAC MIB entry(Walk over the static mac entries with no OIDs) causes the error message. During the walk, it is expected that no entries are read from static mac-db entries, however the EODB is not set to indicate MAC-DB walk has ended. This error log does not have any functional impact on the RPF mib-walk. mib2d[xxx]: MIB2D_RTSLIB_READ_FAILURE: check_rtsock_rc: failed in reading mac_db: 0 (Invalid argument) mib2d[xxx]: SNMP_GET_ERROR1: macStatsEntry getNext failed for interface: index1 ge-*/ */* (Invalid argument) [PR1042610](#)
- On Trio-based platform, when the Layer 3 packets destine to an Integrated Routing and Bridging (IRB) interface and then hit the underlying Layer 2 logical interfaces (IFLs), due to the egress feature list of the Layer 2 IFLs may get skipped, the features under the family bridge (for example, the firewall filter) on the Layer 2 interfaces may not be executed. [PR1073365](#)

General Routing

- scale-subscriber "License Used" filed shows wrong value after GRES [PR980399](#)
- Activation of ADFv4 might be incomplete (that is, only first rule will send to dfwd) in case ERX-Ipv6-Primary-Dns = ERX-Ipv6-Secondary-Dns in the Access-Accept message from the RADIUS server. [PR1041764](#)
- In subscriber management environment, the authentication service process (authd) may crash if the Framed-Route attribute is received from a RADIUS server on MX series routers. When the issue occurs, the authd may keep crashing and no subscribers could login. [PR1057003](#)
- In subscriber Management environment, subscriber login/logout causes multiple daemons to interact each via Inter-Process Communication (IPC). A 32 bit counter "Current time" is used internally for this type of communication to collect stats like elapsed time etc. The error log message, like "invalid current (31902492), baseline (4287968879), time (0)", might be observed when "Current time" is overflowing and holding an invalid time. [PR1079507](#)

High Availability (HA) and Resiliency

- During a router hardware upgrade procedure, in a dual REs system, the newly installed Routing Engine (RE) may overwrite the other RE configuration with the factory default configuration. As a result, both REs may boot-up in "Amnesiac" mode. This situation can occur under following conditions: - RE0 has default factory configuration and, - RE1 has "commit synchronize" enabled - Both RE0 and RE1 boot-up simultaneously, or - RE0 is UP and running and RE1 is restarted. [PR909692](#)

Infrastructure

- When a neighbor frequently fails to answer ipv6 neighbor solicitation messages or does so with delay, there is a potential race condition where the neighbor entry can be marked reachable in the cache but have no MAC address. As a consequence, transit traffic going to this neighbor will be sent with a random destination MAC and be flooded by any switch in the middle to all devices on the LAN for a duration equal to the reachability timeout (30 seconds, add or subtract some jitter value). [PR608439](#)
- Executing the 'show multicast route' CLI command with a crafted argument may cause the CLI process to crash, creating a cli.core.0.gz. The impact of the crash is limited to the CLI process executing the command, so malicious exposure is minimal. [PR939262](#)

Interfaces and Chassis

- With Junos OS Release 11.4 or later and Enhanced SCB installed on a mix of MX Series routers and DPC cards, REG_ERR messages might be reported under certain traffic flow conditions from MX Series routers to the DPC card. On the receiving DPC card fabric cell received out of order will be re-ordered and merged to build the packet. If this out-of-order delivery is too high a reorder event will be triggered and all cells belonging to the packets are dropped. The frequency is low rate. The following syslog entry will be reported Sep 29 20:43:10 node fpc8 ICHIP(3)_REG_ERR:first cell drops in ichip fi rord : 4122 Sep 29 20:43:10 node fpc8 ICHIP(3)_REG_ERR:Non first cell drops in ichip fi rord: 7910 [PR821742](#)
- When "show chassis alarms" command is run, "FPC X Dest Err" alarm is observed, though no Dest Errors are reported in commands "show chassis sibs" OR "show chassis fabric fpcs". Commonly observed on routers where Fabric Blackholing feature (enabled by default) has identified traffic black holing on one or more FPCs due to error(s) observed on single FPC. And feature has attempted to recover the situation by rebooting affected Single FPC. To confirm, run "show chassis fabric reachability detail" command and verify output to match: Fabric reachability action: Fabric reachability action : FPC action Acting on : Single FPC error [PR890598](#)
- PPP interface MTU changes after making any configuration changes to the system [PR897940](#)
- On MX Series router, the physical or logical interfaces (ifd/ift) might be created and marked UP before a resetting FPCs' fabric planes are brought up and ready to forward traffic, as a result, traffic might be black-holed during the time window. This window of traffic black-hole is particular long if the chassis is heavily populated with line-cards, for example, the router has large scale of configuration (routes or subscribers), and coupled with a lot of FPC reset, such as upon a node power up/reset. [PR918324](#)
- Non-Existent leg in AE bundle prevents DHCP subscribers from coming up. [PR918745](#)
- In the large scaled VPLS environment (in this case, 2k+ VPLS sessions), in rare condition, when larger scale route updates during In-Service Software Upgrade (ISSU), one of the FPCs in the router may stuck in Ready state. In normal upgrades, it will not be observed. [PR986264](#)
- On dual RE platforms, when adding the logical interfaces (IFLs) and committing, due to the device control process (dcd) on backup RE may fail to process the configuration and keep it in the memory, in some cases (not happening all the time), it might be observed that the memory of the dcd keeps increasing on backup RE. [PR1014098](#)
- Power-off by pushing Offline button on master RE causes lots of packets lost although GRES/NSR is configured. FPC gets rebooted after RE switchover which also causes traffic loss. [PR1034164](#)
- Service name table PADI delay timer accuracy is off by 2 more seconds than configured delay value. [PR1052632](#)
- In subscriber management environment with scaled PPPoE sessions (in this case, 4K PPPoE sessions), the IPCP/IPv6CP negotiation might be out of sequence due to a timing issue. For both families, a delayed configuration-ack is responded to the

configuration-request. This might lead to multiple IPCP/IPv6CP messages to negotiate a single session. [PR1072062](#)

- After RE switchover with NSR enabled in 14.1, VRRP sessions may flap once. [PR1072086](#)
- MAX-ACCESS value has been changed in jnx-otn.mib for the following oids:
jnxOtnIntervalOdu15minIntervalNumber jnxOtnIntervalOtu15minIntervalNumber
jnxOtnIntervalOtuFec15minIntervalNumber The value has been changed from read-only to not-accessible to be inline with newer MIBs. [PR1080802](#)
- In a VPLS scenario with specific CE mesh group configured, after a RE restart or RE master switchover, the flood NH for the mesh group might not be programmed properly. A complete black-holing for the VPLS instance would be seen as a consequence. [PR1087293](#)
- Deactivating/activating logical interfaces may cause BGP session flapping when BGP is using VRRP VIP as source address. This is caused by a timing issue between dcd and VRRP overlay file. When dcd reads the overlay file, it is not the updated one or yet to be updated. This results in error and dcd stops parsing VRRP overlay file. [PR1089576](#)
- In the dual REs scenario with GRES and ae0 interfaces configuration, if GRES is disabled on system, the backup RE should remove the ae0 bundle, however it doesn't go clean and ae0 remains in backup RE. After switching RE mastership to make other RE as master, the new master RE (which was backup earlier) continues to use invalid MAC address "00:00:00:00:00:00". [PR1089946](#)
- During failure notification state machine, CFM does not correctly transit from DEFECT CLEARING state to RESET once the error indication has been cleared. As a consequence all the forthcoming errors will be considered post errors and will be reported right away without incurring the fngAlarmTime. This is a cosmetic problem. [PR1096346](#)
- When DHCP subscribers are terminated at specific routing-instances and the interface stack is IP demux over vlan-subinterface over AE interface, there might be a memory leak in kernel AE iffamily when subscribers login/logout. [PR1097824](#)
- To ensure that the router or switch is reachable for management purposes while it boots or if the routing protocol process fails to start properly, we can configure a backup router, which is a router that is directly connected to the local router or switch (that is, on the same subnet) through its private management interface (for example, fxp0 or me0). When a backup router running IPv6 and a static route to reach the management network are configured, some invalid IPv6 routes are added to default forwarding-table on the master or the backup Routing Engine (RE). [PR1100981](#)
- The jpppd process might crash and restart due to a stale memory reference. The jpppd process restart results in a minimal impact of system and subscribers. All connected subscribers remain connected and only subscribers are attempting to connect at time of process restart would need to retry. [PR1121326](#)
- In Dynamic PPPoE subscriber management scenario, when the system is overloaded with requests coming, the subscribers might fail to login in a race condition. [PR1130546](#)
- The jpppd process might crash and restart due to a buffer overwrite. The jpppd process restart results in a minimal impact of system and subscribers. All connected subscribers

remain connected and only subscribers are attempting to connect at time of process restart would need to retry. [PR1132373](#)

Layer 2 Features

- In a high-scale VPLS configuration, modification of a tunnel interface through a restart or reconfiguration may cause the packet processing engine to access an invalid interface, resulting in minor packet loss and logging of packet processing engine traps. Existing traffic flows on the packet forwarding engine (PFE) are not affected. The router recovers quickly and normal operation resumes with the new configuration. [PR976972](#)
- RPD cores with 20 CFM session for interface in 10ms with AE configuration on the router [PR1051669](#)

Layer 2 Ethernet Services

- xSTP with NSB not supported on AE interfaces and so may observe multiple seconds of traffic drops during GRES switchover [PR1028313](#)
- After performing ISSU upgrading Junos OS version to afterwards of 13.3R4.2 14.1R3 or 14.2R1, the DHCP subscribers that were connected prior to the ISSU might unable to get traffic any more even though they have been allocated an IP address. Please note the subscribers connected after ISSU are not affected. [PR1031213](#)
- On a router with DHCP local server configured, if there are scaled number of DHCP subscribers connected, most of the subscribers might stuck in "RELEASE" status after performing graceful Routing Engine switchover (GRES). [PR1038385](#)
- There is a racing between DDOS-Protection and MAC-PINNING, when both feature is configured, MAC-PINNING discard notifications might get blocked by DDOS protection. [PR1042671](#)
- there was bug in handling the code while redistributing the xmit and adj entries to pman, when the interface entry was in pending distribution state. [PR1116741](#)

MPLS

- Given a Point-to-Multipoint branch LSP, the value of jnxMplsTeP2mpTunnelTotalUpTime is reported incorrectly after a new instance of the branch LSP is re-signaled at the ingress. [PR543855](#)
- Currently we allow configuration of both fast-reroute and link-protection/node-link-protection on a single LSP. But when customer configure both types of protection on the LSPs, it might cause scaling issues in customer network. Now, there is a change requirement, restrict the configuration to either fast-reroute or link/node-link protection on per LSP basis. [PR860960](#)
- When "soft-preemption" is enabled on ingress router and the preemption is configured with "aggressive" option to preempt RSVP sessions whenever bandwidth is lowered or a new higher-priority session is established, if one LSP is established over AE (aggregation ethernet) interface, in case of link failure to cause the bandwidth insufficient at the ingress, the CSPF is not triggered, owing to missing data in TED, to

establish a new path for more than 30 seconds and eventually the LSP is hard preempted. [PR1030586](#)

- LDP route delete is always communicated to resolver before resolver unresolves BGP nexthop via LDP tunnel. In this particular scenario which seems to stress Junos routing daemon, LDP route delete happens but resolver is yet to know about it. And resolver access the deleted LDP route entry. This leads resolver to access LDP route pointer that is already freed and leads to RPD core. [PR1097642](#)
- If a RSVP LSP has both primary& standby path and link-protection enabled, a /32 bypass route is unhidden when the primary link goes down. This /32 route is supposed to be made hidden again when primary link comes back up. But in some cases, due to software defect, this /32 bypass route remains unhidden forever which causes some issues, for example, BFD session down due to better prefix received from Bypass LSP. [PR1115895](#)
- User is allowed to configure both "load-balance-label-capability" and "no-load-balance-label-capability" together. This is incorrect and confusing. [PR1126439](#)

Multicast

- In multicast point-to-multipoint (P2MP) environment which has route -> indirect -> composite NH forwarding topology, if in case a change for composite NH happens before indirect NH addition itself, replication error is encountered on backup RE with error ENOENT (that is, does not exist), as a consequence, ksyncd process (ksyncd) will crash. [PR1036751](#)

Network Management and Monitoring

- When sessions are coming at a high rate, few of the syslogs are not logged. [PR868812](#)
- In some race conditions with firewall filters change, it is possible that the mib2d process receives a new TRIO filter ADD event before it learns about a non-TRIO filter DELETE event for the same filter index. The mib2d process will crash due to this. [PR1057373](#)
- SNMP engine-id is by default generated from default IP of the management interface. The failure in v3 authentication is due to different engine-ids on the master & backup members of the MX-VC setup. If the snmp engine-id need to be uniform across all the Routing-engines in the MX-VC set up, we could configure "set snmp engine-id local <local-engine-id-suffix>". [PR1059569](#)
- In rare cases, when the mib2d process attempts connection with the snmpd process and there are pending requests waiting to be finished, the mib2d process might crash and the CPU utilization is high around the same time as the crash happens. [PR1076643](#)
- In a certain mib view config where specific MIB OID "instances" are excluded from the MIB view, In a manner that leaves adjacent instances, included and excluded at the end of the MIB view. In this scenario when a snmp bulk-get request is made which covers these adjacent mib objects (at the end of View). The response may be malformed and may get dropped at snmpd. As a workaround specific mib instances must not be excluded at the end of mib view. Either the entire MIB or MIB object must be excluded or if in case a specific instance must be excluded ? it should not be at the end of mib view. [PR1126432](#)

Platform and Infrastructure

- Customer might see following error logs after deployment of Service Now (AIS). It is not service impacting and also not a memory leak issue. cscript process memory has exceeded 85% of RLIMIT_DATA: used 129132 KB Max 131072 KB [PR613704](#)
- The "old-rom-packet-count" command displays information about installed linecards. A non-zero number indicates that the bootrom on that linecard needs to be updated. [PR776703](#)
- Adaptive load-balance functionality is only supported for unicast traffic. If the aggregate bundle contains logical interfaces for bridge or vpls domains, flooded traffic might get dropped. [PR821237](#)
- When scripts are synchronized from one Routing Engine to the other, the destination for the scripts in the other Routing Engine should be based on the configuration on the other Routing Engine. This issue prevents this from happening and destination for scripts depends on the current Routing Engine from which the scripts were synchronized instead of the configuration on the other Routing Engine. [PR841087](#)
- The audit daemon (auditd) is the daemon which handles system accounting events and tries to send them out to configured RADIUS servers. If there is any problem in sending these accounting records to RADIUS (In this case RADIUS servers is unreachable or disconnecting frequently), auditd will spend more time on each accounting record because of the retries, and during this time if there are many accounting events, all those records will be in queue. And at one point of time, queue exceeds its limit and hence auditd crashes. [PR863697](#)
- On M/MX/T-Series platform with service-PIC. Kill the GPRS tunneling protocol controller daemon (gtpcd), the old Global Session Controller (GSC) card might lose contact. (Under normal circumstances, the GSC moves to other board). [PR868342](#)
- After the "show version detail" command is executed, the syslog message "UI_OPEN_TIMEOUT: Timeout connecting to peer" might appear. This message is cosmetic only; you can ignore this message. [PR895320](#)
- A crafted IP packet destined to an MX Series or T4000 router utilizing Trio or Cassis-based PFE (Packet Forwarding Engine) modules can cause the PFE to reboot. Affected modules include MPC1, MPC2, MPC3, and MPC4, integrated MPCs (CHAS-MX*), as well as Type 5 FPCs on the T4000. Please see JSA10621 for more information. [PR904887](#)
- When there is huge logical interface (IFL) scaling on AE (500 or more) with more than 32 member links and when all FPCs are restarted one by one, followed by member link addition to the link aggregation group (LAG), the state dependency evaluation in the kernel will take a long time given the scale involved, causing the FPCs not getting all the states from the Routing Engine (RE). This is a pretty uncommon sequence of events/conditions to happen and the likelihood of this happening in the field is very remote. [PR938592](#)
- Backing up the configuration with transfer-on-commit does not work in a MX-VC environment [PR947444](#)

- With enhanced LAG (As a part of enhanced-ip mode in 114.2) This issue is seen under certain conditions for traffic from IRB to AE [PR961685](#)
- The problem is seen because CFMD is getting a config commit after the MX-VC switch has happened. This commit is deleting the cfmd session and then creating a new session which is causing the old information of action-profile to be deleted which brings the interface back up. This problem fixes by the code correction. [PR974663](#)
- In the dual REs scenario, backup RE doesn't sync up the configuration change while deleting inactivated interface from the master. So after the operation, the inactivated interface still exists on the backup RE. [PR991081](#)
- rate-limit value does not match between RE and pfe [PR1023809](#)
- Still analyzing. [PR1024877](#)
- For MX platform with MPC/MIC interfaces or T4000 FPC5 with the Two-Way Active Measurement Protocol (TWAMP) reflector functionality activated, the sequence numbers generated by the reflector might not get from zero when the sessions start. Furthermore, TWAMP config will cause unrelated nexthop corruption. [PR1031201](#)
- When MX series router acts as the Broadband Remote Access Server (BRAS), if the PPP subscriber interface is configured over Aggregated Ethernet (AE) interface which spans multiple FPCs and there are member links within AE bundle are down, the router might not send out the control packet (such as LCP Echo-Reply packet when receiving the LCP Echo-Request packets from the subscriber), as a result, subscriber session would be torn down due to PPP keepalive failure. [PR1031218](#)
- MSDPC-HTTP redirect stops working [PR1039849](#)
- ISC BIND software included with Junos for SRX series devices is affected by CVE-2014-8500. This may allow a network based attacker to cause a denial of service condition on SRX devices. This issue only affects SRX devices where "set system services dns dns-proxy" has been configured. This is not enabled by default on SRX devices. This issue does not affect other Junos OS based devices as they do not have BIND DNS server feature. This issue has been assigned CVE-2014-8500. Please see <https://kb.juniper.net/JSA10676> [PR1048628](#)
- Software upgrade might cause firewall filters to redirect packets to an incorrect routing instance. This issue only affects Junos OS 12.3R7, 12.3R7-S1, 12.3R7-S2, 12.3R7-S3, and 12.3R8. [PR1057180](#)
- In MX virtual chassis (MXVC) scenario, there are two punt queues in VC ingress host path: high and low. Currently, VCCPd packets, VCP heartbeat packets, various protocol hello packets and TCP chassis control packets are directed to queue high. DDoS threshold for these queues default to 10K pps. If customer deploys large scaled route, during route change, the TCP control traffic would increase. As a result, DDoS threshold is reached, leading to packets drop. These packets drop might lead to VC link outage. [PR1058087](#)
- On MX Series router with frame-relay (FR) CCC to connect FR passport devices. If some of the FR circuits carry traffic without any valid FR encapsulations, the MX Series based Packet Forwarding Engine drops those frames. [PR1059992](#)

- If a firewall filter is configured with action 'then decapsulate gre routing instance <routing-instance>' and without other actions such as 'then [count | sample]' or 'then decapsulate gre [sample <protocol>] [forwarding-class <fc>] [no-decrement-ttl]', traffic loss might be seen if the board reboots or if the configuration for the routing instance and the filter are both applied at the same time. [PR1061227](#)
- In certain instances invalid flow data gets exported on any MX series router running 14.1R4 or 14.1R2 [PR1068598](#)
- On Trio-based line cards, when the firewall filters with prefixes are configured, the heap memory leak issue might be observed. [PR1073911](#)
- When deleting some uncommitted configuration on active RE, backup RE RPD process restart due to "Unable to proceed with commit processing due to SIGHUP not received. Restarting to recover" [PR1075089](#)
- On Trio-based platform, when learning the MAC address from the pseudo IFL (for example, label-switched interface), if the MAC address is aged out in source FPC where the MAC got learnt, due to the delay (around 2 to 3 milliseconds) of MAC address deleting message processed in the source FPC and the egress FPC (destination FPC of the traffic), the MAC address may be deleted first from the egress PFE but get added again during these 2-3 milliseconds time intervals (as there is continuous traffic coming on egress FPC destined to this MAC, the MAC query is generated and sent to RE and source FPC, since source FPC has not yet processed the MAC deleted message, it sends the response, so stale MAC will get added on the egress PFE), in this situation, no L2 flooding would occur for the "unknown" unicast (since the MAC address is present on the egress PFE). [PR1081881](#)
- Aggregate Ethernet (AE) interfaces in combination with shared-bandwidth-policer might lead to PFE policer corruption if there are child member links configured on the same PFE and the AE interface is being reconfigured (add/delete of logical units). This corruption could alter the policer rate programmed in hardware and lead to unexpected policer behavior. A different trigger with physical interface flap illustrating the same symptoms are tracked in PR1035845. [PR1084912](#)
- With MSDPC equipped on BNG, there might be a memory leak in ukernel, which eventually causes MSDPC to crash and restart. [PR1085023](#)
- Log reports "LUCHIP(0) RMC 0 Uncorrectable ECC 0x6db6db6d6db6db6d" and "PLCT INT_STAT 0x00000001 Invalid DMEM Address". The FPC may lose connections and need to be rebooted to clear the condition. [PR1086557](#)
- Under large-scale setup, VPLS MAC might not be aged-out from remote-PFE when local-PFE is MPC3/MPC4/MPC3E/MPC4E, then unknown-unicast frames flood will be seen on local PFE. [PR1099253](#)
- On MX series routers with MS-DPC equipped, traffic may be black holed because service IFLs and channel IFLs are not added correctly when the MS-PIC becomes offline/online. [PR1101432](#)
- Preventing an issue where one could end up with two "<junos:comment>" entries under the [interfaces] stanza. [PR1102086](#)

- The kernel next-hop acknowledgement timeout maximum interval configured (krt-nexthop-ack-timeout) under the CLI hierarchy "routing-options forwarding-table" has been increase to 400 seconds to avoid performance issues with scaled subscribers. [PR1102346](#)
- The following fields have been added to v10 Sampling (IPFIX) template and data packets: - SAMPLING RATE - SAMPLING INACTIVE TIMEOUT - SAMPLING ACTIVE TIMEOUT - TOTAL PACKETS EXPORTED - TOTAL FLOWS EXPORTED [PR1103251](#)
- Improved VTY commands to show internal JNH memory usage. [PR1103660](#)
- When MX2020 or MX2010 is running with FreeBSD10 based 15.1 JUNOS image, I2C error will be seen sporadically. tcabc i2c accelerator error: Group 0xX device 0xXX cmd timeout 984 usecs If the i2c error happens on voltage sensor, and it reaches count limit (9 times), chassis alarm will be shown up like this. 1 alarms currently active Alarm time Class Description 2015-09-10 06:42:40 UTC Minor CB 1 Volt Sensor Fail Those are cosmetic error but there is no way to clear the chassis alarm other than offline/online the FRU. [PR1122821](#)
- On ungraceful exit of telnet (quit/shell logout), perm and env files created by pam were not deleted. [PR1142436](#)

Routing Protocols

- When there are more than one tunnel PICs or tunnel interfaces 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](#)
- Continuous soft core-dump may be observed due to bgp-path-selection code. RPD forks a child and the child asserts to produce a core-dump. The problem is with route-ordering. And it is auto-corrected after collecting this soft-assert-coredump, without any impact to traffic/service. [PR815146](#)
- The multicast nexthop (show multicast nexthop) shown for Master and Backup RE for the same flow could be different if the nexthop is hierarchy MCNH. When doing NSR switch, however, there is no traffic loss caused by this show difference. [PR847586](#)
- In BFD over AE, multihop BFD or VCCV BFD (used for pseudo-wire) scenario, performing Graceful Routing Engine switchover (GRES) or In-Service Software Upgrade (ISSU) may result in the periodic packet management daemon (ppmd) crash. [PR917056](#)
- When a Junos OS router with multicast enabled receives IGMP packets with protocol DVMRP (IGMP_PROTO_DVMRP) to the IGMP port is 0x5 (DVMRP_ASK_NEIGHBORS2), IGMP builds a neighbor list and responds back to the source IP address of the sender. This source IP address can be a unicast address or a multicast address. There is no throttling of responses. The requests are answered at the highest rate possible. Secondary impacts are that the routing protocol daemon (rpd) IGMP utilization goes very high and the host path and interface network control queues can get congested. Refer to KB29553 for more information and mitigation. [PR945215](#)
- In BGP scenario with IPv4 and IPv6 neighbors mixed in the same group, if all of the IPv4 peers flaps but none of the IPv6 peers flaps, a timing issue might happen that one of the IPv4 peers comes up before inet.0 RIB is cleaned up, as a result, the routing protocol daemon (rpd) crash will be seen. [PR986272](#)

- Unified ISSU may fail from (< 12.3R6 and 13.1) to higher release(> 12.3R6, 13.2, 13.3, 14.1, 14.2) on legacy EX platforms if the following are configured 1) V6 access guard for STP (for <12.3R6) 2) Mobile edge (for 13.1) [PR1017987](#)
- The multicast traffic might be pruned with a static IGMP join configuration upon receiving an IGMP leave group message when the interface is not a querier on the corresponding interface. [PR1034270](#)
- Junos exhibits two different next-hop advertisement behaviors for MP_REACH_NLRI on a multi-hop eBGP session, based on whether its loopback peering or physical interface peering. When the routers are peering on their loopback, only the Global IP of the interface (lo0) is advertised where as when the routers are peering through the physical interface, both Global and Link-local address are advertised as the NHs. [PR1115097](#)

Services Applications

- In ipsec-vpn scenario, if backup-remote-gateway under [set services ipsec-vpn rule term then] hierarchy is configured, when the Internet Key Exchange (IKE) security association (SA) negotiation for the primary remote-gateway fails, the IKE tunnel failover might not be initiated hence the IKE tunnels are not established with any router. [PR849758](#)
- 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](#)
- kmd can core when IPsec tunnels are scaled with VRFs and DPD [PR865528](#)
- The SIP ALG does not recognize or translate the rare 'rtcp' attribute in the SDP payload, as a consequence non sequential RTP and RTCP ports are not supported. The RTP flow will be unaffected, not so for the RTCP control flow. [PR880738](#)
- The LCP state for tunneled subscriber is wrongly displayed as "OPENED" (which reflects the LCP state before tunneling) by CLI command "show interfaces pp0.<unit>" on the LAC. As a workaround, we can use "show ppp interface pp0." command to determine the correct LCP state for the subscriber. [PR888478](#)
- When using clear services nat mappings to clear implicit dynamic mappings (created for NAPT44 with PCP at full scale) more than one time, the SCHEDULER OINKER messages may be seen. These messages can be safely ignored. [PR923166](#)
- For PPTP ALG, data session ports are cleared only on CTRL session destroy. Even though data session expire, port is not freed [PR927362](#)
- Repeatedly applying the commands clear services nat mappings or clear services nat flows at maximum scale, immediately followed by scale traffic may result in stability issues or packet loss for short durations when using an MS-DPC, due to memory constraints. We recommend that you avoid such repeated usage of these clear commands at maximum scale. However, these problems do not occur when the commands are used freely at 70% load. [PR934580](#)
- Requires documentation update. When user configured multiple proposals (under IKE and IPSec configuration), the order in which proposals are configured, preference is

given in same order. For example, when peer proposes multiple proposal, the first proposal that matches the configured proposal in local device, will be chosen instead of picking strongest configured proposal. [PR947773](#)

- In L2TP scenario, when the LNS is flooded by high rate L2TP messages from LAC, the CPU on RE might keep too busy to bring up new sessions. [PR990081](#)
- On L2TP network server (LNS) router, if L2TP Control Channel Failure happens, The DestinationSessionCount might not get decremented and when it reaches the Destination maximum sessions, new L2TP subscribers will not be able to login anymore. [PR1025235](#)
- With scaling Layer 2 Tunneling Protocol (L2TP) sessions (for example, 128k sessions), when executing L2TP "show" command in one terminal and "clear" command in another terminal simultaneously, pressing Ctrl-C or closing the terminal on one terminal might cause the jl2tpd process crash. [PR1063207](#)
- With majority of L2TP subscribers login with invalid credentials (75% of new login requests are invalid), low call setup rate (CSR) will be observed for the good login attempt subscribers. [PR1079081](#)
- In Network address translation (NAT) environment, if the translation type is "dynamic-nat44", when processing bursts of packets (for example, packets coming in one after the other at a delay close to the interpacket gap, based on the replication, issue was seen when processing line rate traffic on a 1GE port), because the device may fail to free up dynamic NAT addresses, the address pool may get exhausted quite fast. In this situation, since no new addresses could be allocated, the incoming traffic drop might be seen on the device, also memory leak (not the main issue here, may become a problem way after the addresses leak) may occur on the device. [PR1098583](#)

Subscriber Access Management

- On BNG router, When the router is processing session "idle timeout", the following error message might be seen:
`./.././../src/junos/usr.sbin/authd/acc/authd_aaa_acc.cc:1273 Failed to process the Idle Timeout for session-id:10 Please note it will not affect any services.` [PR1041654](#)
- When using Neighbor Discovery Router Advertisement (NDRA) and DHCPv6 prefix delegation over PPPoE in the subscriber access network, if a local pool is used to allocate the NDRA prefix, when the CPE send DHCPv6 solicit message with both Internet Assigned Numbers Authority (IANA) and Identity Association Prefix Delegation (IAPD) options, the subscriber might get IPv6 prefix from the NDRA pool but not the delegated pool. As a workaround, the CPE should send DHCPv6 solicit message with only IAPD option. [PR1063889](#)
- In subscriber management environment with Remote Authentication Dial In User Service (RADIUS) server configured, when performing scaling subscribers login/logout, the device may stuck in RADIUS communication. [PR1070468](#)
- In subscriber management environment, when dual-stack service is activated by the Change of Authorization (CoA) request from the Radius Server, both families will be activated in the same profile response. Due to a software defect, the service accounting

session id is not generated properly and the Service Accounting Messages and Interim-updates failed to be sent out. [PR1071093](#)

- In large scaled DHCP subscribers environment (in this case, 110k DHCP users), the user executes logout and abort testing. In rare condition, the DHCP subscribers stuck in Terminating state. - logout testing: in this case the traffic generator sends DHCP Release for a subscriber and MX delete it from DHCP relay table, sdb, etc - abort testing: in this case the traffic generator just drop session and does not send DHCP Release message. From MX point of view DHCP subscriber is alive. [PR1073541](#)
- In subscriber management environment, after performing the graceful Routing Engine switchover (GRES), if the RE switchover happens before the Acct-Start response is received, and the timeout on service session happens before timeout on subscriber session, the authentication process (authd) may crash. [PR1074011](#)
- On MX series routers when adding the LI-Action attribute for mirroring the traffic of dual stack subscribers, due to the loop of the service requests lookup and adding, the authentication process (authd) CPU utilization may stay high indefinitely and traffic mirroring is not happening. [PR1077940](#)
- If authentication-order is configured as none under access profile and domain-name servers (DNS) are configured locally under access profile, then the subscriber will login but will not get DNS addresses which were configured locally. [PR1079691](#)

User Interface and Configuration

- 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](#)
- On HTTPS service J-web is not launching the chassis viewer page at Internet Explorer 7. [PR819717](#)
- On configure->clitools->point and click->system->advanced->deletion of saved core context on "No" option is not happening at jweb. [PR888714](#)
- Basic value entry format error check is not present in Configure-->Security-->IPv6 Firewall Filters, but the same is present in IPv4 Firewall Filters. But it will throw error when try to commit the wrong format data entered. [PR1009173](#)
- Customer wants visibility [PR1029477](#)
- When entering "restart r" incomplete command in CLI, even though there are multiple options available, command "restart routing" is executed finally. It should throw an error like "error: invalid daemon: r". [PR1075746](#)

VPNs

- 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](#)

- In NG-MVPN spt-only mode with a PE router acts as the rendezvous point (RP), if there are only local receivers, the unnecessary multicast traffic continuously goes to this RP and dropped though it is not in the shortest-path tree (SPT) path from source to receiver. [PR1087948](#)

Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

- [Resolved Issues 14.1R7 on page 93](#)
- [Resolved Issues: 14.1R6 on page 108](#)
- [Resolved Issues: 14.1R5 on page 121](#)
- [Resolved Issues: 14.1R4 on page 145](#)
- [Resolved Issues: 14.1R3 on page 154](#)
- [Resolved Issues: 14.1R2 on page 170](#)

Resolved Issues 14.1R7

Class of Service (CoS)

- On MX104 platform, when applying the "rate-limit" and the "buffer-size" on the logical tunnel (lt-) interface on the missing MIC (not inserted on MPC), commit failure with error message would occur. As a workaround, this issue could be avoided by applying the "rate-limit and "buffer-size" on inserted MIC, then commit. [PR1142182](#)

Forwarding and Sampling

- The command "clear firewall all" will now clear the policer stats displayed by "show policer __auto_policer_template_1__", ... "show policer __auto_policer_template_8__". [PR1072305](#)
- On MX Series platform with MX-FPC/DPC, M7/10i with Enhance-FEB, M120, M320 with E3-FPC, when there are large sized IPv6 firewall filters (for example, use prefix lists with 64k prefixes each) enabled, commit/commit check would fail and dfwd process

would crash after configuration commit/commit check. There is no operational impact. [PR1120633](#)

- On all Junos OS platforms, when both the filter and the policer are configured for an interface, in rare cases, the policer template may not be received by PFE (from the RE) when it is referenced by the filter term (normally the policer template gets received before the filter term referencing it which is ensured by mechanism in RE kernel). In this situation, the FPC would crash due to this rare timing issue. This issue might be avoided by the recommended steps below: 1. Deactivate the physical interface (IFD) and commit 2. Enable any filter and policer that attached to the interface (e.g. IFL) and commit 3. Activate interface back [PR1128518](#)

General Routing

- In a Layer 3 wholesale configuration, DHCPv6 advertise messages might be sent out with source MAC all zeroes if the subscriber is terminated on the demux interface in a non-default routing instance. For subscribers on default instance there is no such issue observed. [PR972603](#)
- On MX platform, MPC may crash when bringing up the 100-Gigabit Ethernet MIC with CFP2 (model number: MIC6-100G-CFP2) if initialization failure occurs (e.g. when bringing up the MIC6 which has hardware issue). [PR1037661](#)
- On all routing platforms M/MX/T/PTX with BGP configured to carry flow-specification route, in case of deleting a filter term and policer, then add the same term and policer back (it usually happens in race condition when adding/deleting/adding the flow routes), since confirmation from dfwd for the deleting policer might not be received before attempting to add the same policer, the rpd would skip sending an add operation for it to dfwd. As a result, when the filter term is sent to dfwd and tell it to attach to the policer, dfwd had already deleted the policer, and since rpd skipped re-adding it, dfwd will reject the attach filter with policer not found error and rpd will crash correspondingly. [PR1052887](#)
- As a precautionary measure, a periodic sanity check is added to FPC situated on M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX Series with DPC. It checks FPC error conditions and performs the appropriate actions in case of an error. [PR1056161](#)
- When a labeled BGP route resolves over a route with MPLS label (e.g. LDP/RSVP routes), after clearing the LDP/RSVP routes, in the short window before the LDP/RSVP routes restore, if the BGP routes resolves over a direct route (e.g. a one-hop LSP), the rpd process might crash. [PR1063796](#)
- Upon BFD flapping on aggregate interfaces, the Lookup chip (XL) might send illegal packets to the center chip (XMCHIP) and compromise packet forwarding and an FPC restart is needed to recover from this condition. If Fabric path side is affected, the fabric healing process will initiate this process automatically to recover from such conditions. MPC6E/MPC5E/NG-MPC are exposed to this problem. Corrupted parcels from Lookup chip LU/XL to Center Chip (XM) can also compromise packet forwarding and report DRD parcel timeout errors. An additional parcel verification check is added to prevent sending corrupted parcels to the center chip (XM) [PR1067234](#)

- The license-check process may consume more CPU utilization. This is due to a few features trying to register with the license-check daemon which license-check would not be able to handle properly and result in high CPU on Routing Engine (RE). Optimization is done through this fix to handle the situation gracefully so that high CPU will not occur. [PR1077976](#)
- Scheduler: Protect: Parity error for tick table single messages might appear on MPC3E/MPC4E/MPC5E/MPC6E/T4000-FPC5. [PR1083959](#)
- Wrong diagnostic optics info might be seen for GE-LX10 SFP and SFP+ for SumitomoElectric. The issue only for a specific SFP type - "Xcvr vendor part number : SCP6F44-J3-ANEÃ,Ã", it can be seen with "show chassis pic fpc-slot X pic-slot Y". [PR1091063](#)
- High latency might be observed when continuous IPv6 pings are sent to VMX platform. [PR1096403](#)
- In occasionally , AFEB PCI reads from Cortona MIC with ATM OAM traffic might return garbage values even though the actual content in the MIC has the correct value , this corrupted values would lead to AFEB crash , and also PCI error logs such as : afeb0 PCI ERROR: 0:0:0:0 Timestamp 91614 msec. afeb0 PCI ERROR: 0:0:0:0 (0x0006) Status : 0x00004010 afeb0 PCI ERROR: 0:0:0:0 (0x001e) Secondary bus status : 0x00004000 afeb0 PCI ERROR: 0:0:0:0 (0x005e) Link status : 0x00000011 afeb0 PCI ERROR: 0:0:0:0 (0x0130) Root error status : 0x00000054 afeb0 PCI ERROR: 0:0:0:0 (0x0134) Error source ID : 0x02580258 afeb0 PCI ERROR: 0:2:11:0 Timestamp 91614 msec. afeb0 PCI ERROR: 0:2:11:0 (0x0006) Status : 0x00004010 afeb0 PCI ERROR: 0:2:11:0 (0x004a) Device status : 0x00000004 afeb0 PCI ERROR: 0:2:11:0 (0x0052) Link status : 0x00004001 afeb0 PCI ERROR: 0:2:11:0 (0x0104) Uncorrectable error status : 0x00000020 afeb0 PCI ERROR: 0:2:11:0 (0x0118) Advanced error cap & ctl : 0x000001e5 afeb0 PCI ERROR: 0:2:11:0 (0x011c) Header log 0 : 0x00000000 afeb0 PCI ERROR: 0:2:11:0 (0x0120) Header log 1 : 0x00000000 afeb0 PCI ERROR: 0:2:11:0 (0x0124) Header log 2 : 0x00000000 afeb0 PCI ERROR: 0:2:11:0 (0x0128) Header log 3 : 0x00000000 [PR1097424](#)
- When the clock sync process (clksyncd) is stopped and resumed during link flaps, the clksyncd process might get into an inconsistent state with various symptoms, the clock source might be ineligible due to "Interface unit missing" or "Unsupported interface" with no Ethernet Synchronization Message Channel (ESMC) transmit interfaces. [PR1098902](#)
- Fragmenting a special host outbound IP packet with invalid IP header length (IP header length is greater than actual memory buffer packet header length), can trigger NULL mbuf accessing and dereferencing, which may lead to a kernel panic. [PR1102044](#)
- With Nonstop active routing (NSR) enabled, deleting routing-instance/logical system configuration might cause a soft assert of rpd. If NSR is not enabled, after deleting routing-instance/logical system configuration, executing "restart routing" might trigger this issue too. The core files could be seen by executing CLI command "show system core-dumps". This timing issue has no function impact. [PR1102767](#)
- When using "write coredump" to invoke a live coredump on an FPC in T-series, the contents of R/SR ASIC memory (Jtree SRAM) will get dumped. In the situation that there is a parity error present in the SRAM, then the coredump will abort and the FPC

will crash. As a workaround, configuring "set chassis pfe-debug flag disable-asic-sram-dump" before "write coredump" will help to avoid the issue.

[PR1105721](#)

- On MX-VC with heartbeat connection, if it is in a scaled subscribers environment, when power down both VCM REs, there might be a delay (minutes) for backup chassis to be master and during which time, traffic blackhole might be seen. [PR115026](#)
- On a busy MX Series Virtual Chassis platform, for example, with 100k subscribers and 16k subscribers concurrent login/logout, the ksyncd process might crash on Virtual Chassis backup Routing Engines after a local or global graceful Routing Engine switchover (GRES). This issue has no service impact. [PR115922](#)
- For MPC6E with CFP2, there was a race condition between the Interrupt service routine and the periodic, as a result interface up/down will not happen for laser off/on. [PR115989](#)
- On TX/TXP platform, when an LCC hit overtemp situation, it might go offline abruptly without notify SFC and other LCCs, which might cause traffic loss or performance degradation. Now with the fix, the overtemp situation on LCC is handled gracefully. [PR116942](#)
- On MX Series routers containing multiple PFEs (Packet Forwarding Engines) such as MX240/MX480/MX960/MX2010/MX2020, with either MPC3E/MPC4E/MPC5E/MPC6E cards, if the routers have GRE decap, then certain packet sizes coming via these aforementioned line cards, at very high rate can cause these line cards to exhibit a lockup, and one or more of their PFEs corrupt traffic towards the router fabric. [PR117665](#)
- The commit latency will increase along with the increasing lines under [edit system services static-subscribers group <group-name> interface]. Use ranges to create static demux interfaces is a recommend option. e.g. [edit system services static-subscribers group PROFILE-STATIC_INTERFACE] + interface demux0.10001001 upto demux0.10003000; [PR1121876](#)
- For scaled configuration, it may take too much time for commit and session gets hung because there is an unnecessary check to see if family Ethernet-switching co-exists with family bridge for all interfaces having bridge configuration. [PR1122863](#)
- This is a cosmetic issue that vMX firewall logs may show wrong packet length for dropped packets. [PR1124855](#)
- With BGP configured on CE-faced interfaces (in VRFs), doing 'show route' frequently may cause rpd to slowly leak memory. The leak rate will be one memory block of the size necessary to hold the instance name of the routing instance for a BGP neighbor. If the rpd process memory exhausted, the rpd process might crash, and the routing protocols are impacted and traffic disruption will be seen due to loss of routing information. You can check rpd memory usage with "show task memory brief" command. [PR1124923](#)
- In EVPN scenario, the EVPN route table between the master RE and backup RE would be different (unused garbage routes will appear) once RE switchover (e.g. by rebooting the "old" master RE or performing graceful routing engines switchover) is performed, which may cause kernel crash on the new master RE in some cases. [PR1126195](#)

- When Junos OS devices use Link Layer Discovery (LLDP) Protocol, the command 'show lldp neighbors' displays the contents of PortID Type, Length, and Value (TLV) received from the peer in the field 'Port Info', and it could be the neighbor's port identifier or port description. Junos CLI knob can select which 'interface-name' or 'SNMP ifIndex' to generate for the PortID TLV, so we do not have any problem as long as two Junos OS devices are connected for LLDP, but we might have an interoperability issue if other vendor device which can map the configured 'port description' in the PortID TLV is used. In such case, the Junos OS displays the neighbor's PortDescription TLV in the 'Port info' field, and if the peer sets 'port description' whose TLV length is longer than 33 byte(included), Junos OS is not able to accept the LLDP packets then discards packets as errors. The PortID TLV is given as : "the port id tlv length = port description field length + port id subtype(1B)" [PR1126680](#)
- In multi-homing Ethernet VPN (EVPN), if there are two loopback addresses and the router-id and the primary loopback addresses are different on the designated forwarder (DF) PE, when the link between CE and DF PE down, the Type 4 route of old DF are not deleted properly from the backup PE and causing the new DF election failure. The traffic forwarding will be affected. As a workaround, we should configure single primary loopback address and remove "router-id" knob on both multi-homing PEs. [PR1126875](#)
- On M320/T320/T640 with FPC 1/2/3 and their enhanced version (-E2/-E), in multicast scenario and AE interface is within multicast NH (such as, AE interface is the downstream interface for a multicast flow), egress multicast statistics displays incorrectly after flapping of AE member links. [PR1126956](#)
- In current Juniper implementation, the IPv6 multicast Router Advertisement timer is not uniformly distributed value between MinRtrAdvInterval and MaxRtrAdvInterval as described in RFC 4861. [PR1130329](#)
- On MX with MS-MIC (or possibly, MS-MPC is affected as well), changing configuration of sampling input parameters, such as "rate" under forwarding-options is not reflected without restarting the line card. [PR1131227](#)
- On Trio based line card, multiple modifications of firewall filter might cause lookup chip error and traffic blackhole, following jnh_free error messages could help to identify this issue: messages: fpc1 jnh_free(10212): ERROR [FW/3]:1 Paddr 0x006566a9, addr 0x2566a9, part_type 0 call_stack 0x40497574 0x418ffa84 0x41900028 0x418ecf94 0x41861690. [PR1131828](#)
- When customers do changes under "protocol router-advertisement interface X" (such as changing timers etc), they expect that commit would trigger a new router-advertisement being sent out to notify hosts about configuration changes. However it does not seem to be a case unfortunately. It makes the router information to expire on hosts and causes obvious loss of connectivity for the hosts. [PR1132345](#)
- On MX platform with non-Q MPC (for example, MPC2-3D) or Q-MPC with enhanced-queueing off, when traffic has to egress on any one of the dynamic PPPoE (pp0), IP-DEMUX (demux0) and VLAN-DEMUX (demux0) IFLs, the queue mapping might get wrong. The traffic forwarding might be affected. [PR1135862](#)

- MXVC-Same subnet VC-heartbeat polling failed to recover [PR1136119](#)
- Upon either FPC boot with 100G CFP2 MIC installed or upon 100G CFP2 MIC installation post FPC boot, if the MIC is unable to initialize correctly, the MPC6E can crash and then restart. This has only been seen when the MIC has suffered a hardware failure. [PR1148325](#)

High Availability (HA) and Resiliency

- On MX Series Virtual Chassis (MX-VC) with scaled configuration, for example, 110000 DHCP and 11600 PPP subscribers, the unified in-service software upgrade (ISSU) might fail due to the management daemon (MGD) timer expiring before Field-replaceable units (FRUs) update finish. [PR1121826](#)
- With NSR enabled on multiple RE system, when dynamic GRE tunnel is configured, performing RE switchover might causing rpd crash repeatedly on backup RE. [PR1130203](#)

Infrastructure

- When "show version detail" cli command has been executed, it will call a separate gstatd process with parameter "-vvX". Because the gstatd could not recognize these parameters, it will run once without any parameter then exit. [PR1078702](#)
- The Remote NFS Server process (nfsd) is not terminated on the new backup Routing Engine (RE) after RE switchover. As a result, it spawns a new one upon RE switchover until running out of memory. [PR1129631](#)

Interfaces and Chassis

- From Junos OS release 13.3R1, any host outbound traffic go across AE interface, will lead to ifstat memory leak, finally result in kernel crash. Please use the command "show system virtual-memory | match "mem|ifstat "" to monitor the memory utilization. [PR975781](#)
- In the bridge domain configuration with IRB interface environment, the IRB interface INET/ISO MTU is set to 1500. When the MTU on IRB interface is deleted, the MTU wouldn't be changed. [PR990018](#)
- On DPC only chassis, after software upgrade or not graceful Routing Engine switchover, Ethernet OAM related LAG bundles might not come up due to the Link Fault Management (LFM) packets arrive on AE interface instead of physical link interface. [PR1054922](#)
- MS-DPC might crash when allocating chain-composite nexthop in enhanced LAG scenario. [PR1058699](#)
- During subscriber login/logout the below error log might occur on the device configured with GRES/NSR. /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV (pp0.1073751222) [PR1058958](#)
- With "enhanced-ip" mode and AE interface configured, if SCU/DCU accounting is enabled, the MS-DPC might drop all traffic as regular discard. [PR1103669](#)
- The 'optics' option will now display data for VCP ports: show interfaces diagnostics optics vcp-0/0/0 [PR1106105](#)

- On MX240 or MX480 platform with at least two DC modules (PN: 740-027736) equipped, when shutting down one of the PEMs and then turn it on again, even the PEM is functioning, the "PEM Fan Fail" alarm might be observed on the device due to software logic bug. There is no way to clear the ALARM_REASON_PS_FAN_FAIL for I2C_ID_ENH_CALYPSO_DC_PEM once it has been raised. [PR1106998](#)
- On all Junos platform, if the "HDD /var" slice (for example, "/dev/ad1s1f" depending on the type of RE) is not mounted (for example, label missing, file system corrupted beyond repair, HDD/SDD is removed from the boot list, etc), the system may build emergency "/var/", however, no alarm or trap is generated due to the incorrect operation of the ata-controller. Although the boot messages may present the logs, it may not be sufficient enough to identify the issue before encountering other problems (for example, Junos OS upgrade failure and the Routing Engine may hang in a recovery shell). In addition, as a method to check where Routing Engine is running from, a manual check could be done as below, user@re0> show system storage | match " /var\$ " /dev/ad2s1f 34G 18G 13G 57% /var <<<Indicate that "/var" is mounted from the HDD/SSD

user@re0>show system storage | match " /var\$" <<<<No output here, it means that the RE is running from "emergency /var" [PR1112580](#)
- Junos OS now checks ifl information under the ae interface and prints only if it is part of it. [PR1114110](#)
- On MX Series platform, when using Ethernet OAM Connectivity Fault Management (CFM), the CFM process (CFMD) may crash in either of the following scenarios, - Scenario 1 When CFMD is restarted or GRES. There is no specific defined configuration which could cause this crash, but normally this would be seen with VPLS or Bridge domain with multiple Mesh-groups. The crash happens rarely in this scenario. - Scenario 2 When configuring 2 interfaces in the same bridge-domain (BD) or routing-instance, and both interfaces have maintenance association end point (MEP) configuration along with action-profile enabled. Also there is no maintenance association intermediate point (MIP) configuration on that BD or routing-instance. The crash might be seen with the above configurations and when one of the interfaces is flapped or deleted and then re-created. In addition, in this scenario, this issue may not happen always as this depends on the ordering of kernel event. [PR1120387](#)
- On JUNOS platform, an aggregate-ethernet bundle having more-than one member link can show incorrect speed which wouldn't match to the total aggregate bandwidth of all member links. The issue would be seen when LFM is enabled on the aggregate-ethernet bundle. The issue would be triggered when one of the member link flaps. Although after the flap, the current master Routing Engine would show correct aggregate speed, the backup Routing Engine would report incorrect value. In this state, when Routing Engine mastership is switched, the new master Routing Engine (which was backup) will show incorrect value. One of the side-effect of this issue is that RSVP also reflects incorrect Bandwidth availability for the affected aggregate-ethernet bundle, thus can cause under-utilization of the link with LSP having bandwidth constraints. [PR1121631](#)
- Since a bug which was introduced in 15.1R1, loopback sub-interfaces always have a Flag down in the output of CLI command "show interfaces". [PR1123618](#)

Layer 2 Features

- When AE is core facing ifl in ldp-mesh vpls instance with local-switching in it, the traffic is looped back. [PR1138842](#)
- In a VPLS scenario, when "\$junos-underlying-interface-unit" is configured in "dynamic-profiles" hierarchy, which is then implemented in a routing-instance. The upgrade/commit will fail with the following error message, Parse of the dynamic profile (dynamic-profiles) for the interface: \$junos-interface-ifd-name and unit: \$junos-underlying-interface-unit failed! [PR1147990](#)

Layer 2 Ethernet Services

- On MX series platform with none-stop-routing (NSR) enabled and some L2 protocols configured, performing RE switchover might cause layer 2 control protocol daemon (l2cpd) to crash and FPC to be rebooted. [PR1076113](#)
- On MX platform with Dynamic Host Configuration Protocol (DHCP) maintain subscriber feature enabled, after rebooting the FPC hosts the Demux underlying interfaces, the next-hop for some DHCP subscribers might be marked as dead in the forwarding table. When this issue occurs, we can execute CLI command "clear dhcp server binding <address>" to restore. [PR1118421](#)
- For PVSTP/VSTP protocols, when MX/EX92xx router inter-operate with Cisco device, due to the incompatible BPDU format (there are additional 8 Bytes after the required PVID TLV in the BPDU for Cisco device), the MX might drop these BPDUs. [PR1120688](#)
- In the DHCPv4 or DHCPv6 relay environment with large scaled environment (in this case, 50-60K subscribers), and the system is under stress (many simultaneous operations). The subscribers might get stuck in RELEASE state with large negative lease time. [PR1125189](#)
- In some rare scenarios, the MVRP PDU might unable to be transmitted, which could cause memory leak in layer 2 control plane daemon (l2cpd), and finally results in the l2cpd process crash. [PR1127146](#)
- Input/Output pps/bps statistics might not be zero after a member link of AE interface with distributed pppmd was down in M320/T-Series(GIMLET/STOLI based FPC) [PR1132562](#)

Multiprotocol Label Switching (MPLS)

- With egress protection configured for Layer 3 VPN services to protect the services from egress PE node failure in a scenario where the CE site is multihomed with more than one PE router, when the egress-protection is un-configured, the egress-protection route cleanup is not handled properly and still point to the indirect composite nexthop in kernel, but the composite nexthop can be deleted in rpd even the egress protection route is pointing to the composite nexthop. This is resulting in composite nexthop "File exists" error when the egress protection is re-enabled and reuse the composite nexthop (new CNH addition fails as old CNH is still referenced in kernel). [PR954154](#)
- MPLS auto-bandwidth does not reset MAX Avg Bandwidth when overflow or underflow threshold limit is configured. It may lead to wrong bandwidth reservations occasionally. [PR954663](#)

- In next-generation MVPN extranet scenario, if there is a mix of VT interface and LSI (vrf-table-label is used) interface on next-generation MVPN egress node, after changing some vrf policies, the routing protocol process (rpd) might crash and reset. [PR1045523](#)
- In MPLS scenarios, removing the "family mpls" configuration from an outgoing interface may cause inet and/or inet6 nexthops associated with that interface to unexpectedly transit to dead state. Even adding back "family mpls" cannot restore it. [PR1067915](#)
- When an LSP is link-protected and has no-local-reversion configured, if the primary link (link1) is down and LSP on bypass (link2), then another link (link3) is brought up, before the LSP switch to link3, if link1 is enabled and link3 is disabled, the LSP will stuck in bypass LSP forever. This is a timing issue. [PR1091774](#)
- From Junos release 13.2R1 and above, in MPLS L3VPN scenario, when "l3vpn-composite-nexthop" knob is enabled on a PE router and an interface style service set is attached to the ingress interface, the L3VPN packets with the MPLS labels will be sent to the service card and dropped. As a workaround, we should disable "l3vpn-composite-nexthop". [PR1109948](#)
- If "optimize-timer" is configured under P2MP branch LSP, this branch LSP will not be re-established if link flap on egress node. If "optimize-timer" is configured at protocols/mpls level, issue could be avoided. [PR1113634](#)
- For an MPLS L3VPN using LDP-signaled LSPs, in a rare racing condition (e.g. large-scale environment or Routing Engine CPU utilization is high), the rpd process might crash after an LDP neighbor down. [PR1115004](#)
- When multipoint LDP (M-LDP) in-band signaling is enabled to carry multicast traffic across an existing IP/MPLS backbone and routing process is enabled to use 64bit mode, the rpd might crash due to accessing an uninitialized local variables. [PR1118459](#)
- When an PLR is a non-Juniper router, Juniper ingress node might stay on the bypass tunnel and ignore the CSPF result. [PR1138252](#)
- There is no entropy label for LDP route in scenario of LDP tunneling across a single hop RSVP LSP with label 0 (explicit-null) used. As workaround, either remove LDP tunneling or RSVP explicit-null will resolve the issue. [PR1142357](#)
- This issue is related to inter-op between multi vendor scenario. This fix will add sub-object RRO which will help change of label during FRR active scenario. [PR1145627](#)

Network Management and Monitoring

- The SNMPv3 message header has a 4 byte msgID field, which should be in (0....2147483647), when the snmpd process has been running for a long time, the msgID might cross the RFC defined range and causing Net-SNMP errors, "Received bad msgID". [PR1123832](#)

Platform and Infrastructure

- On Trio based line card, when GRE keepalive packets are received on a PFE that is different from the tunnel interface hosted, the keepalive message will apply the firewall filter configured on default instance loopback interface. [PR934654](#)
- Bad udp checksum for incoming DHCPv6 packets as shown in monitor traffic interface output. The UDP packet processing is normal, this is a monitor traffic issue as system decodes checksum=0000. [PR948058](#)
- Under certain conditions the PFE flow export thread and flow update thread might be out of sync resulting in a situation where the update thread might attempt to update a flow record that is being aged-out/deleted by the export thread. As a consequence, PFE traps might be generated during flow processing; the PFE trap signature is very dependent on the operation performed on that particular record: fpc1 PFE Sync TXN Err Trap: Count 3, PC 637f, 0x637f: flow_export_read_src_address_ipv6 LUCHIP(2) PFE_4 Errors sync txtn error Under rare conditions, this can ultimately lead to record corruption. Trying to reuse or update such a record would trigger the following error: [LOG: Err] LUCHIP(2) HASH INT Status FPM Error: [LOG: Err] LUCHIP(2) HASH FPM ERROR: Alloc OMI Ram IF Error, TID=1, FP_ID=0x2. - There is no impact to forwarding. - There may or may not be impact on Jflow. - Its a generic problem for any inline-jflow application including IPv4 and IPv6. With 13.2 release, new fields (min, max TTLs/QinQ values) are added to jflow record. These fields need to be updated (if value changes) per packet in the flow. So the probability of hitting the race condition between export thread (deleting the record) and jflow data path code (updating the same record) and is higher in 13.2 release onwards. [PR968807](#)
- In the dual REs scenario with NSR configuration, the knob "groups re0 interfaces fxp0 unit 0" is configured. If disable interface fxp0, backup RE is unable to proceed with commit processing due to SIGHUP not received, the rpd process on backup RE might crash. [PR974430](#)
- When netconf or Junos OS scripts are used to manage the device, the management process gets stuck in a loop, causing high CPU usage. [PR991616](#)
- On MX Series Virtual Chassis (MX-VC) platform, mirroring of OAM packets may not work as expected if the OAM packet traversing through multiple Packet Forwarding Engines (for example, the mirrored port and VCP port are on separate PFEs). [PR1012542](#)
- when one of the "deny-commands" is incorrectly defined on the profile of TACACS+ server, all "deny-commands" regexes will be ignored, which leads to an over-permissive profile without any warning. [PR1078238](#)
- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)

- When BFD or VRRP is running on a multi LU (lookup chip) PFE (such as MPC3 or MPC4), some incoming BFD or VRRP packets might be incorrectly evaluated by a firewall filter configured on a loopback interface of a different logical system or routing instance. Therefore, packets might be unexpectedly discarded leading to session/mastership flaps. [PR1099608](#)
- On an MPC3E or MPC4E or on an EX9200-2C-8XS line card, when the flow-detection feature is enabled under the [edit system ddos-protection] hierarchy, if suspicious control flows are received, two issues might occur on the device: ? The suspicious control flow might not be detected on the MPC or line card. ? After suspicious control flows are detected, they might never time out, even if traffic flows no longer violate control parameters. [PR1102997](#)
- Junos defines SNMP ifXTable (ifJnxInErrors/ifJnxInL3Incompletes) counter as 64-bit width, but it worked as 32-bit width counter. It works as 64-bit width counter after the fix. [PR1105266](#)
- On Trio-based platform, in MX Series Virtual Chassis (MXVC) environment, if the subscriber logical interface (IFL) index 65793 is created (for example, when carrying 15K DHCPv4 subscribers to exceed IFL index creation 65793) and the IEEE 802.1p rewrite rule is configured (for example, using CoS rewrite rules for host outbound traffic), due to usage of incorrect IFL index, the Virtual Chassis Control Protocol Daemon (vccpd) packets (for example, Hello packets) transmission may get lost on all VC interfaces, which may lead to VC decouple (split brain state, where the cluster breaks into separate parts). As a workaround, either delete the rewrite rule (delete class-of-service host-outbound-traffic ieee-802.1 rewrite-rules), or find the IFL in jnh packet trace that is not completing the vccpd send to other chassis and at RE clear that subscriber interface may resolve the issue. [PR1105929](#)
- CVE-2015-5477 A vulnerability in ISC BIND's handling of queries for TKEY records may allow remote attackers to terminate the daemon process on an assertion failure. See <http://kb.juniper.net/JSA10718>. [PR1108761](#)
- On MX-VC, when traffic with TPID 0x88a8 or 0x9100 is sending over AE interface, the packets which across VCP links might be dropped on egress VCP PFE due to invalid fabric token. [PR1112752](#)
- When inline static NAT translation is used, if two rules defined in two service sets are pointing to the same source-prefix or destination-prefix, changing the prefix of one of the rule and then rolling back the changes is not changing back all the pools correctly. [PR1117197](#)
- On Trio-based line card, the firewall filter may have some issues when matching on Authentication Header (AH) protocol. This can affect VRRP (among others) when authentication is used, and an Routing Engine (RE) firewall filter is matching on protocol AH. As a workaround, we can change the filter to match on other criteria (e.g. source or destination address). [PR1118824](#)
- When the AC Single Phase Power Distribution Module (PDM) is installed on an MX2010 or MX2020 router running Junos release 14.1 or 14.2, the system does not recognize the FRU and alarms are triggered as a result. [PR1121068](#)

- After changing an outer vlan-tags, the ifl is getting programmed with incorrect stp state (discarding), so the traffic is getting dropped. [PR1121564](#)
- With "fast-synchronize" configured, adding a new configuration-group that has configuration relevant to the rpd process and apply it and commit, then any configuration commits might cause the rpd process on the backup Routing Engine (RE) crash. We can reboot the backup RE to restore. [PR1122057](#)
- On Trio-based platform, when fragmented packets go through the inline NAT (including source NAT, destination NAT, and twice NAT), the TCP/UDP checksum would not be correctly updated. In this situation, checksum error would occur on the remote end (inside and outside device). Non-fragmented packets would not be affected by the issue. If possible, this issue could be avoided by either of the following workarounds,
* Enable "ignore-TCP/UDP-Checksum errors" at the inside or outside device which processes TCP/UDP data OR
* Make sure there will not be any fragments subjected to inline NAT functionality by appropriate MTU adjustment or setting [PR1128671](#)
- Parity error at ucode location which has instruction init_xtxn_fields_drop_or_clip will lead to a LU Wedge. LU is lookup ASIC inside the Trio. The LU wedge will cause the fabric self ping to fail which will lead to a FPC reset. This is a transient HW fault, which will be repaired after the FPC reset. There is no RMA needed unless the same location continues to fail multiple times. [PR1129500](#)
- On Junos device with DHCP Relay config but without accounting config, and the accounting license does not exist, when the first DHCP control traffic is received, the following subscriber-accounting license grace period alarms might be triggered:
alarmd[1650]: Alarm set: License color=YELLOW, class=CHASSIS, reason=License grace period for feature subscriber-accounting(30) is about to expire craftd[1592]: Minor alarm set, License grace period for feature subscriber-accounting(30) is about to expire [PR1129552](#)
- For IPv6 packet with "no next header" in Hop-By-Hop header, if the Hop-By-Hop header length field value is large than 112, the router will drop such packet and log the following error: PPE PPE HW Fault Trap: Count 105, PC 60ce, 0x60ce: ipv6_input_finished_parsing LUCHIP(3) PPE_10 Errors lmem addr error [PR1130735](#)
- NTP.org published a security advisory for thirteen vulnerabilities in NTP software on Oct 21st, 2015. These vulnerabilities may allow remote unauthenticated attackers to cause Denial(s) of Service(s), disruption of service(s) by modification of time stamps being issued by the NTP server from malicious NTP crafted packets, including maliciously crafted NTP authentication packets and disclosure of information. This can impact DNS services, as well as certificate chains, such as those used in SSL/https communications and allow attackers to maliciously inject invalid certificates as valid which clients would accept as valid. Refer to JSA10711 for more information. [PR1132181](#)
- Too many duplicate ACK messages are generated from PFE for TCP control connection with RE. This could cause: 1. MX-VC DDoS protection violation for VC-control low queue and makds MXVC split. 2. Cause RE and FPC high CPU utilization. [PR1133293](#)
- With scaled firewall filters attached to interfaces (e.g. 10k+ filters), running "show configuration" command can cause high CPU of the mgd process. As a workaround, we can use "show configuration |display set" command to view the config. [PR1134117](#)

- PPE thread timeout trap may cause XM chip wedge, it will not affect MQ based FPC. [PR1136973](#)
- On MX2020, when we remove whole power of a power zone, and then put the power back to the zone, FANTray LED stays Amber and FANTray LED on craft card stays OFF, and do not revert to green (FANTray LED) or ON (Craft LED) until we reboot the entire chassis system or hot swap that FAN tray. For Zone 0 (PSM 0 to 8), FAN 1 shows the above described behavior. For Zone 1 (PSM 9 to 17), FAN 3 shows the above described behavior. [PR1138209](#)
- When the cli command "show pfe statistics exceptions | match reject" executed CPROD thread in the PFE may hogg the CPU and result in FPC crash [PR1142823](#)

Routing Protocols

- BGP "accepted-prefix-limit" feature might not work as intended when it is configured together with "damping". Root cause of this issue is that when BGP module count the maximum routes accepted from BGP neighbor, it doesn't count the accepted BGP routes which in damping status. So when these damping routes are reused, the total number of received BGP routes exceeds the configured value for "accepted-prefix-limit". [PR897124](#)
- Since Junos 13.3R2 and higher if delegated BFD sessions are flapping continuously, packet buffer memory maybe be leaked. The automatic memory leak detection process will report this within the syslog once certain threshold is reached "fpc7 SHEAF: possible leak, ID 8 (packet(clones)) (10242/128/1024)". Please note BFD sessions operating in centralized mode are not exposed. [PR1003991](#)
- From Junos release 14.1R1 or above, the rpd process might crash while executing CLI command "show isis backup spf results". [PR1037114](#)
- EDITED MP 8/31 When a multicast group in protocol independent multicast (PIM) dense mode has a large number of multicast sources, the RPD process can crash after a routing engine switchover. [PR1069805](#)
- On large scale BGP RIB, advertised-prefixes counter might show the wrong value due to a timing issue. [PR1084125](#)
- When a BGP session supports multiple address families, the inactive route of some of the address families might not be flushed correctly, leading to wrong behaviors for some of the features which need to advertise inactive routes(e.g. advertise-inactive, advertise-external, optimal-route-reflection, etc). [PR1097297](#)
- Due to software bug Junos cannot purge so called doppelganger LSP, if such LSP is received over newly formed adjacency shortly after receiving CSNP from the same neighbor. [PR1100756](#)
- When two (or more) route target communities of MP-BGP route match to two (or more) route target communities in VRF import policy of a RI duplicate routing entries might be installed in the RI. In the output of 'show route table <RI-name>.inet.0 detail' two identical routing entries appear with one being marked as 'Inactive reason: Not Best in its group - No difference'. When such duplicate routing information is to be deleted, rpd process process will crash. [PR1113319](#)

- When the Multicast Source Discovery Protocol (MSDP) is used, if the RP itself is the First-Hop Router (FHR) (i.e. source is local), the MSDP source active (SA) messages are not getting advertised by the RP to MSDP peers after reverse-path forwarding (RPF) change (e.g. the RPF interface is changed). [PR1115494](#)
- When a logical unit of an interface is associated with a Bidirectional Forwarding Detection (BFD) session, if changing the unit number of the interface (for example, change the unit number for a running BFD session from ge-1/0/0.2071 to ge-1/0/0.285), the device may fail to change the name due to the missing check for logical interface (IFL) index change. [PR1118002](#)
- On dual Routing Engine platform with Nonstop active routing (NSR) and authentication of the Bidirectional Forwarding Detection (BFD) session enabled, BFD process (bfdd) memory leak may occur on the master RE and the process may crash periodically once it hits the memory limit (RLIMIT_DATA). The problem does not depend on the scale, but the leak will speed up with more BFD sessions (for instance 50 sessions). As a workaround, if possible, disabling BFD authentication will stop the leak. [PR1127367](#)
- In multicast environment, when the RP is FHR (first hop router) and it has MSDP peers, when the rpf interface on RP changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)
- In multicast environment with Protocol Independent Multicast sparse mode (PIM SM) used, if a upstream router of last-hop router receives the (S,G) SPT join while the shortest-path tree (SPT) is not yet established (only because multicast source is not reachable, a reachable route for SPT which is just not established yet will not cause this issue), when the multicast route get deleted on the router (e.g. receives the (S,G) prune from downstream PIM router), the router would incorrectly stop forwarding the multicast traffic even if rendezvous-point tree (RPT) path exists. [PR1130279](#)
- Pending RN from Development [PR1135205](#)
- In rare condition, mt tunnel interface flap cause backup RE core. The exact root cause is not known. While processing updates on the backup RE (received from master RE), accessing free pointer cause the Core. [PR1135701](#)
- On dual Routing Engine (RE) platform with Bidirectional Forwarding Detection (BFD) protocol enabled, after graceful Routing Engine switchover (GRES), the periodic packet management process (ppmd) might crash on backup RE due to a software defect. [PR1138582](#)
- In the BGP labeled unicast environment, the secondary route is configured with both add-path and advertise-external. If the best route and secondary route are changed in a routing table at the same time, add-path might miss to readvertise the changed route. The old route with the old label is still the last route advertised to one router instead of updating the advertisement with the new route and new label. So the traffic forwarding might be affected. [PR1147126](#)

Routing Policy and Firewall Filters

- When a malformed prefix is used to test policy (command "test policy <policy name> <prefix>"), and the malformed prefix has a dot symbol in the mask field (e.g. x.x.x.x/24), the rpd process might crash. [PR1144161](#)
- From Junos OS release 13.2R1, an attempt to commit a configuration with a dangling conditional policy referring a non-existent/inactive routing-instance will be permitted. If we have a conditional policy referring an active routing-instance, deleting/deactivating this routing-instance and then committing will cause the rpd process crash. As a workaround, we should always make sure that conditional policies are referring active routing-instances. [PR1144766](#)

Services Applications

- Junos 13.3 and above release, when configuring a /31 subnet address under a nat pool, the adaptive services daemon (SPD) will continuously crash. [PR1103237](#)
- If l2tp is configured under access-group hierarchy, during commit or commit check operation, the pppd process might crash (the configuration could commit successfully). It might result in a minimal impact of system, and it will restore automatically. As a workaround, please configure the same under the access profile client hierarchy. [PR1108024](#)
- SIP one way audio calls when using X-Lite SIP Softphone, in case that SIP media is switched to another media gateway through a SIP RE-Invite message [PR1112307](#)
- In CGNAT environment, when a service PIC is in heavy load continuously, there might be a threads yielding loop in CPUs, which will cause the CPU utilization high, and might cause one of the CPUs to be reset. [PR1115277](#)

Software Installation and Upgrade

- In certain conditions, when /var is not mounted from a persistent filesystem, executing a Junos upgrade will have unexpected results. This is caused by an inexact check of whether it is running from an Emergency VAR. [PR1112334](#)

Subscriber Access Management

- For scenarios that are not in a Layer 3 wholesale network environment, we can configure "duplication-vrf" to send duplicate accounting records to a different set of RADIUS servers that reside in either the same or a different routing context. After Routing Engine switchover, the duplicate accounting feature stops work for existing subscribers. [PR1121524](#)
- In subscriber management environment, the authentication process (authd) crash may occur. This issue is not reproduced yet, possibly, it might be seen when generating a CLI Change of Authorization (CoA) request (e.g. via CLI command "request network-access aaa subscriber add service-profile filter-service session-id 10"), then logging out the subscriber (the one with service just activated), if the management CLI session is closed before subscriber entry is reused, the crash may occur. [PR1127362](#)

User Interface and Configuration

- When committing a config with very long as-path, in this case the as-path is almost 12000 characters long, the commitd process might crash. The commitd process restart results in a minimal impact of system. As a workaround, please config as-path less than 4096 characters long. [PR1119529](#)

VPNs

- In a multi-homed source topology in NG-MVPN (applicable to both inter-AS and intra-AS scenario), there are two problems: The first problem is Multicast (S, G) signaling doesn't follow RPF. When the routing table (mvprninstancename.inet0) has two routes, due to the policy configuration, the best route to the source is via the MPLS core, but Multicast (S, G) PIM join and NG-MVPN Type 7 both point to inactive route via local BGP peer. The second problem is when "clear pim join instance NG" is entered, the multicast forwarding entries are wiped out. [PR1099720](#)
- In scenario involving pseudowire redundancy where CE facing interface in the backup neighbor (can be non-standby, standby, hot-standby type), if the virtual circuit (VC) is not present for the CE facing interface, the CE facing interface may go up after committing an unrelated VC interface configuration (e.g. changing description of another VC interface) even though the local pseudowire status is in down state. [PR1101886](#)
- In Global Table Multicast (GTM) scenario (instance-type mpls-internet-multicast), when the GTM instance and master instance are used, if the name of the GTM instance is changed, the routing protocol process (rpd) may crash due to the usage of the incorrect routing table handle. [PR1113461](#)
- In L2circuit environment, if one PE has pseudowire-status-tlv configured but remote hasn't, and at the same time, this PE doesn't support control-word but remote does, then it will not send changed local status code to remote PE, in a rare condition, after enable status-tlv support at remote end, the L2circuit might stuck in "RD" state on remote PE. [PR1125438](#)

Resolved Issues: 14.1R6

- [Class of Service \(CoS\) on page 109](#)
- [Forwarding and Sampling on page 109](#)
- [General Routing on page 109](#)
- [Infrastructure on page 112](#)
- [Interfaces and Chassis on page 112](#)
- [Layer 2 Features on page 115](#)
- [Layer 2 Ethernet Services on page 115](#)
- [MPLS on page 115](#)
- [Network Management and Monitoring on page 116](#)
- [Platform and Infrastructure on page 116](#)
- [Routing Policy and Firewall Filters on page 119](#)

- [Routing Protocols on page 119](#)
- [Services Applications on page 120](#)
- [Software Installation and Upgrade on page 121](#)
- [Subscriber Access Management on page 121](#)
- [VPNs on page 121](#)

Class of Service (CoS)

- After restarting chassisd or doing an in service software upgrade from 13.2R8.2 to 13.3R7.3 results in the following messages seen in syslog:
cosd_remove_ae_ifl_from_snmp_db ae40.0 error 2 Messages appear to be harmless with no functionality impact. [PR1093090](#)
- On MX104 platform, when we configure rate-limit for the logical tunnel (lt-) interface, the commit will fail. As a workaround, we can use firewall filter with policer to achieve the same function. [PR1097078](#)
- When performing the Routing Engine switchover without GRES enabled, due to the fact that the Class-of-Service process (cosd) may fail to delete the traffic control profile state attached to logical interface (IFL) index, the traffic-control-profile may not get programmed after the ifl index is reused by another interface. [PR1099618](#)

Forwarding and Sampling

- The issue is seen while moving an interface from one mesh group to another. [PR1077432](#)
- In rare cases, SSH or telnet traffic might hit incorrect filter related to SCU (Source Class Usage) due to the defect in kernel filter match. This issue comes when the filter has match condition on source class ID. [PR1089382](#)
- In rare cases, MX Series routers might crash while committing inline sampling related configuration for INET6 Family only. [PR1091435](#)

General Routing

- Changing the static route configuration from next-hop to qualified-next-hop might result in static route getting missed from the routing table. Restarting routing process can bring back the routes but with the rpd core. [PR827727](#)
- If with accounting/sampling enabled, an unnecessary update from the routing protocol process (rpd) to the route record database might be triggered by certain configuration change. This process causes jump in CPU utilization of all Packet Forwarding Engines. [PR1002107](#)
- SNMP MIB walk of object "jnxSpSvcSet" gives hardcoded value as "EXT-PKG" for SvcType [PR1017017](#)
- In IP security (IPsec) VPN environment, after performing the Routing Engine switchover, the traffic may fail to be forwarded due to the SAs may not be downloaded to the PIC, or due to some security associations (SAs) on the PIC may incorrectly hold references for old Security Policy Database (SPD) handles while SPD has deleted its entries in the Security Association Database (SAD). [PR1047827](#)

- MPC with Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC (MIC-3D-4COC3-1COC12-CE) might crash. This problem is very difficult to replicate and a preventive fix will be implemented to avoid the crash. [PR1050007](#)
- When "satop-options" is configured on an E1 with Structure-Agnostic TDM over Packet (SAToP) encapsulation, after Automatic Protection Switching (APS) switchover, some SAToP E1s on the previously protect interface (now working) start showing drops. [PR1066100](#)
- When setting the syslog to debug level (any any), you may note reoccurring messages of the form "ifa for this rt ia is not present, consider ifa as ready". These messages are logged for IPv6 enabled interfaces when receiving forwarded packets and cause no harm. Set a higher debug level to avoid seeing them. [PR1067484](#)
- When VMX is deployed, initially there is no management port configuration, so configuration needs to be applied by serial console. The console for VMX is set to 9600 baud rate, with this rate, only a small number of configuration lines can be pasted at a time. [PR1068152](#)
- For Network Address Translation (NAT), Traffic Detection Function (TDF), or IPsec service configured on MX Series platform with MS-MPC/MS-MIC, the received fragmented IPv4/IPv6 packets will be re-assembled and sent out. Under scaled environment, the mspmand process might crash while MS-MPC/MS-MIC is under process of assembling the fragmented packets. [PR1075454](#)
- Traffic throughput test between MPC1/1E/2/2E card and MPC2E/3E NG card, the flowing from MPC1/1E/2/2E card to MPC2E/3E NG card is lesser then from MPC2E/3E NG card to MPC1/1E/2/2E card. [PR1076009](#)
- When a router with AMS infrastructure has MAC flow control enabled, the continuous fragmented packets might crash the NPU and mspmand process (which manages the Multi-Services PIC). [PR1076033](#)
- In subscriber management environment, the PPP daemon (jpppd) might crash repeatedly due to a memory double-free issue. [PR1079511](#)
- On MX Series platform with MS-MPC/MS-MIC, in some mspmand process crash scenarios, after the mspmand coredump is finished or almost finished, PIC kernel also crashes and dumps vmcore. The mspmand cores in these scenario are readable but vmcores are not. [PR1081265](#)
- The rpd process might crash on both master and backup Routing Engines when a routing instance is deleted from configuration if the routing instance is cleaned up before the interface delete is received from device control daemon (dcd). This is a rare timing issue. [PR1083655](#)
- OTN based SNMP Traps such as jnxFruNotifOperStatus and jnxIfOtnNotificationOperStatus are raised by offline/online MIC although no OTN interface is provisioned [PR1084602](#)
- On MX Series routers with MPCs/MICs, if a rlsq interface is receiving continuous fragmented traffic, doing rlsq switchovers couple of times might cause FPC to crash and reboot. [PR1088300](#)

- Wrong ESH checksum computation with non-zero Ethernet Padding in MX Series router. [PR1091396](#)
- In a fib-localization scenario, IPv4 addresses configured on service PICs (SP) will not appear on FIB-remote FPCs although all local (/32) addresses should, regardless of FIB localization role, install on all Packet Forwarding Engines. There is no workaround for this and it implies that traffic destined to this address will need to transit through FIB-local FPC. [PR1092627](#)
- There are entries for PEM in jnxFruEntry in VMX. It is not necessary and is cosmetic. [PR1094888](#)
- Just after the system reboots, rpd process is determining the Routing Engine mastership mode too early before chassisd is determining the mastership, which would cause overload feature not to work properly. [PR1096073](#)
- For Junos OS Release 13.3R1 and later, the DPC card might experience a performance degradation when it's transferring bidirectional short packets (64B) in inline rate. [PR1098357](#)
- On XL-based cards such as MPC5/MPC6, PPE thread timeout errors (resulting in PPE trap files) can be triggered when the FPC allocates illegal memory space for the forwarding state of router operations. - In certain cases, this can result in packet loss depending on how many packets use this forwarding state. [PR1100357](#)
- When the null pointer of jbuf is accessed (jbuf, that is, a message buffer is allocated only when the packet is ready to process. The buffer is freed after the packet completes ALG handling is accessed), for example, when using the Microsoft Remote Procedure Call (MS RPC) (as observed, issue may also happen on Sun Microsystems RPC) Application-level gateway (ALG) with NAT (stateful firewall is used as a part of the service chain), if the traffic matching configured universal unique identifier (UUID) is arrived on the ALG, the mspmand (which manages the Multiservice PIC) crash occurs. [PR1100821](#)
- After Junos OS Release 13.3R1, IPCMON infra is added to debug IPCs between PFEMAN and the Routing Engine. When convergence occurs, string processing of IPCMOM will take added time. Then the slow convergence will be seen. It is a performance issue, it is visible in scaled scenario (for example, more than 100K routes). As a workaround, please execute command "set pfe ipclog filter clear" to disable IPC logging on all FPCs. [PR1100851](#)
- FFP is a generic process that should be called during commit process, and FFP calls the PDB initialization as part of its process. On the PDB-unsupported platforms (MX Series, M10i, M120, M320 are PDB-supported), when committing configuration, some error messages will be seen. [PR1103035](#)
- If fpc offline knob is configured after the presence of Non-recoverable faults, then offline action will not be performed. [PR1103185](#)
- Non-queuing MPC5E might crash continuously if rate-limit under transmit-rate for scheduler is applied. As a workaround, do not configure rate-limit and use firewall policer for forwarding-class instead. MPC5EQ is not exposed. [PR1104495](#)

- An IPv4 filter configured to use the filter block with term that has both "from precedence" and another non 5-tuple (i.e. not port, protocol, address) will cause an XL/EA based board to reboot. Example: set firewall family inet filter FILTER fast-filter-lookup set firewall family inet filter FILTER term TERM from precedence PRECEDENCE set firewall family inet filter FILTER term TERM from tcp-established [PR1112047](#)
- In the scenario that the power get removed from the MS-MPC, but Routing Engine is still online (for example, on MX960 platform with high capacity power supplies which split into two separate power zones, when the power zone for the MS-MPC line card loses power by switch off the PEM that supports the MS-MPC situated slot), if the power goes back (for example, switch on the PEM), the MS-MPC might be seen as "Unresponsive" (checked via CLI command "show chassis fpc") and not coming up back online due to failure of reading memory. [PR1112716](#)
- Under certain conditions, when the Junos OS Routing Engine tries to send an IP packet over a IPIP tunnel, the lookup might end up in an infinite loop between two IPIP tunnels. This is caused by a routing loop causing the tunnel destination for Tunnel#A to be learned through Tunnel#B and the other way around. [PR1112724](#)
- Under certain conditions, when the Junos OS Routing Engine tries to send an IP packet over a GRE tunnel, the lookup might end up in an infinite loop between two GRE tunnels. This is caused by a routing loop causing the tunnel destination for Tunnel#A to be learned through Tunnel#B and the other way round. [PR1113754](#)

Infrastructure

- On dual Routing Engine platform, if GRES is configured (triggered by "on-disk-failure"), when a disk I/O failure occurs on the master Routing Engine due to hardware issue (for example, SSD failure), the graceful Routing Engine switchover might not be triggered immediately after initial IO failure has been detected. As a result, Routing Engine might enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. [PR1102978](#)

Interfaces and Chassis

- After changing the speed of fxp0 interface (the management Ethernet interface) to 1G (the maximum speed), the interface process (dcd) configures the interface but reads the speed even before the change takes effect. Although the hardware speed is updated to 1G, from dcd perspective, the speed is still not changed. Then if you change back to the original speed, the change is ignored by dcd. [PR976825](#)
- On MX Series platform, when an aggregated Ethernet bundle participating as L2 interface within bridge-domain goes down, the following syslog messages could be observed. The messages would be associated with FPC0 even if there are no link(s) from this FPC0 participating in the affected aggregate-ethernet bundle. mib2d[2782]: SNMP_TRAP_LINK_DOWN: ifIndex 636, ifAdminStatus up(1), ifOperStatus down(2), ifName xe-3/3/2 mib2d[2782]: SNMP_TRAP_LINK_DOWN: ifIndex 637, ifAdminStatus up(1), ifOperStatus down(2), ifName xe-3/3/3 mib2d[2782]: SNMP_TRAP_LINK_DOWN: ifIndex 740, ifAdminStatus up(1), ifOperStatus down(2), ifName ae102 fpc0 LUCHIP(0) Congestion Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f:f:f:f:f fpc0 LUCHIP(0) Congestion Detected, Active Zones 2:0:0:0:0:8:a:0:0:0:0:0:8:4:0:a alarmd[1600]: Alarm set: FPC

color=RED, class=CHASSIS, reason=FPC 0 Major Errors craftd[1601]: Major alarm set, FPC 0 Major Errors fpc0 LUCHIP(0) Congestion Detected, Active Zones 2:0:0:0:8:a:0:0:0:0:8:4:0:a alarmd[1600]: Alarm cleared: FPC color=RED, class=CHASSIS, reason=FPC 0 Major Errors craftd[1601]: Major alarm cleared, FPC 0 Major Errors fpc0 LUCHIP(0): Secondary PPE 0 zone 1 timeout. fpc0 PPE Sync XTXN Err Trap: Count 7095, PC 10, 0x0010: trap_nexthop_return fpc0 PPE Thread Timeout Trap: Count 226, PC 34a, 0x034a: nh_ret_last fpc0 PPE PPE Stack Err Trap: Count 15, PC 366, 0x0366: add_default_layer1_overhead fpc0 PPE PPE HW Fault Trap: Count 10, PC 3c9, 0x03c9: bm_label_save_label fpc0 LUCHIP(0) RMC 0 Uninitialized EDMEM[0x3f38b5] Read (0x6db6db6d6db6db6d) fpc0 LUCHIP(0) RMC 1 Uninitialized EDMEM[0x394cdf] Read (0x6db6db6d6db6db6d) fpc0 LUCHIP(0) RMC 2 Uninitialized EDMEM[0x3d9565] Read (0x6db6db6d6db6db6d) fpc0 LUCHIP(0) RMC 3 Uninitialized EDMEM[0x3d81b6] Read (0x6db6db6d6db6db6d) These message would be transient in nature. The discrepancy of nexthop handling that is addressed in this PR can also manifest itself in form of other issues in the system. Basically when the nexthops go out of sync we are bound to see either Packet Forwarding Engine crashes/traps or Routing Engine crashes. The fix in this PR should take care of this behavior and ensure we handle the nexthops correctly to maintain the synchronization between master Routing Engine, backup Routing Engine and all Packet Forwarding Engine peers. [PR990023](#)

- dcd will crash if targeted-distribution applied to ge ifd via dynamic-profile [PR1054145](#)
- Two redundant logical tunnels (rlt) interfaces are configured with "per-unit-mac-disable" enabled. After configure the second one, the first rlt interface goes down. rlt0 { logical-tunnel-options { per-unit-mac-disable; <<<<<< } } [PR1055005](#)
- It is observed that the syslog messages related to kernel and Packet Forwarding Engine may get generated at an excessive rate, especially in subscriber management environment. Most of these messages may appear repeatedly, for example, more than 1.5 million messages may get recorded in 2 hours, and there are only 140 unique messages. Besides, these messages are worthless during normal operation and due to the excessive rate of log generation, high RE CPU consumption (for example, RE CPU utilization can be stuck at 100% for a long time (minutes or hours), it depends on the activity of subscribers (frequency of logins and logouts) and on the AI scripts used by the customer) by event process (eventd) might be observed on the device. [PR1056680](#)
- For Junos release 13.3R1 or later, after multiple (e.g. 26) iterations of graceful Routing Engine switchover (GRES), the TNP address of management interface might be deleted wrongly during switchover, this leads to all FPCs be offline. [PR1060764](#)
- On MX Series routers, INET MTU (PPP payload MTU, that is IP header plus data excluding any L2 overhead) is being set to lowest MRU of either MX (local device) or peer. This behavior is not inline with ERX behavior, which is set to min(local MTU, peer MRU). This might cause the packet drops in the customer network in the downstream path. [PR1061155](#)
- When the Ethernet Link Fault Management (LFM) action profile is configured, if there are some errors (refer to the configuration, for example, frame errors or symbol errors) happening in the past (even a long past), due to the improper handling of error stats

fetching from kernel, the LFM process (lfmd) may generate false event PDUs and send the false alarm to the peer device. [PR107778](#)

- On MX Series Virtual Chassis (MX-VC) platform with "subscriber-management" enabled, after power up/reboot, the VC backup router (VC-B) experiences a rapid sequence of role transitions from no-role to VC master router (VC-M) to VC-B, the expected local GRES and a reboot of the former master Routing Engine might not happen on the VC-B, some of the FPCs on it might be stuck in "present" state and eventually rebooted. [PR1086316](#)
- When an interface on SFPP module in MIC is set disabled, after pulling out the SFPP and then insert it, the remote direct connected interface might get up unexpectedly. [PR1090285](#)
- After removing a child link from AE bundle, in the output of "show interface AE detail", the packets count on the remaining child link spikes, then if add back the previous child link, the count recover to normal. [PR1091425](#)
- In MX Series Virtual Chassis (MXVC) environment, when rebooting the system or the line cards which contain all the Virtual Chassis port (VCP) links, because line cards may fail to complete the rebooting process within 5 minutes, the timer (that is, the amount of time allowed for the LCC to connect to the SCC) started by the master router may expire which may cause the VCP links establishment failure. In addition, this issue is not specific to the line cards type, based on the observation, the timer (5min) may expire on a MX2020 with all 20 FPCs equipped as well. [PR1095563](#)
- On PB-2OC12-ATM2-SMIR PIC, port 0 and port 1 are configured with clock source as external, if Loss of signal (LOS) is inserted on port 0, the port 0 will down, the expect behavior is clock being used from port 1. But in this case, port 0 down will results in port 1 flapping and reporting SONET phase lock loop (PLL) errors. [PR1098540](#)
- Due to the fact that the error injection rate configured by user on Routing Engine via CLI command "bert-error-rate" may not be programmed in the hardware register, the PE-4CHOC3-CE-SFP, PB-4CHOC3-CE-SFP, MIC-3D-4COC3-1COC12-CE, and MIC-4COC3-1COC12-CE-H may fail to inject bit errors during a Bit Error Ratio Test (BERT). [PR1102630](#)
- On MPC-3D-16XGE-SFPP line card, when an optics (for example, 10G-LR-SFP) is disabled and then enabled administratively, if the SFP is not temperature tolerant (non-NEBS compliant), the TX laser may not be turned on due to the fact that the chassis process (chassisd) may keep sending the "disable-non-nebs-optics" command to the optics if the current temperature of FPC reaches the threshold temperature. [PR1107242](#)
- On MX Series platform, continuous error messages might be seen on the MICs (for 10G/40G/100G MICs) from MIC3 onwards (listed as below) when physical interface (IFD) settings are pushed (e.g. booting the MPC). Based on the current observation, the issue may not have any operational impact and the MICs that may encounter this issue are listed as below, - 10G MICs: MIC3-3D-10XGE-SFPP, MIC6-10G, MIC6-10G-OTN, - 40G MICs: MIC3-3D-2X40GE-QSFPP, - 100G MICs: MIC3-3D-1X100GE-CFP, MIC3-3D-1X100GE-CXP, MIC6-100G-CXP, MIC6-100G-CFP2 [PR1108769](#)

Layer 2 Features

- During interface flaps a high amount of TCN (Topology Change Notification) might get propagated causing other switches to get behind due to high amount of TCN flooding. This problem is visible after the change done from 11.4R8 onwards which propagates TCN BPDU immediate and not in the pace of the 2 second BPDU Hello interval to speed up topology change propagation. The root cause is the TCNWHILE timer of 4 seconds is always reset upon receiving TCN notifications causing the high churn TCN propagation. [PR1089580](#)
- In MX Series Virtual Chassis (MXVC) environment, when packets come from an interface (for example, xe-16/0/1.542) situated on one member of VC (for example, VC member 1), if the ingress Packet Forwarding Engine (for example, FPC16 PFE0, who runs hash to determine which interface it should send the packet to) decides that it should send the packet via another interface (for example, xe-4/0/1.670) situated on a different member (for example, VC member 0), it will send the frame to member 0 via the vcp-intf. In case of xe-4/0/1.670 belongs to an AE bundle which has multiple child links, a hash needs to be run on Packet Forwarding Engine carrying the VCP port (receiving side on member 0) to determine which one is the egress Packet Forwarding Engine within member 0 to send the packet out after vcp-intf gets the packet. This hash result should get the same result as the ingress Packet Forwarding Engine. If it is not the case, then the packet would get dropped on Packet Forwarding Engine on member 0. [PR1097973](#)
- With scaled subscribers connected, restarting one of MPCs might cause subscribers unable to log in for about 2 mins. [PR1099237](#)
- In a scenario that BGP based VPLS stitching with L2circuit, with "pseudowire-status-tlv" configured under L2circuit's mesh-group, if L2circuit neighbor doesn't configure "pseudowire-status-tlv", then status of "Negotiated PW status TLV" of VPLS connection is "NO", this will cause BGP based VPLS connection can not up even the L2circuit is up. [PR1108208](#)

Layer 2 Ethernet Services

MPLS

- In Resource Reservation Protocol (RSVP) environment, if CoS-Based Forwarding (CBF) for per LSP (that filter out traffic not related to that LSP) is configured, and either the feature fast-reroute or link-protection is used on the device, when the primary link is down (for example, turning off the laser of the link), due to some next hops of the traffic may be deleted or reassigned to different class of traffic, and the RSVP local repair may fail to process more than 200 LSPs at one time, the traffic may get dropped by the filter on the device before the new next hop is installed. In this situation, the feature (fast reroute or link protection) may take longer time (for example, 1.5 seconds) to function and the traffic loss might be seen at the meantime. In addition, the issue may not be seen if the CBF for per LSP is not configured on the device. [PR1048109](#)
- Junk characters are being displayed in output of show connections extensive command. [PR1081678](#)

- On dual Routing Engine platform with GRES, the kernel synchronization process (ksyncd) may crash on the backup Routing Engine when adding of route pointing to indirect nexthop on system. [PR1102724](#)
- In Junos OS release 13.2R1 and above, in MPLS L3VPN scenario, when "l3vpn-composite-nexthop" statement is enabled on a PE router and an interface style service set is attached to the ingress interface, the L3VPN packets with the MPLS labels will be sent to the service card and dropped. As a workaround, we should disable "l3vpn-composite-nexthop". [PR1109948](#)

Network Management and Monitoring

- Mib2d cores while trying to re-add a lag child into the internal DB. Since the entry is already present in the internal DB. Before adding the child link mib2d does a lookup on the tree, to know if the entry is not already there. However, this lookup returns no results, since the child link is part of snmp filter-interface configuration. [PR1039508](#)

Platform and Infrastructure

- LSI logical interface input packet and byte stats are also added to core logical interface stats, but when the LSI logical interface goes down and the core logical interface stats are polled, there is a dip in stats. The fix is to restore LSI logical interface stats to core logical interface before deleting the LSI logical interface. [PR1020175](#)
- The Priority code point (PCP) and Drop eligible indicator (DEI) bit in 802.1Q header are preserved while packet gets routed within the same Packet Forwarding Engine . The expected behavior is resetting the PCP and DEI bit when the packet is routed. [PR1036756](#)
- Due to a defect in the Junos OS software, when a telnet user experiences some undefined network disconnect, .perm and .env files under /var/run are left behind. This scenario happens only under certain unknown ungraceful network disconnects. When considerable number of .perm/.env files get accumulated under /var/run, issue is seen with telnet users, that they are not able to perform permitted operations on the router, post-login. [PR1047609](#)
- If a Radius server is configured as accounting server, when it is non-reachable, the auditd process might stressed with huge number of audit logs to be sent to the accounting server, which might cause auditd to crash. [PR1062016](#)
- VRRP advertisements might be dropped after enable delegate-processing on the logical tunnel (lt) interface. It would result in VRRP master state observed on both routers. [PR1073090](#)
- Problem: It tries to check allotted power for all the FPCs, here in the CHASSISD_I2CS_READBACK_ERROR logs it shows for the FPCs which are not present in chassis. It just calls i2cs_readback() to read i2c device and fails there as these FPCs? slots are blank and prints those readback errors. Also the errors are harmless: "CHASSISD_I2CS_READBACK_ERROR: Readback error from I2C slave for FPC" Fix: Code to check 'if power has been allotted to this FPC', needs to be executed only if the FPC is present. [PR1075643](#)
- When a MX Series chassis network-services is "enhanced-ip" and an AE is part of a Layer 2 bridge (bridge-domain or VPLS), there is a possibility that an incorrect

forwarding path may be installed causing traffic loss. This could happen when first applying the configuration, restarting the system or restarting the line card. [PR1081999](#)

- If with both MPC/MSDPC and other type of DPCs equipped, for local switching at mesh group level, split horizon on PW interfaces won't work and this would cause packets to loop back to same PW interface. [PR1084130](#)
- In Junos OS Releases 13.3R3, 14.1R1, 14.2R1, there is a new feature, an extra TLV term is added to accommodate the default action for the "next-interface" when the corresponding next-interface is down. While doing a unified ISSU from an image without the feature to an image with this feature, all MPCs might crash. [PR1085357](#)
- Issue is specific to 64-Bit RPD and config-groups wildcard config specific as in below case: set groups TEST routing-instances <*> routing-options multicast forwarding-cache family inet threshold suppress 200 set routing-instances vrf1 apply-groups TEST set routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 600 With this daemon(rpd) reads suppressed value ?200? (i.e. coming from groups) instead of reading value ?600? from foreground and customer sees unexpected behavior with respect to threshold-suppress. Workaround: They can replace wildcard with actual routing-instance name as in below example: set groups TEST routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 200 set routing-instances vrf1 apply-groups TEST set routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 600 [PR1089994](#)
- On MX series router, if ifl (logical interface) is configured with VID of 0 and parent ifd (physical interface) with native-vlan-id of 0, when sending L2 traffic received on the ifl to Routing Engine, the VID 0 will not imposed, causing the frames to get dropped at Routing Engine. [PR1090718](#)
- When an interface on MQ-based FPC is going to link down state, in-flight packet on interface transmit path will be stuck on the interface and never drained until the interface comes up again. As a result, small number of such stacked packets will be sent out when the interface is going to UP state. No other major impact should not be seen after those packets are drained. [PR1093569](#)
- On MX2020/2010 router, an SPMB core file will be seen if there are bad XF chips (fabric chip) on SFB, which might trigger RE/CB switchover. [PR1096455](#)
- When a P2MP LSP is added or deleted at ingress LSR, traffic loss is seen to existing sub-LSP(s) at transit LSR which replicates and forwards packet to egress PEs. This issue only affects MX Series routers with MPCs/MICs. [PR1097806](#)
- The "shared-bandwidth-policer" knob is used to enable configuration of interface-specific policers applied on an aggregated Ethernet bundle to match the effective bandwidth and burst-size to user-configured values. But this feature is broken from Junos release 14.1R1 when "enhanced-ip" is configured on MX Series platform with pure trio-based line cards. The bandwidth/burst-size of policers attached to Aggregated Ethernet interfaces are not dynamically updated upon member link adding or deletion. [PR1098486](#)

- On Trio-based platform, when the type of the IPv6 traffic is non-TCP or non-UDP (for example, next header field is GRE or No Next Header for IPv6), if the traffic rate is high (for instance, higher than 3.5Mpps), the packet re-ordering may occur. [PR1098776](#)
- On MX Series-based line cards, when the prefix-length is modified from higher value to lower value for an existing prefix-action, heap gets corrupted. Due to this corruption, the FPC might crash anytime when further configurations are added/deleted. The following operations might be considered as a workaround: Step 1. Delete the existing prefix-action and commit Step 2. Then re-create the prefix-action with newer prefix-length [PR1098870](#)
- In an MPLS L3VPN network with a dual-homed CE router connected to different PE routers, a protection path should be configured between the CE router and an alternate PE router to protect the best path. When BFD is enabled on the BGP session between the CE and the primary PE router, with local traffic flowing from another CE connected with the primary PE to this CE, after bring down the interface on the best path, the local repair will be triggered by BFD session down but it might fail due to a timing issue. This will cause slow converge and unexpected traffic drop. [PR1098961](#)
- On Trio-based platform, before creating a new unicast nexthop, there is a check to see if there is at least 512k DoubleWords (DW) free. So, even the attempting NH requires only a small amount of memory (for example, < 100 DWs), if there is not enough free DWs (that is, 512k), the check will fail and the end result is that the control plane will quit adding this NH prematurely - stopping at ~80% of capacity. With the fix, it will check for 64k free DWs which is lower reference watermark for available resource, thereby ensuring that can allocate resource. [PR1099753](#)
- From Junos release 14.1 and above, IPv6 mobility packets with Heartbeat option that the length of the mobility header (including the ethernet encapsulation and main IPv6 header) extends beyond 128 Bytes will be discarded as bad IPv6 option packet due to a logic error in packet handling. [PR1100442](#)
- Large scaled inline BFD session (in this case, 6000 inline BFD sessions) are loaded with the minimum-interval value 50ms. If FPC restarts, some BFD sessions might flap. [PR1102116](#)
- On MPC3E/MPC4E line card, when the feature "flow-detection" is enabled (under "ddos-protection" hierarchy), if suspicious control flow is received, two issues may occur on the device: Issue 1: sometimes, the suspicious control flow may not get detected on the line cards Issue 2: once the suspicious control flows are detected, they may never time out even if the corresponding packets stop [PR1102997](#)
- Customer is having the similar issue [PR1103771](#)
- On T4000 platform with FPC Type-5 equipped, after performing unified ISSU, due to the fact that only 6 out of 16 temperature sensors may get initialized, the temperature reading for the line card may be shown as "Absent". [PR1104240](#)
- Any configuration or logical interface (IFL) change will introduce 160 bytes memory leak on MPC heap memory when we have any type of inline sampling configured (ipfix or version 9). Only trigger of issue is the configuration of inline sampling, even without traffic being sampled. The leak is more evident in a subscriber management scenario

when we have many IFL addition/deletion. Rebooting MPC in a controlled maintenance window is the only way to restore memory. [PR1105644](#)

- When a common scheduler is shared by multiple scheduler maps which applies to different VLANs of an Aggregated Ethernet (AE) interface, if the knob "member-link-scheduler" is configured as "scale", for some VLANs, the scheduler parameters are wrongly scaled among AE member links. As a workaround, we should explicitly configure different schedulers under the scheduler maps. [PR1107013](#)
- Due to a software defect found in 13.3R7.3 and 14.1R5.4 inclusively, Juniper Networks strongly discourage the use of Junos software version 13.3R7.3 on routers with MQ-based MPC. This includes MX-Series with MPC1, MPC2; all mid-range MX-Series; and some of EX9200 line cards. [PR1108826](#)
- DHCP End options (option 255) is missing by DHCP-relay agent (where 20 bytes DHCP options 82 inserted) for client DHCP discover message with 19bytes padding. [PR1110939](#)
- On Trio-base FPC, when MPLS-labeled fragmented IPv6 packets arriving at PE router (usually seen in 6PE and 6VPE scenario), the Packet Forwarding Engine might mistakenly detect such IPv6 header and then drop these packets as "L3 incompletes" in the output of "show interface extensive". [PR1117064](#)

Routing Policy and Firewall Filters

- In Class-of-Service (CoS) environment, there is a possibility (happened twice so far and not reproducible in the lab) that routing protocol process (rpd) may crash because the CoS memory may get incorrectly freed and then allocated again. [PR1062616](#)
- On the platform that M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120, and MX with DPC, when the flood filter is configured in VPLS instance on the Packet Forwarding Engine, if the Packet Forwarding Engine receives a filter change (for example, FPC reboot occur and comes up), the line card may fail to program the filter. [PR1099257](#)

Routing Protocols

- In mutli-topologies IS-IS scenario, there is huge difference between estimated free bytes and actual free bytes when generating LSP with IPv6 Prefix. It might cause LSP fragment exhaustion. [PR1074891](#)
- There are two issues in the PR: (1) In multicast environment, Incoming interface list (IIF) list has only RPF interface, designated forwarder (DF) winners are not added in the list in backup Routing Engine. (2) "Number of downstream interfaces" in show pim join extensive is not accounting Pseudo-VXLAN interface. [PR1082362](#)
- When removing BGP Prefix-Independent Convergence (PIC) from the configuration, the expect behavior is that any protected path would become unprotected. But in this case, the multipath entry that contains the protection path (which is supposed to be removed) remains active, until BGP session flaps or the route itself flaps. As a workaround, we can use "commit full" command to correct or to commit. [PR1092049](#)
- In BGP environment, when configuring RIB copy of routes from primary routing table to secondary routing table (for example, by using the CLI command "import-rib [inet.0 XX.inet.0]") and if the second route-table's instance is type "forwarding", due to the BGP routes in secondary routing table may get deleted and not correctly re-created,

the routes may be gone on every commit (even commit of unrelated changes). As a workaround, for re-creating the BGP routes in secondary route table, use CLI command "commit full" to make configuration changes. [PR1093317](#)

- In Junos OS Release 9.1 and later, RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. In this case, when AS4_AGGREGATOR attribute (18) is received from a 2-byte AS peer (note AS4_AGGREGATOR attribute is only received when the aggregator has 4-byte AS but this peer only supports 2-byte AS), NSR synchronization with standby Routing Engine would fail, causing session constantly bouncing on standby Routing Engine (hogging CPU). [PR1093615](#)
- With this change the default label hold timer was increased for 10 seconds to 60 seconds. [PR1093638](#)
- The rpd process might crash when resolve-vpn and rib inet.3 are configured under separate levels (BGP global, group and peer). The fix is If anybody configures a family at a lower level, reset the state created by either of knobs from higher levels. This behavior conforms with our current behavior of family config - which is that any config at a lower level is honored and the higher level config is reset. [PR1094499](#)
- When BGP routes has multiple protocol nexthops including discard/reject and other IGP nexthops, the discard/reject nexthop will be selected as BGP nexthop, which will cause traffic loss. [PR1096363](#)
- When polling SNMP OID isisPacketCounterTable 1.3.6.1.2.1.138.1.5.3, the rpd process might crash. [PR1101080](#)
- When the ISIS configurations has been removed, the ISIS LSDB contents got flushed. If at the same time of this deletions process, there is an SPF executions (that is, try to access the data structures at same time when/a fraction of secs after freeing its content), routing protocol process (rpd) crash occurs. [PR1103631](#)

Services Applications

- 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](#)
- In IPsec environment, after performing the Routing Engine switchover (for example, performing Graceful Routing Engine Switchover) or chassis reboot (that is, whole device is powered down and powered UP again), due to the key management daemon (kmd) may be launched before the Routing Engine mastership is finalized, it may stop running on the new master Routing Engine. [PR863413](#)
- In the L2TP scenario with dual Routing Engines. After subscriber management infrastructure daemon (smid) being restarted, because the delete notification to backup RE might be lost, the subscriber database (SDB) information does not synchronize between master RE and standby RE. After RE switchover is executed, the Layer 2 Tunneling Protocol daemon (jl2tpd) might crash, and new L2TP subscribers are unable to dial. [PR968947](#)

- On an L2TP access concentrator (LAC) device with more than 8K L2TP sessions up, if execute command "clear services l2tp session all" and then stop the command by using ctrl-C, the Layer 2 Tunneling Protocol process (jl2tpd) might crash. [PR1009679](#)
- On MX-series router that configured as L2TP tunnel switch (LTS), after receiving a Call-Disconnect-Notify (CDN) message on LNS interface from remote LNS, the L2TP daemon (l2tpd) might crash and dump a core file. [PR1021881](#)
- On Trivial File Transfer Protocol (TFTP) Application Layer Gateway (ALG) with NAT translation type "dynamic-nat44" configured, MS-DPC/MS-MPC/MS-MIC might crash when processes the TFTP packets. [PR1091179](#)
- On M series platform, in Layer 2 Tunneling Protocol (L2TP) network server (LNS) environment, not all attributes (Missing NAS-Identifier, NAS-Port-Type, Service-Type, Framed-Protocol attributes) within Accounting-Request packet are sending to the RADIUS server. [PR1095315](#)
- If MS-DPC is used in CG-NAT environment, in a very rare condition, when the MS-DPC tries to delete a NAT mapping entry (e.g. entry timeout), error might occur and the MS-DPC might get rebooted and then dump a core file. [PR1095396](#)
- Some values of MIB object jnxSrcNatStatsEntry might be doubled when AMS (or rsp) interface and NAT are configured together. [PR1095713](#)

Software Installation and Upgrade

- Add "on <host>" argument to "request system software validate" to allow validation on a remote host/RE running Junos. [PR1066150](#)

Subscriber Access Management

- In scaled DHCP subscribers environment, the authd process might crash and generate a core file after clearing DHCP binding or logout subscribers. [PR1094674](#)

VPNs

- In BGP MVPN scenario, an MSDP timeout on the PE might occur causing the source to be removed even if it is local. This will cause type-5 flaps and traffic loss of 30 to 40 seconds. The issue showed up in a scaling MSDP configuration where the KA timer periodically expires with a CE (not PE) acting as RP. The fix has been provided to add a check for local source (even if not local RP) before withdrawing the type-5 route. [PR1011124](#)
- In Internet multicast over an MPLS network by using next-generation Layer 3 VPN multicast (NG-MVPN) environment, when rib-groups are configured to use inet.2 as RPF rib for Global Table Multicast (GTM, internet multicast) instance, the ingress PE may fail to add P-tunnel as downstream even after receiving BGP type-7 routes. In addition, this issue only affects GTM. [PR1104676](#)

Resolved Issues: 14.1R5

- [Class of Service \(CoS\) on page 122](#)
- [Forwarding and Sampling on page 123](#)

- [General Routing on page 124](#)
- [Infrastructure on page 128](#)
- [Interfaces and Chassis on page 129](#)
- [Layer 2 Features on page 132](#)
- [MPLS on page 133](#)
- [Network Management and Monitoring on page 134](#)
- [Platform and Infrastructure on page 135](#)
- [Routing Protocols on page 139](#)
- [Routing Policy and Firewall Filters on page 142](#)
- [Services Applications on page 142](#)
- [Software Installation and Upgrade on page 143](#)
- [Subscriber Access Management on page 143](#)
- [VPNs on page 143](#)

Class of Service (CoS)

- This error message "only per-unit and 2-level hierarchical scheduler are supported on this interface" is a cosmetic regression issue without any functional impact. [PR1050512](#)
- Forwarding class accounting stops working after Routing Engine switchover. This behavior has been corrected in Releases 13.3X2, 13.3R7, 14.1R5, 14.2R3, 13.3R6, and 15.1. Issue comes when MPC reboots for any reason with forwarding-class-accounting configured on AE/AS interface. In forwarding-class-accounting feature, counters are allocated based on number of forwarding classes configured in MPC. In error case on MPC reboot, AE interface is getting created before the message for configuring number of forwarding classes in MPC comes. As a result while enabling forwarding-class-accounting feature on AE interface, number of forwarding classes value in MPC is 0 and counters are not allocated causing issue. Cause: Race condition when on MPC reboot AE interface getting created before number of forwarding classes are configured. Fix: When number of forwarding classes are set after MPC reboot, check for any AE interface with forwarding-class-accounting configured and reprogram it. [PR1060637](#)
- Add chassis scheduler map support on gr interface on MS-PIC, which means there will be no commit error if scheduler-map-chassis is applied on gr interface. [PR1066735](#)
- 1. With "hierarchical-scheduler" configured at IFD level 2. Under class-of-service hierarchy "output traffic control profile" configured at "interface-set" as well as IFD level, for the same IFD/IFL. With the above two conditions met, when a Junos OS upgrade is performed on a dual Routing Engine system the configuration validation check would fail on the Routing Engine that is upgraded later with this error message. Error message: "cannot configure a traffic control profile for this ifl when a parent has a traffic control profile that references a scheduler map: ifl xe-11/0/0.5000 refers to traffic-control-profile TCP_PE-CE_30M. It is also a member of interface set xe-11/0/0_OTag=80 which has traffic-control-profile TCP_PE-CE_80M which references scheduler-map SM_PE-CE" condition-1: lab-re1> show configuration interfaces xe-11/0/0 { hierarchical-scheduler; <<< Condition-2: lab-re1> show configuration interfaces interface-set xe-11/0/0_OTag=80 { interface xe-11/0/0 { <...>; } } lab-re1> show

```
configuration class-of-service interfaces interface-set xe-11/0/0_OTag=80 {
  output-traffic-control-profile TCP_PE-CE_80M; <<< } <..> xe-11/0/0 {
  output-traffic-control-profile TCP_Maxbuff; unit 5000 { output-traffic-control-profile
  TCP_PE-CE_30M <<< } } PR1069477
```

- On MX Series platform configured for IP network-services (default) and with MS-DPC/Tunnel-Interface, virtual-tunnel (vt) interfaces are created automatically to support ultimate-hop-popping upon enabling "protocol rsvp". These interfaces are associated with default IP and MPLS classifiers along with MPLS re-write rule. When "protocol rsvp" is disabled/enabled or MS-DPC/FPC (with tunnel-service) restarts, the vt interfaces are deleted and re-added to the system. However during the deletion, these interfaces are not getting released from cosd process and thus leads to memory leak in cosd. [PR1071349](#)
- On MX Series platform, when aggregated Ethernet (AE) interface is in link aggregation group (LAG) Enhanced mode, after deactivating and then activating one child link of the LAG, the feature that runs on AE interface rather than on the child link (for example, IEEE-802.1ad rewrite rule) may fail to be executed. [PR1080448](#)

Forwarding and Sampling

- On MX Series routers with MPCs/MICs, when deleting firewall filter and the routing instance it is attached to, in some race conditions, the filter might not be deleted and remains in resolved state indefinitely. [PR937258](#)
- When a firewall filter, which is used to de-encapsulate the IPv4 packets encapsulated in IPv6 GRE header, is attached to interface hosts on MX Series MPC/MIC, the IPv6 GRE header would be de-encapsulated but the inner IPv4 packet would end up getting dropped and not forwarded. This issue affects the packet with IPv4 over IPv6 GRE header only, and those packets with IPv6 over IPv6 GRE header are not affected. [PR1054039](#)
- If the template of the policer is changed (for example, change the bandwidth-limit value of policer), shared-bandwidth-policer knob may not function properly anymore. [PR1056098](#)
- In some rare cases, SNMP might get Output bytes of Local statistics instead of the Traffic statistics when retrieving Output bytes of Traffic statistics on a logical interface. [PR1083246](#)

General Routing

- On dual Routing Engine platforms, after performing unified graceful Routing Engine switchover (GRES) with 8K subscribers, the ksyncd process may crash due to the replication error on a next hop change operation. The issue is hit when there's memory pressure condition on the Routing Engine and in that case, it may lead to null pointer de-reference and ksyncd crash. Or in some case, the kernel on the new master Routing Engine might crash after Routing Engine switchover if Routing Engine is under memory pressure due to missing null check when trying to add a next hop and the next hop is not found at the time. [PR942524](#)
- When the mirrored interface and mirror destination interface are hosted on different Virtual Chassis (VC) members, the ingress MPLS packets are not getting mirrored to the mirror destination. [PR979888](#)
- In point-to-point (P2P) SONET/SDH interface environment, there is a destination route with this interface as next-hop. When this interface is disabled, the destination route is still kept in the forwarding table and might cause ping fails with "Can't assign requested address" error. [PR984623](#)
- The knob 'gratuitous-arp-on-ifup' should send a gratuitous arp on each unit of a physical interface, but in Release 12.3 and later versions, only the first unit is seeing the configured behavior. [PR986262](#)
- Optics lane#3 and lane#4 TX, RX power alarm data was ignored but the lane#1 and lane#2 data was used for lane#3 and lane#4 respectively. Causing incorrect/false alarm on lane#3 and lane#4 [PR1001670](#)
- When there are no services configured, datapath-traced daemon is not running. In the PIC the plugin continues to try for the connection, and continuous connection failure logs are seen. [PR1003714](#)
- A raw IP packet with invalid Memory Buffer(mbuf) length may trigger a kernel crash. The invalid mbuf length might be set incorrectly by other daemons. [PR1006320](#)
- During Wan Link flaps, ASIC streams in the Packet Forwarding Engines are disabled/enabled on the fly when traffic is in flight. This is normal and will result in the Cell drops, PKTR ICELL signature errors and SLOUT errors. However under certain rare conditions, Lout IP -Pkt Len Mismatch error is observed which sometimes trigger automatic restart of the FPC. On TXP, TXP-3D in FPC Type 4-ES can experience automatic restart during wan interface flaps. [PR1013522](#)
- If you issue the "show services nat mappings details" command with a large number of service sets configured (such as 1000 service sets) and one or two NAT mappings specified, the command takes a certain amount of time to display the output. During this period, if you deactivate or activate the services, a multiservices PIC management daemon core file is generated. [PR1019996](#)
- Total CPU Utilization and Interrupt CPU Utilization are displayed incorrectly for MPC3E-3D-NG-Q, MPC3E-3D-NG, MPC2E-3D-NG-Q, and MPC2E-3D-NG. This is because the router incorrectly calculates CPU utilization from system startup rather than the CPU utilization at a particular point of time. [PR1024150](#)

- On MPC5E line card, if a firewall filter with large-scale terms (more than 1300) is attached to an interface, traffic drop might be seen. [PR1027516](#)
- In the scenario where router acts as both egress LSP for core network and BRAS for subscribers, RSVP-TE sends PathErr to ingress router due to matching to subscriber interfaces wrongly when checking the explicit route object (ERO), if subscribers are associated with same lo0 address as used by RSVP LSP egress address. [PR1031513](#)
- On the Type 5 PIC, when the "hold-time down" of the interface is configured less than 2 seconds and the loss of signal (LOS) is set and cleared repeatedly in a short period (for example, performing ring path switchover within 50ms), the "hold-time down" may fail to keep the interface in "up" state within the configured time period. [PR1032272](#)
- Problem scenario: Issue is seen only with vMX. It will be seen when the PPPoE session's Keep-alive timer expiry happens. This might be due to non-graceful termination of remote side or due to communication path failure with the remote end. Problem statement: When PPPoE session Keep-alive timer expires the local PPPoE session is NOT closed/logout. [PR1034520](#)
- When the CPU usage is very high (e.g. 100%) on Routing Engine, the MS-MIC might get stuck due to kernel deadlock, which triggers the card to crash and generate a core file. [PR1038026](#)
- If default-address-selection knob is configured on MX-VC, VC-heartbeat connection between member chassis may be unable to come up. [PR1041194](#)
- CLNS ping fails for l3vpn over ethernet scenario. IS-IS routes are being considered as ARP routes which leads to this problem. [PR1041251](#)
- For MLPPP interface on MX Series routers with MPCs/MICs, in some very rare conditions, the received fragmented packets might be dropped. [PR1041412](#)
- This issue is applicable to a case which inline NAT configured on an interface belongs to either an MPCE or an MP3E/MPC4E/T4000-FPC5. Ingress and egress traffic traversing between an MPCE and these cards may cause the router to drop packets. [PR1042742](#)
- On T Series platform running Junos OS Release 12.1 or later, for interface connected via optical system like DWDM, when the interface is admin disabled, there might be a delay (300-400msec) for system to detect the event and during which time, traffic blackhole might be seen. Please note if disable the interface by breaking the Rx or Tx link, issue will not happen. [PR1043762](#)
- Queue stats on LSQ interfaces are not properly cleaned up when queuing is enabled on the IFD and the queues hosted at IFD level. This happens when a subsequent delete and create of the LSQ interface (not always though) - 14.1R4.10 [PR1044340](#)
- On MX Series platform with one of the following protocols configuration, flapping the protocols will trigger the Composite Next-hop change operation. In rare condition, since it is not properly programmed, the FPC might crash. This is a day-1 issue. - LDP - MPLS - Point-to-multipoint LSP - RSVP - Static LSPs [PR1045794](#)
- Once default route 0.0.0.0/0 is added, deleted or changed, the PFEMAN thread running on the MPC/FPC5 needs more than 600msec to program such changes. This is long

enough to trigger LFM or BFD flap. Junos OS Release 13.3R2 or later is exposed to this symptom. [PR1045828](#)

- On MX Series platforms, the unilist next-hop member will become 'replaced' status on Packet Forwarding Engine after interface flapping with ARP timeout. While the problem is happening, routing-table will display all right next-hop status but can not forward traffic since forwarding next-hop in Packet Forwarding Engine is in 'replaced' status and no longer active. [PR1046778](#)
- On T Series FPC 1-3 and M320 except E3-FPC with fib-local configuration. If there are multiple FIB local FPCs or the FIB local is a multiple Packet Forwarding Engine FPC, the TCP packets might get out of order, and packets re-ordering would occur. It reduces the application-level throughput for any protocols running over TCP. [PR1049613](#)
- In the PPP dual-stack subscribers environment, in rare condition, if bringing up 1000 dual-stack subscribers quickly, the PPP negotiation might fail. Then PPP retries negotiation, all subscribers fully establish. [PR1050415](#)
- Incorrect flow count is reported in the field 'count' of V9 header in all the packets sent to the collector. [PR1050543](#)
- This problem is because of a race condition, where other FPCs are not able to drain "which is 1 second" Fabric Streams connecting to FPC which is getting offline. With this situation - even when FPC comes online, other FPCs which have observed message "xmchip_dstat_stream_wait_to_drain" will not be able to send traffic to that particular FPC over fabric. There is no workaround. Rebooting FPCs which observed error message "xmchip_dstat_stream_wait_to_drain" is a recovery. [PR1052472](#)
- On all Junos OS based platforms, there are two different types of memory blocks that might be leaked. The first issue is rpd-trace memory block leak. There is one block each for any trace files opened for rpd. They could be leaked for each time a configuration commit is done. Around 40 bytes are leaked per operation. The issue does not occur in Junos OS releases prior to 14.1. The second issue is rt_parse_memory block leak which could happen during the configuration of aggregate routes, configuration information might not be freed. Around 16384 bytes are leaked per operation. This issue is a day-1 issue. [PR1052614](#)
- In subscriber management environment, the Berkeley Database (DB) may get into deadlock state. It is brought on by multiple daemons attempting to simultaneously access or update the same subscriber or service record. In this case, due to the access to DB were blocked by device control daemon (dcd), the subscriber management infrastructure daemon (smid) fails to recover the DB. Consequently, the router may stop responding to all the login/logout request as well as statistics activity. This timing related issue is most likely to occur during login or logout and when the system is busy. [PR1054292](#)
- In subscriber management scenario, when dynamic VLAN (DVLAN) demux interface is configured on MX Series router, the interface may get in stuck state. It could be observed that the statistics of demux0 may stop incrementing. This is because Session Database (SDB) may incorrectly calculate the number of subscribers over DVLANS. When the issue occurs, for example, the router may not able to process any PPPoE Active Discovery Initiation (PADI) packet, and fail to establish the PPPoE session. [PR1054914](#)

- OpenSSL project has published a security advisory for vulnerabilities resolved in the OpenSSL library on January 8th 2015: CVE-2014-3569, CVE-2014-3570, CVE-2014-3572, CVE-2014-8275, CVE-2015-0204, CVE-2015-0205. Refer to JSA10679 for more information. [PR1055295](#)
- In IPv6 environment, after enabling feature "solicit-router-advertisement-unicast", the IPv6 router may fail to reply the Router Advertisement (RA) to the IPv6 host as unicast only. To be exact, the IPv6 router may not only reply to the IPv6 host an RA as unicast to its link local address, but also send the RA as multicast to all nodes group (Multicast Address: ff02::1). The sample configuration could be found as follows: user@router> show configuration protocols router-advertisement interface ge-1/0/3.400 { ... solicit-router-advertisement-unicast; <<<<< "solicit-router-advertisement-unicast" feature is enabled ... } [PR1056599](#)
- IFCM error messages may occur in logs when it is not used. We lowered the severity of the message to avoid confusion. [PR1057712](#)
- In LDP tunneling over single hop RSVP based LSP environment, after enabling "chained-composite-next-hop", the router may fail to create the chained composite next hops if the label value of VPN is equal with the label value of LDP. [PR1058146](#)
- On MX Series routers, the interrupt-driven basis link down detection (an interrupt-driven link-down notification is generated to trigger locally attached systems to declare the interface down within a few milliseconds of failure) may fail after performing unified in-service software upgrade (ISSU). The interrupt might got prevented after performing unified ISSU due to disable the interrupt registers before unified ISSU but never restored after. [PR1059098](#)
- In an IPsec load-balancing environment using MS-MPC cards, the ICMP request and ICMP reply can go through two different IPsec tunnels due to asymmetric routing; that is, ICMP request goes through one PIC, and ICMP reply goes through another PIC. Because of this, the ICMP reply will get dropped and never reach the other side of the IPsec tunnel. [PR1059940](#)
- When enabling pseudowire subscribers the "show subscribers extensive" command does not display CoS policies applied to the subscriber interface. [PR1060036](#)
- Output MTU counter shows incorrect data in the show pfe statistics traffic command output [PR1061111](#)
- Due to incomplete fix, in releases containing PR869773 fix, rate limit drops are seen for Ingress queuing even though rate-limit is not configured or supported for ingress. [PR1061256](#)
- If Bidirectional Forwarding Detection (BFD) protocol is enabled via site-to-site IPsec tunnel, the BFD session may fail to come up. It is because, when the BFD protocol is trying to exchange the packet via IPsec tunnel, the value of the TTL in inner IP header for packet may be decremented, hence the BFD packet gets dropped on the peer side and no BFD session would come up. [PR1061342](#)
- On MX Series router with MPC2E-3D-NG/MPC3E-3D-NG/MPC5/MPC6 linecards, the Ethernet frame loss measurement (ETH-LM) feature does not work. [PR1064994](#)

- If there are application-sets matching conditions in the NAT rule, NAT port might leak after deleting applications under application-set in live network. [PR1069642](#)
- Higher baseline CPU utilization and periodic CPU spikes might be seen on XM-based MPC as compared to MPC-3D-16XGE-SFPP Cards due to below reasons: On XM-based MPC, low priority threads which monitor various things in the background on a periodic basis such as voltage, temperature, stats counters, hardware status etc are existed. When the system is idle these threads are allowed to take more of the load and that is why higher baseline CPU/CPU spikes are seen. This does not prevent other higher priority threads from running when they have to, as these are non-critical activities being done in the background and hence is a non impacting issue. [PR1071408](#)
- overhead-accounting frame-mode command does not work on 100GbE CFP MIC, 100GbE CXP MIC, 2x40GbE QSFP MIC, and 10x10GbE SFPP MIC on MPC3E-3D-NG-Q, MPC3E-3D-NG, MPC2E-3D-NG-Q, and MPC2E-3D-NG [PR1072001](#)
- On MX Series routers, the CLI command **set interfaces interface-name speed auto-10m-100m** is not supported. [PR1077020](#)
- From Junos OS Release 14.1R1, if the hidden knob "layer-4 validity-check" is configured, the Layer 4 hashing will be disabled for fragmented IP traffic. Due to a defect, the Multicast MAC rewrite is skipped in this case, the fragmented multicast packets will be sent with incorrect destination MAC. [PR1079219](#)
- The "inactivity-timeout" knob under [edit applications application application-name] hierarchy does not take effect for TCP based protocols. [PR1080464](#)
- 'show interfaces queue <ifl>' stats are not correct with RLSQ warm-standby mode. Issue seen on MX Series routers with MPCs or MICs as well in 14.1R4.10 [PR1082417](#)
- On a device with lt and ams interfaces configured, walking ifOutOctets or other similar OID's may cause a "if_pfe_ams_ifdstat" message to print. This is a cosmetic debug-level entry, which was incorrectly set to critical-level. [PR1085926](#)
- In some rare conditions, depending on the order in which configuration steps were performed or the order in which hardware modules were inserted or activated, if PTP master and PTP slave are configured on different MPCs on MX Series router acting as BC, it might happen that clock is not properly propagated between MPCs. This PR fixes this issue. [PR1085994](#)

Infrastructure

- Fsck flags -C flag would prevent fsck from fixing the filesystem. Removing this flag as per this PR. [PR1023164](#)
- On all Junos OS platforms, when the gstatd triggers false positives, this would result in unnecessary Routing Engine switchover happening. Thus a config option is added to prevent gstatd from initiating Routing Engine unnecessary switchover or Routing Engine relinquishing the mastership. The following error messages are expected to be seen: gstatd: [ad2] average write duration of 1021.34 crossed threshold of 1000.00 /kernel: mastership: routing engine 1 relinquishing as master: voluntarily requested [PR1024515](#)

- Mirroring to next-hop-subgroup stops working, when there is a change in the next-hop-subgroup configuration [PR1049631](#)

Interfaces and Chassis

- PR fix corrected jnxoptIfOTNPMFECIntervalTimeStamp, jnxPMIntTimeStamp, and jnxoptIfOTNPMIntervalTimeStamp reporting incorrect values around sytem-local midnight time as reported in PR 1065110. It also corrected the “SNMP PM Interval - incomplete date format without UTC offset”. [PR946014](#)
- On MX Series router, in rare condition, the kernel might crash and the router will go in db prompt when router reboots. [PR993978](#)
- MX-VC multiservices MIC and MPC: Offline is not initiated in ISSU if Sparks MIC is in the VC-B chassis from where VC ISSU is started. [PR997255](#)
- On MX Series routers with MPCs and MICs, in virtual private LAN service (VPLS) environment, the next hop in the kernel allocated by connectivity-fault management process (cfmd) may not be freed even after the CFM session has been removed (for example, deactivating the routing-instance). In this situation, after activating the routing-instance back, the interface within the routing instance would fail to come up because the nexthop is not freed by the cfmd application and hence the VPLS connection is down. [PR1000060](#)
- On standalone T series router or TX platform, during Routing Engine rebooting, a bad (or busy) I2C device on Switch Interface Board (SIB) might cause Switch Processor Mezzanine Board (SPMB) to crash. Please note the TXP platform might also experience same issue due the bad I2C, and it has been addressed in another PR, which has been fixed on Junos OS version 13.1R5 13.2R6 13.3R1 13.3R4 14.1R3 14.2R1 and 15.1R1. [PR1010505](#)
- Some duplicate entries are reported in jnx-chas-defines.mib. This patch removes the duplicate entries to fix the issue. [PR1036026](#)
- On Ethernet PICs with longer hold down timer configured, flapping interface within the hold time might cause traffic loss longer than the hold period. [PR1040229](#)
- At the end of the unified ISSU process, a link flap may be observed on SFP-T (tri-rate) interfaces. Now handling of SFP-T (tri-rate) interfaces has been improved to avoid link flap. [PR1040977](#)
- In case of the IQ2 or IQ2E PIC are working in tunnel-only mode, rebooting the tunnel PIC while the traffic is passing through the tunnel might cause the tunnel PIC to not transfer traffic any more. [PR1041811](#)
- “clear interfaces interface-set statistics all” command fails due to memory limitation. [PR1045683](#)
- On MX Series routers (platforms) with Enhanced Switch Control Board (SCBE), when the fan tray is inserted or pulled out, the chassisd process might crash. [PR1048021](#)
- When configuring the Virtual Router Redundancy Protocol (VRRP) on an interface which is included in a routing-instance via applying groups setting, if changes are made to the interface, the VRRP process (vrrpd) memory leak might be observed on the device. [PR1049007](#)

- dcd is cored by configuring IPv6 address on fxp0.0 with master-only option under interfaces configuration. [PR1049450](#)
- When Inherit is part of lower IFL Unit, vrrpd parses it before Active. In this case, vrrpd attaches a dummy Active to the Inherit, with the assumption that the Active will be available soon and then replication of information from Active to Inherit will take place. However, the replication of the priority might not be done correctly due to which the Inherit group gets stuck with priority of 0. [PR1051135](#)
- The "show chassis network-services" command might not show the correct configured value when executed on the backup Routing Engine. This command should only be executed on the master Routing Engine. [PR1054915](#)
- After performing a unified in-service software upgrade (ISSU) on MX Series Virtual Chassis (MX-VC) platform, all physical interfaces may go down. And the interfaces remain down until a graceful Routing Engine switchover (GRES) is performed. [PR1055327](#)
- There is a mismatch in mac statistics, few frames go unaccounted. This is a day-1 issue with the software fetching of mac statistics. The snap and clear bits were set together on pm3393 chip driver software, so it used to happen that even before the copy of stats to shadow registers happened, clear was happening which used to go unaccounted. Now rollover mechanism has been implemented and tested for two continuous days and everything is fine. [PR1056232](#)
- When a dynamic PPPoE subscriber with targeted-distribution configured on a dynamic vlan demux interface over aggregated Ethernet, the device control daemon (dcd) process might crash during a commit if the vlan demux has mistakenly been removed. The end users can't visit Internet after the crash. This is a rare issue and not easy to be reproduced. [PR1056675](#)
- In subscriber management environment, PPP client process (jpppd) might crash as a result of a memory allocation problem. [PR1056893](#)
- In a Multichassis link aggregation groups (MC-LAGs) environment, the MC-LAG peers have the MAC and port information and can forward the traffic appropriately. If a single VLAN on ICL interface is modified to a different VLAN, and then the administrator rolls back the VLAN configuration to the original one, the remote MAC might be stuck in the "Pending" state and not be installed in the bridge MAC-table, which causes the traffic forwarding to be affected. [PR1059453](#)
- For transit traffic on INLINE LSQ redundancy (rlsq) interface, the input firewall-filter counters are logging zero packet count regardless of traffic flow. Output filter counters are logging correctly. For host-bound traffic, the firewall output counter will get double accounted on Classical rlsq and triple accounted on INLINE rlsq. [PR1060659](#)
- In scaling PPP subscriber environment, when the device is under a high load condition (for example, high CPU utilization with 90% and above), the long delay in session timeout may occur. In this situation, the device may fail to terminate the subscriber session (PPP or PPPoE) immediately after three Link Control Protocol (LCP) keepalive packets are missed. As a result, subscriber fails in reconnect due to old PPP session and corresponding Access-Internal route are still active for some time. In addition to

this, it is observed that the server is still sending KA packets after the session timed out. [PR1060704](#)

- For multichassis link aggregation groups (MC-LAGs) running in active-active mode with back-to-back square topology, when the Inter chassis Control Protocol (ICCP) is broken between any MC-LAG devices, the non preferred device reverts to its own local system ID. But its Link Aggregation Control Protocol (LACP) partner on the remote side does not remove the flap link from AE bundle and it remains UP. This might cause a network wide loop resulting in traffic outage until manual intervention. [PR1061460](#)
- In connectivity fault management (CFM) environment, if an AE interface is included in MEP interfaces, and if there is another AE interface configured without any child link (even this AE is not participating in OAM), the CFM sessions might not come up after Routing Engine restart or switchover. [PR1063962](#)
- Error message is continuously logged every second after a particular copper-SFP [P/N:740-013111] is plugged into a disabled port on MIC. ***** error message *****
mic_sfp_phy_program_phy: ge-*/*/* - Fail to init PHY link mic_periodic_raw: MIC(*/*)
- Error in PHY periodic function PQ3_IIC(WR): no target ack on byte 0 (wait spins 2)
PQ3_IIC(WR): I/O error (i2c_stat=0xa3, i2c_ctl[1]=0xb0, bus_addr=0x56)
mic_i2c_reg_set - write fails with bus 86 reg 29 mic_sfp_phy_write:MIC(*/*) - Failed to write SFP PHY link 0, loc 29 mic_sfp_phy_mdio_sgmii_lnk_op: Failed to write: ifd = 140
ge-*/*/*, phy_addr: 0, phy_reg: 29 ala88e1111_reg_write: Failed (20) to write register:
phy_addr 0x0, reg 0x1d Fails in function ala88e1111_link_init [PR1066951](#)
- In PPPoE over AE subscribers management scenario, if "targeted-distribution" is enabled for subscribers IFL, the dcd process might crash and reboot when try to deactivate the AE interface. [PR1067062](#)
- When adding new VCP port MX-VC, some traffic drops are seen. [PR1067111](#)
- On MX Series Virtual Chassis (MX-VC) platform, due to a timing issue, the physical interface (ifd) on the same Modular Interface Card (MIC) with Virtual Chassis port (VCP) might not be created or take a very long time to be created after reboot the hosted Modular Port Concentrator (MPC). [PR1080032](#)
- In an MX Virtual Chassis (MX-VC) scenario, during a unified ISSU operation, the new master Routing Engine does not have the MX-VC SCC's system MAC address. It just has its local system MAC address. The address is not replicated between local Routing Engines, and the new master Routing Engine has not yet connected to the MX-VC SCC to receive it. Hence, the possibility exists to overwrite the FPC with an address that does not match the previous address. [PR1084561](#)
- The VRRP preempt hold time is not being honored during NTP time sync, and system time is changed. [PR1086230](#)

Layer 2 Features

- In DHCP dynamic subscriber management scenario, when maintain DHCP subscribers during interface delete is configured, some interface indices might be reused by a new interface if system is under stress (such as high connection speed, many clients and individual log files configured to be larger than 100M). In this case, it might result in subscriber being associated with an interface that no longer exists. [PR1044002](#)
- If the ppmd does not send replies to lacpd's periodic request to gather port statistics, the lacpd process may crash and restart due to the process memory consumption being slowly increased and finally reaching RLIMIT_DATA value which is 128MB. [PR1045004](#)
- On multiple Routing Engines system with NSR enabled, if the FEC129 VPLS instance has "no-tunnel-service" configured, the VPLS might show status as "OL" (no outgoing label) after performing Routing Engine switchover. [PR1050744](#)
- The Layer 2 Control Protocol process (l2cpd) leaks memory when interface config is applied to LLDP-enabled interfaces using 'apply-groups'. Size of the leak is ~700 bytes per commit. [PR1052846](#)
- After changing the way of getting the site ID of VPLS from fixed site-id to automatic-site-id on one site while other sites are still using the fixed site-id in the network, the rpd process might crash because the site ID gotten by "automatic-site-id" may conflict with the site ID which was configured as fixed site ID on other sites. [PR1054985](#)
- When MX Series router acts as the Virtual Extensible Local Area Network (VXLAN) Layer 3 gateway, the integrated routing and bridging (IRB) interfaces are configured to connect the VXLANs. The VXLAN packets are dropped when the route to reach a remote virtual tunnel endpoint (VTEP) interface is over an IRB interface. [PR1057005](#)
- BGP peer configured between two routers over It (logical tunnel) interface, if deactivating and activating scaled configuration a few times, in rare condition, the It interface might reject all the ARP reply packets, and hence the ARP resolution does not happen over this interface, so the unicast routes are not in the correct states, and pinging to such an It interface will fail. [PR1059662](#)
- With Dynamic Host Configuration Protocol (DHCP) maintain subscriber feature enabled, when the subscriber's incoming interface index is changed, for example, the interfaces go away and come back after changing the MTU configuration of interface, the existing subscribers may get dropped and new subscribers fail in connection. [PR1059999](#)
- After FPC is rebooted, the filter under Packet Forwarding Engine of the ERPS bridge domain may program wrong ifl index, it will cause the router to not be able to receive any ERPS packets. [PR1070791](#)
- LACP partner system ID is shown incorrectly when the AE member link is connected to a different device, this might misguide while troubleshooting the LAG issues. [PR1075436](#)
- On MX Series routers, when configuring the dynamic access routes for DHCP subscribers based on the Framed-Route RADIUS attribute, the access route may be created on

the device, however, the framed routes may not be installed for subscriber interface (under the "Family Inet Source Prefixes"). [PR1083871](#)

- MTU change is not advised on the Ethernet ring protection (ERP) ring interfaces unless ring is in idle condition. Changing ring interface MTU while ring is not in idle state may result in change in the forwarding state of the interface and which can lead to loop in the ring. [PR1083889](#)

MPLS

- When RSVP label-switched-path (LSP) optimize is enabled, RSVP LSP might stay down after a graceful Routing Engine switchover (GRES). To resolve the problem, the corresponding label-switched-path configuration needs to be deactivated, then, be activated again. [PR1025413](#)
- On the P2MP LSP transit router with link protection enabled, if the LSP is the last subLSP, tearing the last subLSP (for example, a RESV tear message is received from downstream router) might crash the routing protocol process (rpd). [PR1036452](#)
- LDP is not distributing a label for BGP FEC/prefix to downstream on demand (DoD) sessions when Forwarding Equivalence Class (FEC)/prefix learned this from IBGP peer to whom ldp-tunneling is configured. [PR1049329](#)
- In LDP link protection which is protected by dynamic RSVP LSP scenario, when flap the interface having LDP link-protection enabled, the rpd process might crash on backup Routing Engine as soon as the bypass LSP is established. [PR1053426](#)
- On M/MX/T Series routers, dynamic-rsvp-lsp is configured under interface link-protection hierarchy level. After interface flap, the bypass LSP does not come up. [PR1054155](#)
- With BGP prefix-independent convergence (PIC) edge feature enabled, more than one BGP next-hop association will be installed in the Packet Forwarding Engine for MPLS VPN and Internet transit traffic. Deactivating/activating the IGP protocol (IS-IS or OSPF) might cause the backup session to stay down on the Packet Forwarding Engine. [PR1058190](#)
- With graceful-restart configured, an inter-domain point-to-multipoint (P2MP) label switched path (LSP) with ERO defined and CSPF enabled might fail to come up after rpd process restart. [PR1058271](#)
- With BGP labeled-unicast egress protection is enabled in a Layer 3 VPN, the protected node advertises primary BGP labeled unicast routes that need protection. When there is next-hop change for a labeled route, for example, deactivating/activating egress-protection knob or route churn, the memory might be exhausted and leads to an rpd process crash. [PR1061840](#)
- This is a regression issue on all Junos operating systems related to a timing factor. When LDP session flaps, over which entropy label TLV or any unknown TLV is received, the LDP speaker might not send label withdraw for some prefixes to some neighbors. As a result, these neighbors will still use stale labels for the affected prefixes. [PR1062727](#)
- The point-to-multipoint (P2MP) label-switched path (LSP) is unable to re-establish after certain links are down. This issue might be encountered when the links are those

that contain the primary and backup LSPs for the P2MP LSP. The P2MP LSP can be restored after the links are up. [PR1064710](#)

- In Junos OS Release 14.1 and later, the "load-balance-label-capability" knob is introduced to enable the router to push and pop the load balancing label, which causes LDP and RSVP to advertise the entropy label TLV to neighboring routers. MX, T4000 and PTX platform have the capability and it is reflected in their default forwarding-options configuration. However, there is a software defect in the way how Entropy Label Capability (ELC) TLV is encoded in the LDP label mapping message. It might cause the LDP session between routers to go down. [PR1065338](#)
- Bypass enabled with optimize-timer will flap during every re-optimization event. [PR1066794](#)
- When CSPF computes the path for node-protected bypass, it considers only the SRLG group configured on next-hop interface along the primary path. However it doesn't consider the SRLG group on next-to-next-hop interface to adequately provide diverse path between primary and node-protected bypass. [PR1068197](#)
- When a primary LSP gets re-routed due to better metric, Link/Node protection for this LSP is expected to come up within 7 seconds provided the bypass-lsp protecting the next-hop link/node is already available. However in some corner cases, the Link/Node protection for re-routed primary LSP will not come up within 7sec even with bypass-lsp availability. The PR fixes this issue and reduces the delay of associating bypass-lsp with primary-lsp from 7 seconds to 2 seconds. [PR1072781](#)
- In scaled l2circuits environment, the rpd process might crash due to a corruption in the LDP binding database. [PR1074145](#)
- In race conditions, the rpd process on backup Routing Engine might crash when BGP routes are exported into LDP by egress-policy and configuration changes during the rpd process synchronizing the state to backup rpd process. [PR1077804](#)

Network Management and Monitoring

- When syslog server is configured using hostname, after Routing Engine switchover router stopped sending the syslogs to external syslog server. Immediately after switchover, DNS was not accessible because it will take some time to learn route to DNS. System stopped retrying DNS resolution and syslogging stopped. System was running GRES (no NSR). [PR947869](#)
- If the size of interface control process (dcd) trace file is configured to large (e.g. 500M), restarting the dcd process when it is doing log rotation (due to size limit reached) might cause dcd process unable to start any more, in the same time, interface ports will be down and "thrashing" error will be seen. [PR1047330](#)
- SNMP mib walk jnxMac does not return value with et- interfaces on MPC3/MPC4/MPC5/MPC6 [PR1051960](#)

- There is no specific counter name in the MIB2D_COUNTER_DECREASING syslog message [PR1061225](#)
- SNMP queries for LAG MIB tables while LAG child interface is flapping, may cause mib2d grow in size and eventually crash with a core file. Mib2d will restart, and recover by itself. [PR1062177](#)

Platform and Infrastructure

- When Network Configuration Protocol (NETCONF) service is used on the device, after the NETCONF session is established, because all the output that contain <error> tag might be incorrectly converted into <rpc error>, the management daemon (mgd) may crash on the device. As the example below, the output that contains <error> tag may lead to the crash. `user@re0> show subscribers address 1000 | display xml .. <<<<<<` The output contain <error> tag and may trigger the crash. [PR975284](#)
- If the system has service related configurations, Error message generated by mountd might be seen: "can't delete exports for /packages/mnt/jbase: Bad address". [PR991814](#)
- For inline BFD over aggregated Ethernet (AE) interface which member links are hosted on different FPCs, BFD packets coming on ingress line card will be steered to anchor Packet Forwarding Engine through fabric. If FPC reconnect to master Routing Engine (such as Routing Engine switchover operation), the inline BFD session punts the BFD packet to host, the BFD packet should go through loopback interface filter of VRF on which it is received. But in this case, the BFD packet might hit the wrong loopback interface filter from wrong routing-instance since the VRF information is not carried across fabric. [PR993882](#)
- CPQ RLDRAM ECC single and double bit error will generate CM alarm. "show chassis alarms" command can be used to view CM alarm. Details ===== 1> CPQ RLDRAM ECC single bit error in last 10 secs will raise minor CM alarm. 2> No CPQ RLDRAM ECC single bit error in last 10 secs will clear minor CM alarm. 3> CPQ RLDRAM ECC double bit error will raise Major CM alarm (this alarm will not be cleared until the FPC is restarted) [PR1023146](#)
- On MX platform with scaled set-up, after deactivate/activate or renaming a bridge domain (BD) which has irb interface associated, the IGMP snooping configured under the BD might not work any more. Please note it happens only when the router is in "network-services enhanced-ip" mode. [PR1024613](#)
- Recurring local memory (LMEM) data errors may cause LU chip (lookup chip on Trio based FPC) wedge and eventually FPC crash. [PR1033660](#)
- The Priority code point (PCP) and Drop eligible indicator (DEI) bit in 802.1Q header are preserved while packet gets routed within the same Packet Forwarding Engine. The expected behaviour is resetting the PCP and DEI bit when the packet is routed. [PR1036756](#)
- Presence of /8 prefix in two terms results in incorrect filter processing and unexpected behavior. [PR1042889](#)
- When IRB interface is configured with VRRP in Layer 2 VPLS/bridge-domain, in corner cases IRB interface may not respond to ARP request targeting to IRB sub-interface IP address. [PR1043571](#)

- On MX Series platform with Extensible Subscriber Services Management (ESSM) subscribers configured using Junos OS commit script, after performing sequence of operations repeatedly with same set of configuration (subscribers apply-macros'), like adding subscribers, then deleting same subscribers again, then adding, then deleting again and again like so, the memory would leak on mgd process. In a generic scenario where a script just commits transient change and then exits, the issue will not be experienced. [PR1048770](#)
- By default, after Agent Smith boards comes up about 75% of queues were allocated to support rich queuing with MQ chip. Such allocation causes MQ driver software module to poll stats. Polling stats causes this rise in CPU usage. [PR1048947](#)
- For a Routing Matrix, if different Routing Engine models are used on switch-card chassis (SCC)/switch-fabric chassis (SFC) and line-card chassis (LCC) (for example, RE-1600 on SCC/SFC and RE-DUO-C1800 on LCC), where the out-of-band (OoB) management interfaces are named differently (for example, fxp0 on SCC/SFC RE and em0 on LCC RE), then the OoB management interface configuration for LCC RE will not be propagated from SCC/SFC RE during commit. [PR1050743](#)
- NTP.org has published a security advisory for multiple vulnerabilities resolved in ntpd (NTP daemon) that have been assigned four CVE IDs. Junos has been confirmed to be vulnerable to one of the buffer overflow vulnerabilities assigned CVE-2014-9295 which may allow remote unauthenticated attackers to execute code with the privileges of ntpd or cause a denial of service condition. Refer to JSA10663 for more information. [PR1051815](#)
- Values for the "input-traffic-control-profile" statement get reset after deactivating/activating the traffic-manager mode. [PR1052785](#)
- Change of MED value fails when configure private is issued [PR1055178](#)
- On Trio-based line cards, in forwarding topology that has unilist next-hop and the unilist next-hop is pointing to indirect next-hop or another unilist next-hop, if there is a change happens on one of the downstream paths, the change may fail to be propagated to the unilist next-hop. As a result, the unilist selector table is not updated. This might lead to traffic blackholing. [PR1056150](#)
- While using certain 14.1 daily builds and the JAM package users might notice erroneous outputs in the fields of Domain ID and FlowSet ID. This issue is fixed in Junos 14.1R4. [PR1057450](#)
- Under very rare situations, Packet Forwarding Engines on the following linecards, as well as the compact MX80/40/10/5 series, may stop forwarding transit traffic: - 16x10GE MPC - MPC1, MPC2 This occurs due to a software defect that slowly leaks the resources necessary for packet forwarding. Interfaces handled by the Packet Forwarding Engine under duress may exhibit incrementing 'Resource errors' in consecutive output of 'show interfaces extensive' output. A Packet Forwarding Engine reboot via the associated linecard or chassis reload is required to correct the condition. [PR1058197](#)
- With the configuration "extend-size", if user loads and commits scaled configuration (in this case, 250K Unique Prefix list policy options), then deletes the knob "extend-size", the dfwd process might crash. [PR1058579](#)

- After committing the Network Time Protocol (NTP) configuration, if the number of routing-instances per source-address exceeds 18, it may cause NTP daemon (ntpd) crash. In this scenario, the NTP feature may not be functional. For example there are 19 routing-instance names per source address statement in the sample configuration below. `ntp { server X.X.X.X; source-address X.X.X.X routing-instance [X1 X2 X3 X4 X5 X6 X7 X8 X9 X10 X11 X12 X13 X14 X15 X16 X17 X18 X19]; (19 routing-instance names) }` [PR1058614](#)
- On MX Series line cards with MPCs and MICs with Junos OS Release 12.3R3 and later, the system does not push the configured Tag Protocol ID (TPID) value (for instance, 0x88a8) to the packets while sending out the packets; instead it pushes default TPID 0x8100. This might lead to traffic drop on the peer device if it is expecting a particular TPID (for instance, 0x88a8) whereas it receives a different one. [PR1059225](#)
- Modifying IEEE-802.1ad rewrite-rule on the fly might make it unable to change IEEE-802.1p ToS values for inner VLAN in QinQ. [PR1062817](#)
- When MX platform acts as Virtual Extensible Local Area Network (VXLAN) gateway, if there are multiple Packet Forwarding Engines, VXLAN packets will be distributed to available Packet Forwarding Engines in the chassis to perform VXLAN encapsulation/decapsulation, this is not expected (Expect behavior: VXLAN packet processing will be done on the same Packet Forwarding Engine on which it is received). This might result in unexpected packet drop and also overlay ping/traceroute not working. [PR1063456](#)
- Observation domain ID in exported flow records is incorrect in MPC 3E cards. and 200G 40x10G MPC and 200G 4x100G MPC for the MX Series. [PR1066319](#)
- On MX Series routers with MPCs and T4000 routers with Type 5 FPCs, the feature "enhanced-hash-key" is configured to select data used in the hash key for enhanced IP forwarding engines. If "type-of-service" is configured at the [edit forwarding-options enhanced-hash-key family inet] hierarchy level, or "traffic-class" is configured at the [edit forwarding-options enhanced-hash-key family inet6] hierarchy level, the last significant 2 bits of the TOS/TC bytes under the IPv4/IPv6 header are extracted incorrectly as load sharing input parameters, this might cause unexpected load balancing result. [PR1066751](#)
- StartTime and EndTime of the flow in inline-jflow (version 9) has future time-stamp [PR1067307](#)
- Firewall filters that have a prefix-action cannot be configured under [edit logical-system <name> firewall family inet] because the Packet Forwarding Engine won't be programmed for the filter. [PR1067482](#)
- An FPC with interfaces configured as part of an Aggregated Ethernet bundle may crash and reboot when the shared-bandwidth-policer is configured as part of the firewall policer. [PR1069763](#)
- If with about 1M routes on MX series router, there might be more than 1 second (about 1.3s) packets dark window during unified ISSU. [PR1070217](#)
- On MX series routers, when using MX Series based FPC with feature inline sampling activated, memory partition error messages and memory leak might be observed on the FPC. In some cases, this issue only affects sample route-records but not regular

Packet Forwarding Engine routes or next-hops. However, in the extreme case, it is also possible to cause the Packet Forwarding Engine failing in installing routes into forwarding next-hops and hence traffic drop. On MX series routers, when using MX Series based FPCs, Junos OS 13.3R5 14.1R4 14.2R1 or higher is exposed. On T4k or TXP-3D routers, when using FPC-3D FPC's, Junos OS 14.2R1 or later is exposed.

[PR1071289](#)

- VPLS filter applied under forwarding-options might drop VPLS frame unexpectedly when it's coming from an lt- interface. [PR1071340](#)
- When inline-sampling is enabled, in race conditions, if packet gets corrupted and the corrupted packet length shows 0, which may cause "PPE_xErrors thread timeout error" and eventually cause MPC card to crash. [PR1072136](#)
- After IPv6 RPM(real-time performance monitor) support, snmp server cannot receive some of IPv6 PING-MIB info. For example, snmp server receives "pingCtlRowStatus(23)" and "pingCtlAdminStatus(8)" error and cannot get "pingResultsTable" and "pingProbeHistoryTable" info. << example >> ** The following logs are snmp server logs. "snmpset -v 2c -c xxxxxx" commands are used. ---pingCtlRowStatus(23) error info. Error in packet. Reason: inconsistentValue (The set value is illegal or unsupported in some way) Failed object:
SNMPv2-SMI::mib-2.80.1.2.1.23.779.87.78.69.82.95.65.6.84.69.83.84.95.65
---pingCtlAdminStatus(8) error info. Error in packet. Reason: inconsistentValue (The set value is illegal or unsupported in some way) Failed object:
SNMPv2-SMI::mib-2.80.1.2.1.8.779.87.78.69.82.95.65.6.84.69.83.84.95.65 ** The following logs are snmp server logs. "snmpwalk -v 2c -c xxxxxx" commands are used. pingResultsTable(3) SNMPv2-SMI::mib-2.80.1.3 = No Such Object available on this agent at this OID pingProbeHistoryTable(4) SNMPv2-SMI::mib-2.80.1.4 = No Such Object available on this agent at this OID [PR1072320](#)
- Problem: MAC filter ff:ff:ff:ff:ff:ff is cleared from the Packet Forwarding Engine hardware mac table. So arp requests are not forwarded to irb. Fix: Not all mac entries pointing to invalid l2 token are candidates for being deleted. Static mac entries are managed by control plane only. So Packet Forwarding Engine cannot delete these entries. The logic for skipping mac deletion for static mac entries done earlier is not proper Packet Forwarding Engine. Fixed the same. [PR1073536](#)
- When an MX Series chassis network-services is "enhanced-ip" and an aggregated Ethernet with "family bridge" configuration is first committed, there is a possibility that an incorrect forwarding path may be installed causing traffic loss. [PR1081999](#)
- LMEM is an internal memory in LU/XL ASIC chip. It has private and shared regions for Packet Processing Engines. LMEM data errors are very rare events caused by environmental factors (this is not created by software). Due to a software defect, an error in the shared LMEM region will result in corruption of critical data structures of Packet Processing Engines that causes unpredictable communication of LU/XL ASIC chip with MQ/XM ASIC chip. These events will corrupt the state in MQ/XM and lead to a MQ/XM wedge. The MQ/XM wedge would cause fabric blackhole and finally reboot the line card. [PR1082932](#)

- On MX Series router with MPCs/MICs, the "RPF-loose-mode-discard" feature is not working when configured within a Virtual Router routing instance. The feature is working only when configured in the main instance. [PR1084715](#)
- With MX Series based FPC, load balance hash seed will be changed after ISSU. Since the hash seed value will be reverted to original value by rebooting FPC, there would be hash value inconsistency in the system which might introduce blackholing on multicast flavor traffic (mcast or BUM on vpls/l2-bridge). Affected versions (Other versions do not have the issue) 12.3R7, 13.1R5 and later, 13.2R4 thru 13.2R5, 13.3R2 thru 13.3R3, 14.1R1 and later, 14.2R1 and later. [PR1086286](#)
- The prompt for SSH password changed in Junos OS 13.3, from "user@host's password:" to "Password:". This change breaks the logic in "JUNOS/Access/ssh.pm" which is located in /usr/local/share/perl/5.18.2/ on Ubuntu Linux, for example. [PR1088033](#)

Routing Protocols

- If with both BGP Prefix-Independent Convergence (PIC) edge and "routing-options multipath" configured, when the primary path fails, the protection provided by BGP PIC edge might not work correctly. [PR1011596](#)
- If with BGP PIC edge feature enabled and OSPF protocol as IGP, when the primary route changed, there is a chance that the Packet Forwarding Engine forwarding entry will stay in reroute state which causes session down. [PR1015598](#)
- When BGP add-path feature is enabled on BGP route-reflector (RR) router, and if the RR router has mix of add-path receive-enabled client and add-path receive-disabled (which is default) client, due to a timing issue, the rpd process on RR might crash when routes update/withdraw. [PR1024813](#)
- RIP is applying the RIB import-policy for the primary RIB table, as per the policy configured evaluation fails and routes are removed from primary RIB. But import-policy is applied only for secondary tables. RIP should apply only the protocol import policy and add routes to primary RIB. Routes are leaked to secondary routing table according to import-policy. Fix: As suggested by rpd infrastructure team, removed the import policy filter application to primary routing table by protocol rip. Now import policy application is handled by policy module within RPD. [PR1024946](#)
- When a BGP peer goes down, the route for this peer should be withdrawn. If it happens that an enqueued BGP route update for this peer has not been sent out, issuing the CLI command "show route advertising-protocol bgp <peer-addr>" might crash the routing protocol process (rpd). This is a very corner issue and hardly to be experienced. [PR1028390](#)
- If precision-timers and traceoptions are enabled for BGP then both main-thread and precision-timers pthread try to rotate the same tracefile without taking any locks. As a result all the status commands for rpd and krt may timed-out. [PR1044141](#)
- If labeled BGP routes are leaked from inet.3 table to inet.0, then activation of BGP "add-path" feature might crash the routing process (rpd). [PR1044221](#)
- BFD session might reset on commit if version is configured. The adaptive RX interval gets set to 0 which results in the reset. A sample configuration of BFD version is as

following: protocols { bgp { bfd-liveness-detection { version 1; minimum-interval 1000; transmit-interval { minimum-interval 1000; } } } } [PR1045037](#)

- When BGP and ICCP are the client of the same multi-hop BFD session, BFD runs in centralized (non-distributed) mode. But if nonstop-routing configuration is added and enabled, runing mode of BFD is changed to distributed mode. This behavior is incorrect but it would not affect to protocols which is client of the BFD session. However, if Routing Engine switchover is performed after enabling NSR, the BFD session will get unstable and all the client protocols also get unstable. [PR1046755](#)
- The Junos OS Multicast Source Discovery Protocol (MSDP) implementation is closing an established MSDP session and underlying TCP session on reception of source-active TLV from the peer when this source-active TLV has an "Entry Count" field of zero. "Entry Count" is a field within SA message which defines how many source/group tuples are present within SA message. [PR1052381](#)
- Either "rib inet.3" or "resolve-vpn" feature is available to be configured in the lower hierarchy for BGP labeled-unicast family routes. These two features are mutually exclusive and only one of them could be used at a single BGP group. If the administrator swaps the two features (for example, using the "resolve-vpn" first, then deactivate it and using "rib inet.3" instead, then use "resolve-vpn" back), the secondary routes (routes in inet.3 which including the ones from this BGP group and from other BGP groups) may got accidentally removed every time on "commit" operation take place. [PR1052884](#)
- After deactivating/deleting BFD configuration, Packet Forwarding Engine receives BFD session down event and it marks corresponding nexthops as down and traffic drops consequently. [PR1053016](#)
- The BGP session sending add-path prefixes can cause an rpd crash when the add-path IDs that it allocates roll over from 65535 to 0. If the routes contributing add-path prefixes are changing, the allocated path-id can eventually reach this value. This fix changes the allocation scheme to always use the lowest available free path-id, so a rollover will never occur. [PR1053339](#)
- The routing protocol process (rpd) might crash when static reverse-path forwarding (RPF) selection is configured and the upstream interface in the VRF routing instance disabled. [PR1054913](#)
- After multicast traffic source incoming interface and source ip RPF (reverse path forwarding) route switching to a different interface, the multicast route cache upstream interface might not be refreshed to be in sync with the pim join upstream interface. This is incorrect and will cause packet blackhole for the affected multicast stream. [PR1057023](#)
- Deletion of a routing-instances may lead to a routing daemon crash. This may happen if routing-instance's Routing Information Bases (RIB) is referenced in an active policy-option configuration. As a workaround, when deactivating the routing-instance, all associated configurations using the route-table names in the routing-instance should also be deactivated. [PR1057431](#)

- When running Simple Network Management Protocol (SNMP) polling to specific ISIS Management Information Base (MIB) with invalid variable, it will cause routing protocol process (rpd) crash. [PR1060485](#)
- In PIM environment, Bootstrap Router (BSR) can be used only between PIMv2 enabled devices. When deactivating all the interfaces which are running PIM bootstrap, the system changes to operate in PIMv1. At this time, all the information learned about/from the current BSR should be cleaned, but actually, BSR state is not cleaned. If the interface which was the previous "elected BSR" is activated, BSR state is PIM_BSR_ELECTED(should be cleaned previously) and the system assumes the BSR timer is still here. When the system tries to access the null BSR timer, the rpd process might crash. [PR1062133](#)
- In Protocol Independent Multicast (PIM) sparse mode environment, in the situation that the router is being used as the rendezvous point (RP) also the last hop router, when the (*,G) entry is present on the RP and a discard multicast route (for example, due to receiving multicast traffic from non-RPF interface) is already existed, if the (S,G) entry is learnt after receiving source-active (SA) of the Multicast Source Discovery Protocol (MSDP), the SPT cutover may fail to be triggered. There is no traffic impact as receivers still can get the traffic due to (*,G) route. [PR1073773](#)
- In an MPLS L3VPN Core network, enable BGP Prefix-Independent Convergence (PIC) Edge feature on a PE router, if the same VPN route is received with different multiple exit discriminator (MED) via two route reflectors (RR), when BGP PIC evaluates those two routes, it disregards the one with higher MED hence fails to build a multipath protection/backup path entry. [PR1079949](#)
- When removing scale BGP configuration, if the BGP session are holding stale routes for the benefit of a restarting peer, the routing protocol process (rpd) may crash. As a workaround, the administrator may use CLI command "show route receive-protocol bgp <peer address> extensive | match STALE" to find the existing stale routes. If there are none, then removing the BGP configuration may not cause the rpd crash. [PR1081460](#)
- If a policy statement referred to a routing-table, but the corresponding routing instance is not fully configured (ie. no instance-type), commit such configuration might cause the rpd process to crash. [PR1083257](#)
- With Multicast Source Discovery Protocol (MSDP) and nonstop active routing (NSR) configured on the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP), the rpd process might permanently get stuck when multicast traffic received shortly after Routing Engines switchover. [PR1083385](#)
- When there are a number of secondary BGP routes in inet.0, an SNMP walk of inet.0 by the bgp4 MIB can cause a core if the corresponding primary routes are being deleted. [PR1083988](#)
- When BGP route is leaked to a routing-instance and there is an import policy to overwrite the route preference, if damping is also configured in BGP, the BGP routes which were copied to second table can't be deleted after routes were deleted in master table. This is a day-1 issue. [PR1090760](#)

Routing Policy and Firewall Filters

- When configuring the unsupported IPv6 flow specification feature, that is, when configuring inet6 address as source/destination of inet-flow route, the configuration can pass the commit check and being committed. But it can cause rpd process crash eventually when trying to program this route to firewall process (dfwd, which manages compilation and downloading of Junos firewall filters). If a flow route is received from a BGP neighbor and prefix-length for source/destination is greater than 32, it can lead to rpd process crash too. [PR1059542](#)

Services Applications

- On M/MX/T Series routers (platforms) with Services PIC, the incoming interface is a services interface. If the services interface receives "ICMP MTU Exceeded" message, the message might be dropped. [PR977627](#)
- Added support to bring up Tunnel-switched sessions when tunnel-group is not configured at LTS and tunnel attributes are returned from RADIUS. [PR1030799](#)
- When using both Port Control Protocol (PCP) and traditional NAT (e.g. DSLITE), if you try to setup two pools with overlapping address ranges, this can lead to MS-DPC to crash and generate a core file. [PR1036459](#)
- On M/MX/T Series routers with Multiservices 100, Multiservices 400, or Multiservices 500 PICs with "dump-on-flow-control" configured, if prolonged flow control failure, the coredump file might generate failure. [PR1039340](#)
- Inline IPv6 L2TP on MPC subscriber terminated at a LNS breaks adaptive services SP unicast nexthops on MS-DPC. Even one subscriber causes the issue. [PR1054589](#)
- When the tunnel between L2TP access concentrator (LAC) and L2TP network server (LNS) is destroyed, the tunnel information will be maintained until destruct-timeout expire (if the destruct-timeout is not configured, the default value is 300 seconds). If the same tunnel is restarted within the destruct-timeout expire, the LNS will use the previously negotiated non default UDP port, which might lead to the tunnel negotiation failure. [PR1060310](#)
- A Layer 2 Tunneling Protocol daemon (l2tpd) crash is seen sometimes when the L2TP service interface unit number is configured higher than 8192. A restriction has been added to force unit numbers below 8192. [PR1062947](#)
- On MX Series routers which are acting as LNS to provide tunnel endpoints, it is observed that the service-interfaces are not usable if an MIC corresponding to them is not physically installed on the FPC. If only those service interfaces that belong to the removed PIC are added to service-device-pool, this results in no LNS subscribers able to login. Note that once the MIC is inserted into the FPC, the features could be used. [PR1063024](#)
- When configuring RADIUS authentication for Layer 2 Tunneling Protocol (L2TP), the RADIUS server cannot be recognized because the source address is not being read correctly. As a result, the L2TP session cannot be established. [PR1064817](#)
- L2TP daemon will core in LTS scenario while the subscriber logs out. This happens when the subscriber has "Called Number AVP" attribute. The "Called Number AVP"

was not getting relayed correctly across LTS boundary, hence daemon cores.

[PR1065002](#)

- Service PIC daemon (spd) might crash with core-dumps due to CGNAT pool's snmp-trap-thresholds configuration. [PR1070370](#)
- In CG-NAT or stateful firewall environment, due to a null pointer check bug, the MS-DPC might crashed every few hours. Note that this is a regression issue. [PR1079981](#)
- The crash happens if in a http flow, the flow structure is allocated at a particular memory region. There is no workaround but the chances of hitting this issue are very low [PR1080749](#)
- On Layer 2 Tunnel Protocol (L2TP) network server (LNS), during L2TP session establishment, when receiving Incoming-Call-Connected (ICCN) messages with Last Sent LCP CONFREQ Attribute Value Pair (AVP) but without Initial Received LCP CONFREQ and Last Received LCP CONFREQ AVPs, the jl2tpd process might crash. [PR1082673](#)
- In a L2TP tunnel-switching scenario, if a tunnel-switched tunnel is cleared with "clear services l2tp tunnel peer-gateway" AND an incoming ICRQ is received simultaneously from the LAC side destined for this tunnel-switched tunnel, this leads to jl2tpd crash. This defect has now been rectified. [PR1088355](#)

Software Installation and Upgrade

- Due to a software defect found in 14.1R5.4, Juniper Networks strongly discourage the use of Junos software version 14.1R5.4 on routers with MQ-based MPC. This includes MX Series with MPC1, MPC2; and all mid-range MX-Series. [PR1108826](#)

Subscriber Access Management

- This issue was introduced as part of another fix. Please contact JTAC for the recommended release for your deployment. [PR1049955](#)

VPNs

- For VPLS over VPLS topology, when the VPLS payload has two labels (Customer-VPLS-label and Customer-MPLS-label), the frame might be dropped by the core facing interface hosted on IQ2 PIC with "L2 mismatch timeout" error. This particular scenario is fixed. But there are some other worse scenarios which might hit this issue again due to the system architecture limitation, which are not fixed but need to avoid: * Addition of VLAN tags on Service provider's or CE's VPLS payload e.g. configuring QinQ. * Addition of MPLS tags on Service provider or CE's VPLS payload. * Enabling VPLS payload load balancing on Service provider's PE router. [PR1038103](#)
- In next-generation MVPN, after the route to C-RP flaps, traffic loss might be seen for a short period of time. [PR1049294](#)
- In next-generation MVPN scenario, when a source is directly connected to a PE that is acting as an RP stops sending the traffic, the PE never withdraws the Type 5 route. This causes the Type 7 routes and forwarding routes to remain on the egress and ingress PEs. [PR1051799](#)

- In L2VPN scenario with local switching enabled, in corner cases, the rpd process might crash after flapping the PE-CE link. For example, if the L2VPN connection type changes from remote to local after link flaps, for a brief period of time, two route entries (for old remote VC connection and for the new local VC connection) might exist for the same egress route (with interface name as destination prefix). In that case, when deleting remote VC connection and route entry associated with that remote connection, the rpd might crash due to trying to reset an internal variable which is already reset during route addition for the new local VC connection. [PR1053887](#)
- In the l2circuit environment, when l2ckt configuration has backup-neighbor, the flow-label operation is blocked at the configuration level. [PR1056777](#)
- With static selective point-to-multipoint LSP configured for an MBGP MVPN, when sending Type 3 S-PMSI A-D BGP route, the Juniper Networks implementation uses the values taken from the selective P-Tunnel configuration, which is not compliant with RFC 6514 section 4.3, which specifies that the source and group length values in Type-3 must be the same as the host prefix length, that is, if the Multicast Source field contains an IPv4 address, then the value of the Multicast Source Length field is 32; if the Multicast Source field contains an IPv6 address, then the value of the Multicast Source Length field is 128. The same is true for group length. [PR1058193](#)
- In a next-generation MVPN scenario, while traffic is not being generated by source for at least 3 and a half minutes and a routing or other multicast issue prevents the multicast traffic from reaching the receiver PE, after the multicast data starts flowing again for about 6 minutes, the Type-7 and Type-5 routes might be withdrawn which causes a discard route to remain present on the RP facing PE and causes the traffic not to be forwarded even if there is state and flowing traffic for that group. [PR1058574](#)
- The rpd process might crash when deactivating a logical system with nonstop active routing (NSR) enabled and BGP multicast virtual private network (MVPN) configured. [PR1059057](#)
- In MVPN RPT-SPT mode, with a mix of local and remote receivers all using (*,g) joins (spt-threshold infinity), the downstream interfaces may not get updated properly and there may be a stuck (s,g) forwarding route. This issue can occur with the following sequence of events: 1. Local receivers are joined 2. Traffic starts, then stops, and the route times out. 3. Remote receiver joins. Both a (*,g) and an (s,g) forwarding route are created. 4. Another local receiver is joined, or an existing one is pruned. 5. In the (*,g) route the downstream interface list reflects the update, but in the (s,g) route the downstream interface list does not. 6. When traffic starts again, the (s,g) route -- which has the wrong interface list -- is used. The traffic flows to the wrong set of receivers. [PR1061501](#)
- Problem, trigger and symptom: On a dual Routing Engine, if mvpn protocol itself is not configured, and nonstop active routing is enabled, the show command "show task replication" on master Routing Engine will list MVPN protocol even though it is not configured. Other than the misleading show output which may be slightly confusing to the user/customer, there is no functional impact due to this issue. There is no workaround available. [PR1078305](#)

Resolved Issues: 14.1R4

- [Class of Service \(CoS\) on page 145](#)
- [General Routing on page 145](#)
- [High Availability \(HA\) and Resiliency on page 149](#)
- [Interfaces and Chassis on page 149](#)
- [Layer 2 Features on page 150](#)
- [MPLS on page 150](#)
- [Network Management and Monitoring on page 151](#)
- [Platform and Infrastructure on page 151](#)
- [Routing Policy and Firewall Filters on page 153](#)
- [Routing Protocols on page 153](#)
- [Services Applications on page 153](#)
- [VPNs on page 154](#)

Class of Service (CoS)

- For an ATM interface configured with hierarchical scheduling, when a traffic-control-profile attached at ifd (physical interface) level and another output traffic-control-profile at ifl (logical interface) level, flapping the interface might crash the FPC. [PR1000952](#)
- Sometimes MX Series might respond with "no such instance" of the second OID when two CoS OIDs in the single SNMP packet. [PR1015342](#)
- This issue specific to rate-limit on trunk port in DPC due to a software issue that installing rate-limit variables to egress Packet Forwarding Engine does not work normally. [PR1022966](#)
- For ichip based platform, IQ2 pic expects FC index in the cookie from ichip for packet queuing. For Transit traffic, fc index is coming in cookie where are for host outbound traffic, queue number is coming in cookie to IQ2 pic. As IQ2 pic is not aware whether traffic is transit or host outbound, it treats value received in cookie as FC value and looks into fc_to_q table to fetch queue number. This is causing issue in queueing of host outbound traffic in IQ2 PIC in incorrect queue. This is a day one issue and will come if in FC to Queue mapping, fc id and queue number are not same. [PR1033572](#)
- This error message "only per-unit and 2-level hierarchical scheduler are supported on this interface" is a cosmetic regression issue without any functional impact. [PR1050512](#)

General Routing

- **show services accounting usage** does not populate CPU utilization for XLP-based cards. Use **show services service-sets cpu-usage**. [PR864104](#)
- On MX Series platform with enhanced DPCs equipped, after router rebooted, the IRB broadcast channel is not enabled, and all the broadcast packets that are received in the IRB interface will get dropped. Also when ping is given the below L2Channel error increases as ping packets are sent: user@router>show interfaces ge-*/*/ extensive

| match channel L3 incompletes: 0, L2 channel errors: 10, L2 mismatch timeouts: 0
[PR876456](#)

- When mirror destination interface is a next-hop-subgroup and enhanced-ip chassis knob is enabled, family any mirroring applied on L3 interfaces (inet/inet6) might not work in certain scenarios. [PR972138](#)
- In the dual Routing Engines scenario with large scale nexthops (in this case, more than 1-million nexthops and around 8K VRFs). In rare condition, kernel might crash on backup and/or master Routing Engine due to exhaustion of nexthop index space. [PR976117](#)
- On MX Series routers, delete an interface A from routing-instance VRF1; then create routing-instance VRF2 and interface A is added to VRF2 with qualified-next-hop configured; finally, delete VRF1. Commit the entire above configuration once, in rare condition, rpd might crash. [PR985085](#)
- In the dual Routing Engines scenario, in rare condition, while executing GRES and deleting interfaces at the same time, it is possible that a nexthop delete message is not sent to rpd process, causing rpd to keep a nexthop index (NHID) that kernel has already deleted. Later when kernel allocates this NHID for next new nexthop and sends it to rpd process, rpd process might crash due to duplicate NHID. [PR987102](#)
- An EVPN with support for inter-subnet routing using an irb interface may experience a crash and restart of rpd, leaving a core file for analysis. In this case, EVPN MAC routes contain MAC+IP, and this IP/32 is installed in VRF table on egress router. Core is triggered in the IP/32 route installation flow. There is no special trigger point- it's a timing issue with basic irb configurations. [PR992059](#)
- In Ethernet VPN (EVPN) routing and bridging (IRB) deployment, when all the access interfaces go down under an EVPN bridge domain, the IRB interface in the bridge domain remains up and causing the issue of IRB subnet remaining being advertised in L3 routing which in turn attracts all L3 VPN traffic for the subnet. [PR994909](#)
- MX960/480/240 fantray red alarm temperature changed from 75C to 80C. [PR995225](#)
- In the dual Routing Engines scenario with NSR configuration, backup peer proxy thread is hogging CPU for more than 1 second if there are multiple updates (>5000) going from master Routing Engine to backup Routing Engine. This leads to FPC socket disconnections. The traffic forwarding might be affected. [PR996720](#)
- On MX104 router with SONET/SDH OC3/STM1 (Multi-Rate) MIC. In rare condition, if the MIC is plugged out from MX104, the Packet Forwarding Engine might crash, the traffic forwarding will be affected. These MICs as below belong to SONET/SDH OC3/STM1 (Multi-Rate) MIC: * MIC-3D-8OC3OC12-4OC48 * MIC-3D-4OC3OC12-1OC48 * MIC-3D-8CHOC3-4CHOC12 * MIC-3D-4CHOC3-2CHOC12 * MIC-3D-8DS3-E3 * MIC-3D-8CHDS3-E3-B * MIC-3D-1OC192-XFP [PR997821](#)
- If encapsulation type is "ppp" for the SONET interface on IQE PIC, sometimes the MTU change might not work. [PR1001880](#)
- If the connection with an OpenFlow controller goes down then comes back up repeatedly, an OpenFlow interface on a QFX5100 switch might send an OFPT_ERROR packet with an XID ID 0 but no data to explain why the error packet was sent. [PR1003538](#)

- On TXP with GRES enabled, when performing graceful switchover on all chassis (include line-card chassis (LCC)) from master Routing Engine to backup Routing Engine, minimal IPv4 traffic loss around 0.04% to 0.05% will be observed on aggregated 100GE PIC on FPC type 4. [PR1014420](#)
- If the service option configured on aggregated Multiservices (AMS) interface is different from its member interface, conflict would happen which might cause some serious issue. After this fix, service-options configuration (which includes timeouts/sessios-limit etc.) should only be configured on all members interfaces when configure AMS bundle. [PR1014898](#)
- Under corner cases, if there are multiple back-to-back Virtual Chassis port (VCP) related CLI commands, Network Processing Card (NPC) core may be observed and FPC hosting the VC ports might reboot. [PR1017901](#)
- Enabling sampling on an ms- interface is not supported configuration, if 'forwarding-opions sampling sample-once' is subsequently deactivated the FPC may reboot. [PR1021946](#)
- MQCHIP(0) mqchip_get_q_forwarded_stats() invalid q_sys 0 q_num messages continuously show in logs. It will cause two GE or XGE interfaces to not forward traffic. [PR1021951](#)
- On Offline/Online cycle of a 40GE QSFP card, a 40GE Interface port's Physical Link might remain down. Few events which will result into the Offline/Online cycle of a 40GE QSFP card are router reboot, FPC reboot, or chassis-control restart or 40GE Card offline request followed by a 40GE Card online request. [PR1026088](#)
- The host MPC might continuously crash when trying to online a faulty MS-MIC after discovering the hardware failure. [PR1026310](#)
- Configuring a routing policy with the "no-route-localize" option to ensure that the routes matching a specified filter are installed on the FIB-remote Packet Forwarding Engines, after removing the routing policy and changing the next-hop for the routes, the previously installed routes using "no-route-localize" policy might not get removed from some Packet Forwarding Engines. Then the Packet Forwarding Engines will not forward received packets to the FIB-local Packet Forwarding Engines to perform full IP table lookup but using the staled routes instead. [PR1027106](#)
- On MPC5E line card, if a firewall filter with large-scale terms (more than 1300 etc.) is attached to an interface, traffic drop might be seen. [PR1027516](#)
- In a rare case, rdd core is reported under /usr/sbin/rdd as soon as applying the group and commit is performed. [PR1029810](#)
- On MX Series platform with MS-MPC card, after performing switchover from master RE0 to backup RE1, 2 internal ARP entries for Routing Engine address (128.0.0.1) on MS-MPC PICs pointing to two eth interfaces connect to CB0 and CB1 separately might be wrongly created. Then if pull out RE0/CB0, the MS-PIC would still select the eth interface connects to CB0, which results in loss of connectivity because that path is not available anymore. [PR1030119](#)

- In VMX, the speed of 10GE interface was not being displayed correctly in "show interface" command. This PR fix allows one to configure the speed on the interface. [PR1031286](#)
- With an unrecognized or unsupported Control Board (CB), mismatch link speed might be seen between fabric and FPCs, which results in FPCs CRC/destination errors and fabric planes offline. Second issue is in a race condition, Fabric Manager (FM) might process the stale destination disable event but the error is cleared indeed, it will result in the unnecessary FPC offline and not allowing Fabric Hardening action to trigger and recover. [PR1031561](#)
- This issue only affects OC-48 MICs. If an SFP is inserted into an OC-48 MIC port that has been disabled the SFP will not show up in a >show chassis hardware command. The issue is fixed with a patch. Contact JTAC to find out which version is best for you. [PR1031851](#)
- The Software Development Kit (SDK) Service process (ssd), which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS, might crash after Routing Engine switchover and following reboot of both Routing Engines. Since the ssd acts as the broker daemon for Applications connecting to Juniper distributed application framework (JDAF) services, the applications will lose JDAF connectivity when ssd restarts. [PR1031860](#)
- With VPLS BGP control word configured, intermittent packet loss might be seen in one direction on VPLS circuit due to the control-word not being programmed at Packet Forwarding Engine after member DPC reboot. The problem can happen on below conditions: 1. LSI interface exists across two or more physical interfaces. 2. Those physical interfaces located in different FPCs. 3. Those physical interfaces consist of equal-cost paths. So, LSI will not be flapped with one member FPC down. 4. Flap the member DPC where one of physical interfaces situated. [PR1031863](#)
- In rare cases, the AUTHD daemon may crash and cause a corruption of subscriber dynamic profiles. In-use profiles may be incorrectly marked as not in use. Any subscribers that reference that profile are forced to remain in Terminating state, until the router is rebooted. Daemon restarts and GRES switches are ineffective in working around this situation. [PR1032548](#)
- On the virtual MX (vMX) platform with high rate data (in this case, 50Mbps). In rare condition, the IPv6 Neighbor Discovery Protocol (NDP) packet might lose, the traffic forwarding will be affected. [PR1035852](#)
- If an IFL is used as the qualified-next-hop (which implies the IFL has unnumbered-address configured), and there are changes in the IFL filter configuration, then the static route might disappear from routing table. To make it reappear, need to delete it from the configuration and add it back. [PR1035598](#)
- Sometimes AE vlan ifl output byte counters are shown as large value and it is a generic issue. [PR1036813](#)
- Using jnxoptIfOTNPMFECIntervalTable and jnxOpticsPMIntervalTable it is not possible to walk these tables from the middle before this fix. [PR1039030](#)
- In a subscriber scenario with auto-sensed VLAN configured, after scaled subscribers (in this case, 16K subscribers) login/logout for several times, the subscriber management

process might stuck and not able to restart due to a Session Database (SDB) deadlock issue. [PR1041094](#)

High Availability (HA) and Resiliency

- Configuring the maximum segment size (MSS) for the TCP connection for BGP neighbors, if "mtu-discovery" and "path-mtu-discovery" knobs are removed, the default MSS value of 512 will be used instead. This is not an expect behavior. [PR835220](#)

Interfaces and Chassis

- Refer to the following topology. If we set interface ge-1/0/8 disable, interface xe-2/0/0 and xe-2/1/0 become down status because "asynchronous-notification" feature. However after 3 or 4 seconds, ether OAM detects link-fault status changed to good. And then, interface xe-2/0/0 and xe-2/1/0 change link status from down to up. The condition is the following. 1. Configure MPLS circuit with ether CCC. 2. Configure "asynchronous-notification" on CE facing interface in both PEs. 3. Configure ether OAM to one of PE, CE pair. 4. Use DPC 10 giga-interface on DTU. * This behavior did not occur with MPC and DPC 1 giga-interface. << topology >>

```
***** local
link remote link DPC 10ge | xe-2/0/0 V ge-1/0/6 ge-1/0/8 [ CE ]-----[ PE ]-----[
PE ]-----[ CE ] xe-2/1/0 ge-1/0/7 ge-1/0/9 (DTU) <-----> <-----> <----->
ether CCC MPLS ether CCC asynchronous-notification asynchronous-notification
<-----> ether OAM *CE:MX240 PE:MX240
*****
```

[PR973840](#)

- IS-IS Adjacency may flap after unified ISSU. This behavior is being further analyzed and is planned to be fixed in further releases. [PR1015895](#)
- On 10GE interface on MIC (e.g. 3D 4x 10GE XFP and 3D 2x 10GE XFP MIC), when "link-down" event under "optics-options alarm low-light-alarm" is configured and the "hold-time down" timer is set greater than 0, the status of the interface will remain up, even when the light power exceed low alarm threshold and traffic being interrupted. [PR1018076](#)
- With vrf-table-label configured on the routing-instances, when a FPC with Enhanced IQ (IQE) PIC is sharing the same Forwarding Engine Board (FEB) with another FPC, and the FEB has two core-facing interfaces configured with the family mpls on aforementioned FPCs separately, the Label-Switched Interface (LSI) interfaces might be removed incorrectly on the working FPC when the other FPC with IQE PIC is set to offline. [PR1027034](#)
- if DPCE 20x 1GE + 2x 10GE X card is present in the chassis, BFD sessions over AE interfaces may not be distributed [PR1032604](#)
- With heartbeat connection for an MX Series Virtual Chassis (MX-VC) enabled, if the heartbeat connection detects that the Virtual Chassis master router (VC-M) is still operating and able to respond during a split caused by a failure of all the Virtual Chassis port (VCP) interfaces, the Virtual Chassis backup router (VC-B) should go offline after the heartbeat timeout period expires. But VC-B retains VC backup role and never go offline although its FPCs went into PRESENT state. In addition to fix the deviation from the expected functionality, the output of CLI command "show virtual-chassis heartbeat

[detail]" is enhanced to more clearly indicate the successful detection of the peer MX-VC member chassis over the heartbeat connection when the chassis loses all VCP adjacency links. A unique "detected" state is provided when MX-VC splits and last heartbeat pulse response is successfully received. [PR1034096](#)

- Some duplicate entries are reported in jnx-chas-defines.mib. This patch removes the duplicate entries to fix the issue. [PR1036026](#)
- FRR switching time is much higher than 50ms (e.g. might be 400-900 ms) when protected links are located on MX Series Gigabit Ethernet enhanced and hardened MICs (i.e. MIC model name end with -E or -EH, currently, the supported MICs are MIC-3D-20GE-SFP-E and MIC-3D-20GE-SFP-EH). [PR1038999](#)
- For Ethernet OAM/CFM, if Maintenance-association (MA) ICC format name of length less than 13 characters (13 byte) is used, deactivate/activate of 'protocol oam' may cause CFM operation failures. 'Cross-connect CCM received' alarm will be seen. There can be other triggers also. ITU CARRIER CODE format uses fix length size for MA NAME (13 octets). Junos OS creates and maintains actual size configured by user. However the length it maintains is 13 octets. For lower size MA name the value accessed is not deterministic. It would work fine if the subsequent memory is initialized to zero. Else would declare cross connect error as the accessed MA name will be different compared to remote end. [PR1041482](#)

Layer 2 Features

- If a customer is using SNMP and performs an snmpwalk on the dhcp binding table, not all of the entries may be displayed. This fix resolves that issue so that bindings for all ip addresses are displayed. [PR1033158](#)

MPLS

- Error "tag_icmp_route:failed to find a chain composite ahead of fwd nh" might be observed when doing traceroute. [PR999034](#)
- When the size of a Routing Engine generated packet going over an MPLS LSP is larger than MTU (i.e. MTU minus its header size) of an underlying interface, and the extra bytes leading to IP-fragmentation is as small as <8 bytes, then that small-fragment will be dropped by kernel and lead to packet drop with kernel message "tag_attach_labels(): m_pullup() failed". For example - If SNMP Response with specific size fall into above mentioned condition then small fragment will be dropped by kernel and eventually the SNMP response will fail. [PR1011548](#)
- TED link information of protocol from highest credibility level is used irrespective of the level at which CSPF is computing. i.e., cspf-metric in "show mpls lsp extensive" would have the sum of te-metric of IGP with highest credibility at each hop in ERO. This has been corrected and the cspf-metric will be sum of te-metric of current credibility at each hop. [PR1021593](#)
- When configuring point-to-multipoint (P2MP) Label Distribution Protocol (LDP) label-switched paths (LSPs), the labels will never be freed even they are no longer

needed. This could lead to the MPLS label exhaustion eventually. To clear the state, the rpd process will restart with core dumps. [PR1032061](#)

- When a LDP enabled router receives a LDP label mapping message which includes an unknown TLVs with unknown and forward bit set, the unknown TLV will be re-advertised along with the LDP message to upstream LSR. However, due to merge issue, Junos appends these unknown TLVs multiple times during construction of label mapping message and will has a unknown TLV(0x0000) with length 0 among the appended unknown TLVs, thereby causing the LDP session with the peer that receives this message flapping. [PR1037917](#)

Network Management and Monitoring

- jnxcpic 380 and jnxcpic 381 definitions has been added in the "mib-jnx-chas-define" file from 14.1R4 release. [PR1036706](#)

Platform and Infrastructure

- With inline jflow enabled, when the flow is exported once and got reinserted, if the low 12 bits of the packet counter are zero (0x000), the packetDeltaCount counter might be incorrect in inline jflow records. There is no traffic impact but may impact billing. [PR886222](#)
- When apply-groups are used in the configuration, the expansion of interfaces <*> apply-groups will be done against all interfaces during the configuration validation process, even if the apply-group is configured only under a specific interface stanza. [PR967233](#)
- BFD session within default routing-instance are not coming up once inline-services pic is configured and fixed class-of-service forwarding-class is assigned. BFD session operating in no-delegate-processing are not affected. [PR999647](#)
- On TX Matrix Plus routers or TX Matrix Plus routers with 3D SIBs, all the incoming interfaces on an FPC are deactivated when none of the fabric planes are functional. By default, the interfaces remain activated. You can enable the deactivation of interfaces by using the fpc-restart configuration statement at the edit chassis fabric degraded hierarchy level. [PR1008726](#)
- On MX Series router with MPCs or MICs, with igmp-snooping enabled and a multicast route with integrated routing and bridging (IRB) as a downstream interface, a multicast composite nexthop is created with a list of L3 and corresponding L2 nexthops. In a rare corner case, the corresponding L2 nexthop to the L3 IRB nexthop is a DISCARD nexthop and will cause the FPC to crash. [PR1026124](#)
- On MX Series router with MPCs or MICs, when the packets are queued for several seconds due to interface congestion and get aged, the ICHIP might not able to detect those aged packets and thus fail to drain the queue out, which results in the FPC showing CRC errors and going into wedge condition. [PR1028769](#)
- MX Series router with MPCs and MICs might crash when trying to install the composite next-hop used for the next-hop-group configuration related to port mirroring of traffic over IRB to an LSI attached to VPLS instance for a remote host. [PR1029070](#)

- For BFD over aggregated Ethernet (AE) interfaces on MX Series routers with MS-MPC that have configured the enhanced-ip option, the BFD distribution to Packet Forwarding Engine for AE interface might not happen. [PR1031916](#)
- This check (log message) has been added as part an enhancement in the JNH error report. For FC accounting on AE interface, ingress FC accounting is enabled on AE interface nexthops and egress FC accounting is enabled on AE child member next hops. While fetching stats for AE, both member child IFL and AE IFL stats are fetched and added for result. If ingress FC accounting is enabled on AE IFL, while fetching statistics for child member links this error trace is coming because of this newly added JNH error trace. The fix is to put a check to not call for child member FC statistics when egress accounting is not enabled on AE bundle. [PR1032952](#)
- On MX Series router with MPC, when there is a congested Packet Forwarding Engine destination, the non-congested Packet Forwarding Engine destinations might experience an unexpected packet drop. [PR1033071](#)
- When the 'enhanced-hash-key services-loadbalancing' feature is used by MX Series router with MPCs or MICs, load balancing of flows across multiple service PICs via the source-address across does not work when internal BGP (IBGP) is used to steer traffic to the inside service-interface. For example the operator will see on the stateful firewall that the same source-address has flows across multiple service interfaces. [PR1034770](#)
- sa-multicast load sharing method under [chassis <> fpc <> pic <> forwarding-mode] is not working on 100GE interface on TRIO FPC. [PR1035180](#)
- Presence of /8 prefix in two terms results in incorrect filter processing and unexpected behavior. [PR1042889](#)
- In a scaled subscriber management environment, the output of CLI command "show subscribers" and its sub flavors might print more pages and has to be terminated by "Ctrl+c" or "q". But this was not closing the back end Session Database (SDB) connection properly. Over a period of time, this will cause inconsistency and the subscriber management infrastructure daemon (smid) fails to register and no new subscribers could connect. [PR1045820](#)
- On T4000 and FPC Type 5-3D or TXP-3D platforms , BFD sessions operating in 100msec interval with default multiplier of 3 might randomly flap after the enhancements implemented via PR967013. BFD sessions with lower intervals of 100msec or higher intervals are not exposed. The internal FPC thread, monitoring the High Speed Fabric links had a run time of longer then 100msec. [PR1047229](#)
- By default, after 16x 10GE MPC boards come up about 75% of queues were allocated to support rich queuing with MQ chip. Such allocation causes MQ driver software module to poll stats. Polling stats causes this rise in CPU usage. [PR1048947](#)

Routing Policy and Firewall Filters

- In the BGP environment, if operator "!" exists in the regex for as-path, the commit operation fails. [PR1040719](#)

Routing Protocols

- Under following combination of events: * graceful-restart is enabled and * bidirectional PIM is enabled and * rpd is restarted, and * multicast traffic for bidir rp group hits the box. Pim creates the discard route and this traffic is pruned. [PR1019560](#)
- When BGP is doing path selection with default behavior, soft-asserts requests are introduced. If BGP route flap a lot, it need to do path selection frequently, because of which a great deal soft-asserts might be produced which will cause unnecessary high CPU and some service issues, such as SNMP can not respond and even rpd core. [PR1030272](#)
- When policy LFA is being used and backup path selection is first based on the root-metric criteria, the root-metric should be taken from the link metric connecting source to the backup neighbor (the one-hop neighbor or a remote router such as an RSVP backup LSP tail-end router), but it is now taken from the shortest-path-first (SPF) metric from source to backup neighbor if root-metric highest is configured. In some topologies, if the two metrics are different, IS-IS might select incorrect backup next-hop. [PR1031408](#)
- In distributed BFD (which is enabled by default), if the CLIENT session (for example BGP) flaps due to any reason, the multi-hop BFD session that comes Up after the flap would not be delegated to FPC. [PR1032617](#)
- When "clear bfd session" is issued immediately(before the Poll - Final sequence is completed) post config check-in for interval change from higher to lower minimum-interval value, BFD sessions don't revert to lower interval. [PR1033231](#)
- Issue in populating isisRouterTable values. Some entries are not filled correctly. This does not block/affect the functionality of IS-IS or other components. [PR1040234](#)

Services Applications

- The show CLI command "service nat pool detail" always displays the Port range starting from 1024 even when privileged ports are used. [PR1019783](#)
- The session-limit-per-prefix feature for the MX Series DS-Lite server does not take Software flow into account when calculating the flow limit. [PR1023439](#)
- In Network Address Translation (NAT) scenario with Endpoint-Independent Mapping (EIM) configured on service PIC, when a new ICMP session is created which matches an existing EIM mapping, the service PIC might crash. [PR1028142](#)
- For T Series or M320 router containing Dynamic Flow Capture (DFC) PIC (either a Monitoring Services III PIC or Multiservices 400 PIC), there are two issues for DFC feature. The first one is the value of "timeout-remaining" for some filters installed on the DFC pic are too huge. The second issue is for some filters, there won't be any flows to which they are attached when forwarding traffic to the content-destination during random DTCP ADDs. [PR1029004](#)

- When NAT has multiple terms that refer to the same NAT Pool, the command 'show snmp mib walk jnxSvcsMibRoot ascii' always print out jnxNatPoolTransHits for the count of jnxNatRuleTransHits in the first term. [PR1035635](#)
- The cause of the KMD crash is not known. This is not due to SA(Security Associations) memory corruption. The code looks that SA is getting freed without clearing the table entry. [PR1036023](#)
- In the context of DS-Lite softwire scenario, the MS-PIC/MS-DPC might crash in rare occasions when the Dual-Stack Lite (DS Lite) softwire concentrator receiving a high volume of outer IPv6 fragmented packets. [PR1044143](#)

VPNs

- Problem Description The problem is that MSDP is periodically polling PIM for S,G's to determine if the S,G is still active. This check helps MSDP determine if the source is active and therefore the SA still be sent. There is a possibility that PIM will return that the S,G is no longer active which causes MSDP to remove the MSDP state and notify MVPN to remove the Type 5. One of the checks PIM makes is to determine if it is the local RP for the S,G. During a re-configuration period where any commit is done, PIM re-evaluates whether it is a local RP. It waits until all the configuration is read and all the interfaces have come up before making this determination. The local rp state is cleared out early in this RP re-evaluation process, however, which allows for a window of time where the local RP state was cleared out but it has not yet been re-evaluated. During this window PIM may believe it is not the local rp and return FALSE to MSDP for the given source. If MSDP makes the call into PIM during this window after a configuration change(commit), then it is possible that the Source Active(Type 5) state will be removed. Fix The fix will be to clear out the local rp state right before it is re-evaluated ie after it reads configuration for all interfaces; to not allow any time gap where it could be inconsistent. [PR1015155](#)
- On MX-VC platform, if with scaled number of MVPN routes, after adding a new interface in the MVPN instance or changing the traceoptions related configuration, the CPU on Routing Engine might experience a high utilization for about 10min. [PR1027596](#)
- Selective provider tunnel might flap few seconds after Routing Engine switchover, type 3 & 4 routes also refreshed, traffic fall to inclusive provider tunnel for a while [PR1049352](#)

Resolved Issues: 14.1R3

- [Class of Service \(CoS\) on page 155](#)
- [Forwarding and Sampling on page 155](#)
- [General Routing on page 156](#)
- [High Availability \(HA\) and Resiliency on page 159](#)
- [Infrastructure on page 159](#)
- [Interfaces and Chassis on page 160](#)
- [J-Web on page 161](#)
- [Layer 2 Features on page 161](#)
- [Layer 2 Ethernet Services on page 161](#)

- [MPLS on page 162](#)
- [Multicast on page 163](#)
- [Network Management and Monitoring on page 163](#)
- [Platform and Infrastructure on page 163](#)
- [Routing Protocols on page 166](#)
- [Services Applications on page 169](#)
- [Subscriber Access Management on page 170](#)
- [User Interface and Configuration on page 170](#)
- [VPNs on page 170](#)

Class of Service (CoS)

- SNMP get-request for OID `jnxCosIngressQstatTxedBytes` (ingress queue) might return the value of `jnxCosQstatTxedBytes` (egress queue). But SNMP walk works fine since it uses get-next-request. [PR1011641](#)
- Sometimes MX Series routers might respond with "no such instance" of the second OID when two CoS OIDs in the single SNMP packet. [PR1015342](#)

Forwarding and Sampling

- On the 32-bit Junos OS, when a very big burst-size-limit value (2147492676 and above) is configured in the ingress interface policer, the kernel may drop Routing Engine destined traffic. [PR1010008](#)
- Deactivating Inline Jflow configuration doesn't make memory release normally. [PR1013320](#)
- When an ARP policer is applied to an interface, it appears commented out in the configuration with the following message: "invalid path element 'disable_arp_policer'". [PR1014598](#)
- On MX Series routers with MPCs/MICs, if layer 2 hierarchical policer is configured, upon committing it, the firewall daemon (dfwd) might crash. [PR1015190](#)
- Remote vtep interface is not created despite sending traffic from inter segment, after vtep router reboots or chassisd is restarted. It causes dropping packets. [PR1016446](#)
- When a TRIO specific filter is configured on an interface located on a DPC, the filter is not being installed and no warning message is logged on the message log file. [PR1022836](#)
- Adding "fast-lookup-filter" knob to a firewall filter using one or more terms with "next-term" action could cause dfwc crash during commit (commit check phase). Hence because of this bug, this disallows use of "fast-lookup-filter" feature on firewall filters with terms using "next-term". [PR1029761](#)

General Routing

- On TXP/TXP-3D platform, a bad I2C device on SFC Switch Interface Board (SIB) might cause Switch Processor Mezzanine Board (SPMB) to crash and all SIBs unable to online. [PR846679](#)
- A few particular sequence of member failures in an AMS with HA-enabled and with NAPT-44 configured, can cause sessions to reset after a GRES (or SPD restart). [PR910802](#)
- In this scenario the CPCD (captive-portal-content-delivery) is configured for HTTP-REDIRECT for Subscriber Management clients using MS-DPC. When services sessions start to redirect the HTTP traffic, the memory-usage consistently increments for MSPMAND on the multi-service PIC. The memory limit then might cause packets loss. [PR954079](#)
- 1) Due to a previous fix chassisd on the protocol master Routing Engine and the protocol backup Routing Engine connect to the main snmpd on the protocol master using the following methods. a) Chassisd on the protocol master Routing Engine connects using a local socket since snmpd is running locally. b) Chassisd on the protocol backup Routing Engine connects using a TNP socket since snmpd is not local. 2) However this fix changed the way the other daemons connect to snmpd. All important daemons run on the protocol master and should connect to snmpd using a local socket. However the fix changed it so that all daemons that ran on the protocol master (other than chassisd) tried to connect using the TNP socket. SNMPD does not accept these connections. As a fix, in an MX-VC, we made sure that chassisd connects to all processes which run on the protocol master using internal socket while the chassisd process on the protocol backup and protocol linecard connect connect using TNP socket. [PR986009](#)
- In the dual Routing Engines scenario, in rare condition, while executing GRES and deleting interfaces at the same time, it is possible that a nexthop delete message is not sent to rpd process, causing rpd to keep a nexthop index (NHID) that kernel has already deleted. Later when kernel allocates this NHID for next new nexthop and sends it to rpd process, rpd process might crash due to duplicate NHID. [PR987102](#)
- In the VPLS environment with control-word configuration, when the "control-word" is changed to "no-control-word", there are 5 minutes service outage. [PR987216](#)
- Mirroring of CCC traffic would be broken for a very small duration when Routing Engine switchover is happening. Post switch-over, CCC mirroring would work as expected. [PR987593](#)
- In 6PE scenario, when PE router is sending IPv6 TCP traffic to MPLS core, in rare occasions, the kernel might crash and reboot with a vmcore file created. [PR988418](#)
- OpenFlow v1.0 running on an MX Series router does not respond reliably to interface up or down events within a specified time interval. Per a fix implemented in Junos OS Release 13.3R3.6, OpenFlow v1.0 running on an MX Series router responds reliably to interface up or down events if the echo interval timeout is set to 11 seconds or more. [PR989308](#)

- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX Series routers with DPC. If "no-local-switching" is present in the bridge domain, then the IGMP-snooping is not functioning and client can't see the multicast traffic. [PR989755](#)
- Configure the global and interface limit to allow for the maximum configuration of the macs in the overlay network. [PR992084](#)
- Commit error need to be reported when using unsupported NAPT44 nat-options max-sessions-per-subscriber config with MS-MIC/MS-MPC. [PR993320](#)
- MX960/480/240 fantray red alarm temp changed from 75C to 80C. [PR995225](#)
- On T4000 router with type5 FPC. After FPC rebooting, if chassisd process does not get FPC ready/FPC online ACK message from FPC in 360 seconds, the FPC might reset again. [PR998075](#)
- On M Series, MX Series, and T Series routers with Network Address Port Translation (NAPT) configuration. When the router receives the packet whose value of protocol field in the IPv4 header is 61, the router erroneously does NAPT44 translation. In the correct situation, the packet should not be translated and forwarded. [PR999265](#)
- By default, the syslog utility exports 800,000 logs per second to a remote syslog server. You can modify the number of syslogs to be sent by including the message-rate-limit statement at the [edit interfaces interface-name services-options syslog] hierarchy level to suit your deployment needs. The rate at which syslog messages can be sent to the Routing Engine is 10,000 logs per second. [PR1001201](#)
- On MX240/MX480/MX960 routers running as Precision Time Protocol (PTP) master when interconnect with MX104 series routers running as slave, the PTP clocking state might stuck in "INITIALIZING" for the first created PTP port and not be aligned to clocking state. There is another issue that when issue command "show ptp clock", wrong "slot" number might be seen on MX104 slave. [PR1001282](#)
- "Syslog generated for session-open will have nat port information only if it is different from the original source port". [PR1001912](#)
- If issue the command "show services nat mappings endpoint-independent" or "show services nat mappings address-pooling-paired" or "show services sessions" and kill it immediately when using EIM/APP feature with too many EIM/APP entries present in the system, lots of ipc message reply failure messages may be seen in the syslog. [PR1002683](#)
- Multi-Services PIC could crash and restart on receiving a stray SIGQUIT signal due to it not handling the signal. [PR1004195](#)
- During ISSU early stage, when Mm is arming packages on other three Routing Engines, Mm will not copy config/ssh files to Bm, and Bm will not fork mgd to copy the files. This should not be a problem. During ISSU, when backup chassis switchover done, the original Bm (new Bs) will copy the files from the original Bs(new Bm, now it has the latest config files from Mm). So the original Bm could always get the latest config/ssh files. [PR1004766](#)
- When several PICs are set up as an aggregated multi-services (AMS) doing load-balancing, if one PIC of the AMS bundle get offline and then get online, a 30 to 40 seconds momentary traffic loss might be seen. [PR1005665](#)

- The l2cpd process might crash if there are multiple unknown type, length, and value (TLV) elements included in received LLDP PDUs. [PR1007223](#)
- MS-DPC memory leak on system service set when HTTP Redirect attempts to process none-HTTP traffic with HTTP ports (80/8080/443). [PR1008332](#)
- Ingress queuing is not supported on MPC5 (With Q-MPC) when Optical Transport Network (OTN) is enabled. Enabling ingress queuing with OTN would lead to line card crash. [PR1008569](#)
- With more than 8 service-sets configured, when using SNMP mibwalk for service-set (object "jnxSpSvcSetTable") info, the mspmand process (which manages the Multi-Services PIC) might crash. [PR1009138](#)
- Add "protocol evpn" configuration under an existing virtual-switch routing-instance might cause EVPN neighborship unable to establish. [PR1009339](#)
- Whenever a FPC goes down suddenly due to hardware failure, the data traffic in transit towards this FPC from the other FPCs could be stuck in the fabric queue thereby triggering fabric drops due to the lack of buffers to transmit the data to active destination FPCs. [PR1009777](#)
- PIC state mgmt is not available in 14.1R2. [PR1013480](#)
- Unknown unicast flood seen with interface flap after router reboot & with static mac,no-mac-learning,interface-mac-limit config for a virtual-switch. [PR1014222](#)
- The routing protocol daemon (rpd) might crash continuously with core-dumps upon adding a sub-interface with "disable" configuration to a MC-LAG interface. [PR1014300](#)
- On TXP with GRES enabled, when performing graceful switchover on all chassis (include line-card chassis (LCC)) from master Routing Engine to backup Routing Engine, minimal IPv4 traffic loss around 0.04% to 0.05% will be observed on aggregated 100GE PIC on FPC type 4. [PR1014420](#)
- A new global knob is added at the top level CLI "set forwarding-options port-mirroring [no-preserve-ingress-tag]" By default the system behavior would remain as it is today where ingress mirrored copy would contain VLAN content exactly as what came in wire over ingress. However, if this knob is configured, if any VLAN modification happens to packet as part of its datapath processing, that would get retained in the ingress mirrored copy that is, we will not restore VLAN to what came in ingress on wire. [PR1015149](#)
- Hash-key command is no longer treated as a hidden command and considered invalid input in 12.3 for small footprint routers (these platforms don't support the hash-key feature), this could cause configuration failure during a software upgrade if hash-key command is configured prior to the upgrade. This PR reverses the above change and allows hash-key command to be ignored on unsupported platforms: show configuration forwarding-options ### Warning: configuration block ignored: unsupported platform (mx80) ## hash-key { family inet { layer-3; } } [PR1016339](#)
- MAC accounting support was added for 40G and 100G interfaces on MPC3 and MPC4 cards. [PR1017595](#)

- With Enhanced IP network service mode configured, traffic might fail to be sent out over the inline LSQ bundle interface. [PR1018887](#)
- Traffic destined to the Broadcast or Network address of a NAT pool using the address prefix setting for the MS-MPC card causes a traffic loop that spikes the CPU. [PR1019354](#)
- No performance or functional impact. Can be safely ignored. "Ignore the PTP message (2) as this MPC doesn't support EEC" should be moved from notice to debug level. [PR1020161](#)
- When source address is configured under ms interface, and the service-set has syslog host as local the FPC slot is printed as -ve. [PR1020854](#)
- For M320 or T series FPCs (M320 non-E3 FPC and T Series non-FPC5) with queuing PIC, if the configured total buffer size temporal values exceed the supported maximum scheduler buffer size for the PIC (e.g. For PD-5-10XGE-SFPP PIC, the maximum temporal buffer size that can be configured for a scheduler is 40,000 microseconds), the default scheduler [95,0,0,5] is applied instead of the default chassis scheduler [25,25,25,25], which might result in the packet drops on Q1 and Q2. [PR1027547](#)
- On MX Series routers with MS-MPC card, after performing switchover from master Routing Engine0 to backup Routing Engine1, two internal ARP entries for Routing Engine address (128.0.0.1) on MS-MPC PICs pointing to two eth interfaces connect to CB0 and CB1 separately might be incorrectly created. Then if pull out RE0/CB0, the MS-PIC would still select the eth interface connects to CB0, which results in loss of connectivity because that path is not available anymore. [PR1030119](#)
- PCS statistics counter is now displayed interfaces in this command: **monitor interface <intf>**. [PR1030819](#)

High Availability (HA) and Resiliency

- This issue occurs in rare condition. In the dual Routing Engines scenario, doing interface flap after Routing Engine switchover. If this action is repeated many times, the stale indirect nexthop entry might be seen in kernel, which leads to traffic blackhole. [PR987959](#)
- In an MX Series Virtual Chassis configuration, a unified in-service software upgrade (ISSU) from Junos OS Release 14.1 or 14.1R2 to Junos OS Release 14.2 fails with traffic loss. [PR1014295](#)

Infrastructure

- SNMP socket sequence error log. [PR986613](#)
- A reboot is needed if "chassis network services enhanced-ip" is configured on MX Series 3D routers or on T4000 Routers with type 5 FPCs. Without the reboot, performing ISSU might cause new master Routing Engine to crash and go to db> prompt. [PR1013262](#)

Interfaces and Chassis

- When the GE port is configured with WAN PHY mode, a "Zero length TLV" message might be reported from the port. [PR673937](#)
- Error message CHASSISD_IPC_DAEMON_WRITE_ERROR is seen in the messages log when there is a Routing Engine mastership change (system reboot, Routing Engine reboot, GRES switchover CLI command), which causes a restart of alarmd, which breaks the IPC connection between alarmd and chassisd. Chassisd does not detect that the IPC connection has been broken, because it is busy processing the mastership change, and then tries to send alarm information to alarmd during this time. So it encounters a write error (broken pipe) and logs the message. [PR908822](#)
- If dynamic VLAN subscriber interface is over a physical interface (IFD), and there are active subscribers over the interface, when deactivate the dynamic VLAN related configuration under the IFD and add the IFD to an aggregated Ethernet (AE) interface which has LACP enabled, the Routing Engine might crash and get rebooted. [PR931028](#)
- In the dynamic-profile environment with preferred-source-address configuration. If subscribers stuck in terminating state, it is impossible to commit changes. [PR978156](#)
- In the PPPoE environment, when the subscriber logs in successfully but profile activate fails, due to code processing error, the address entry is not deleted in the authd's DAP pool. So when the subscriber tries to log in again, it connects fails. [PR995543](#)
- In the demux interfaces over aggregated Ethernet (AE) environment with targeted-distribution configuration. The index of AE interface is confusion when the index is more than 100. It copy only 4 bytes from interface name. (e.g. If bind demux interface to ae110, it will be bound to ae11 at the same time). The traffic forwarding might be affected. [PR998906](#)
- IGMP joins do not work for PPP subscribers that are using MLPPP and LNS. [PR1001214](#)
- In L2 circuit, with async notification configured on a client facing interface goes down, then on the remote PE the corresponding CE interface shows up in show interface terse output while in log snmp reports interface down. [PR1001547](#)
- Fabric Blackholing logic recovery for certain cases will be done with different action (Phase 1/2/3) based on the problem. [PR1009502](#)
- As current Junos OS multichassis link aggregation groups (MC-LAGs) design, the ARP entry will not sync when learning ARP via ARP request but not Gratuitous ARP/ARP reply. In some specific scenarios (e.g. a host changes its MAC address without sending a Gratuitous ARP), traffic loss might occur. [PR1009591](#)
- Here is the expected behavior for CFM CCM: 1. UP MEP CFM session a. If there is a manually configured ieee-802.1 classifier attached to the interface, then forwarding class of the CCM injected should match the respective classifier. b. If there is an interface in which CFM is configured has no ieee-802.1 based 1p classified, then the forwarding class of the CCM will take as configured in "host-outbound-traffic". c. In case if there is no "host-outbound-classifier" present then packets will be treated as network control (Q3). 2. Down MEP CFM session a. forwarding class of the CCM will always depend on the FC classified based on "host-outbound-traffic". If it is not configured, then it will always take Q3. [PR1010929](#)

- IS-IS Adjacency may flap after ISSU. This behavior is being further analyzed and is planned to be fixed in further releases. [PR1015895](#)
- VRRP daemon (vrrpd) memory leak might be observed in "show system processes extensive" when VRRP is set with routing-instance and then change any configuration. [PR1022400](#)
- **set forwarding-options enhanced-hash-key symmetric** command will not get applied on MX104 Packet Forwarding Engine. [PR1028931](#)
- In addition to fixing the reported deviation from the expected functionality, the show virtual-chassis heartbeat [detail] command output is enhanced to more clearly indicate the successful detection of the peer MX-VC member chassis over the heartbeat connection when the chassis loses all VCP adjacency links. [PR1034096](#)

J-Web

- An insufficient validation vulnerability in J-Web can allow an authenticated user to execute arbitrary commands. This may allow a user with low privilege (such as read only access) to get complete administrative access. This scope of this vulnerability is limited to only those users with valid, authenticated login credentials. Refer to JSA10560 for more information. [PR826518](#)

Layer 2 Features

- In BGP signaled VPLS/VPWS scenario, rpd process memory leak might occur when a groups with wildcard configuration is applied to the routing instance. [PR987727](#)
- After configuration change or convergence events, kernel may report ifl_index_alloc failures for LSI interfaces and cause KRT queue ENOMEM issue, eventually preventing new IFL's from being added to the system. This condition always recovers on its own once convergence is completed. [PR997015](#)
- In a mixed VPLS instance where both ldp and bgp flavors are present, any cli change in that instance will result in RPD crash. [PR1025885](#)

Layer 2 Ethernet Services

- In MX Series Virtual Chassis (MXVC) scenario with LACP configuration. In rare condition, after VC-M chassis power down, the LACP state getting stuck in ATTACHED state, all traffic carried over these affected access LAGs are blackholed. [PR959041](#)
- When "system no-redirect" is configured, I2 descriptor destination MAC address gets overwritten and causes "DA rejects" on next-hop router. [PR989323](#)
- In the Ethernet ring protection switching (ERPS) environment, once graceful Routing Engine switchover happens on the ring protection links (RPLs) owner node, there will be a ~30s Ring automatic protection switching (R-APS) message storm in the ring, which in turn cause some VPLS instance flapping. [PR1004066](#)
- On MX Series routers with DHCP service enabled, issuing CLI command "show dhcp-security binding" might result in jdhcpd process crash. [PR1007577](#)

- With scaled VPLS instances configured, aggressively flapping the interfaces belonging to the VPLS instances might result in l2cpd process memory leak. When l2cpd reaches its max memory limit, l2cpd process crash will be seen. [PR1009952](#)
- Commit is failing on backup Routing Engine with ethernet-ring configuration under "protocol protection-group" hierarchy. user@R1# commit synchronize re0: configuration check succeeds re1: [edit protocols protection-group] 'ethernet-ring vkm1' L2CPD : INVALID node-id configured for pg vkm1 error: configuration check-out failed re0: error: remote commit-configuration failed on re1 Node id value is not available to backup Routing Engine, when it is not configured. As in such case it is derived from chassis mac and on backup Routing Engine chassis-mac remains as 00:00:00:00:00:00. Fix: validation check for node-id value will not be done on backup Routing Engine. [PR1011441](#)
- If "maintain-subscriber" knob is enabled on the router, DHCPv6 server/relay might be unable to process any packet if deactivate and then activate the routing instance, which means the subscribers can not get the IPv6 addresses. Note, even with the fix, the results of this scenario is also expected if with "maintain-subscriber" knob enabled. Consider using the workaround to avoid this issue. [PR1018131](#)
- After FPC restart, bridge domain (BD) implicit filters for Ethernet ring protection switching (ERPS) might get reprogrammed with wrong logical interface (ifl) index, which cause ERPS to not work correctly. [PR1021795](#)

MPLS

- Although NSR does not support MPLSOAMD and it does not run on backup Routing Engine, backup RPD is attempting to do task_connect to MPLSOAMD. This behavior causes periodical message popping up on backup Routing Engine. Feb 21 15:14:13.306 2014 mx480-re1 rpd[2840]: task_connect: task MPLSOAMD I/O./var/run/mplsoamd_control addr /var/run/mplsoamd_control: No such file or directory. [PR938284](#)
- In the MPLS environment with no-cspf & strict ERO configuration. In race condition, if a PATH message with routing loop error is received before standby Routing Engine has resolved the correct PATH message with no loop, some of LSP are not replicated on standby Routing Engine. If Routing Engine switchover occurs, the forwarding traffic might be affected. [PR986714](#)
- BGP may reevaluate all its routes if a protocol mpls stanza is configured, but an egress-protection stanza is not. On a scaled setup, this can keep RPD CPU high for several minutes after each commit. [PR1000550](#)
- Interoperability issue between Junos OS and IOS-XRv, the virtual IOS-XR. It is related to the max_pdu TLV in LDP. IOS-XRv only supports max_pdu 1000 or below. On the other hand, Junos OS only supports max_pdu 1200 or above. So LDP session never comes up successfully. There are fixes in both vendors. On the Junos OS side, max_pdu 1000 is accepted after the fix and the session comes up. [PR1007096](#)
- When the size of a Routing Engine generated packet going over an MPLS LSP is larger than MTU (i.e. MTU minus its header size) of an underlying interface, and the extra bytes leading to IP-fragmentation is as small as <8 bytes, then that small-fragment will be dropped by kernel and lead to packet drop with kernel message "tag_attach_labels(): m_pullup() failed". For example - If SNMP Response with specific

size falls into above mentioned condition, then small fragment will be dropped by kernel and eventually the SNMP response will fail. [PR1011548](#)

- The entropy label value allocated at times falls in the reserved mpls label range(0-15). The label value is calculated based on load balancing information and hence only certain mpls flows may encounter this issue. [PR1014263](#)
- In MPLS scenario with TX/TXP router acting as the transit node, performing MPLS LSP ping or traceroute from ingress node might cause kernel crash on the transit node due to improper timer initialization between SCC and LCC chassis. [PR1020021](#)

Multicast

- In multicast environment, if GRES is performed immediately after a routing-instance being deleted, the krt (kernel routing table) queue might get stuck after adding back the routing-instances which were deleted. [PR1001122](#)

Network Management and Monitoring

- Due to a communication error between the master agent (snmpd process) and the subagent (mib2d process), it might cause a failure to register some MIBs. For example: There is no output while running below commands: user@hostname> show snmp mib walk ifTable When user tries polling the device for ifAlias. The following messages might be seen: user@hostname:~\$ snmpwalk -v 2c -c snmp@exp X.X.X.X ifAlias IF-MIB::ifAlias = No Such Object available on this agent at this OID This means that there's no OID registered. [PR978535](#)
- The Packet Forwarding Engine local protocol statistics are 32-bit counters. If there is a rollover (typical candidates are ARP/LACP), those counters start from zero. mib2d will add all counters again if one of the Packet Forwarding Engine statistics traffic counter is less than the previous collected counter, causing the multiplication affect. [PR986712](#)
- Alarm management daemon runs on master and backup Routing Engines on dual Routing Engine systems. There is a 80 megabyte alarm.db file that is copied over from master Routing Engine to backup Routing Engine when the alarm-management daemon has come up on both the routing engines. The basic issue is that alarm-management daemon is trying to copy the alarm.db file over and over again in an infinite loop on the system, causing CPU utilization shooting up after every 20 seconds or so. [PR988969](#)
- The snmpd process becomes unresponsive for ~30 minutes after performing GRES when SNMPv3 notify type is configured to be "inform". [PR1021943](#)

Platform and Infrastructure

- MPLS traceroute causes "rttable-mismatch" syslog messages. [PR960493](#)
- When apply-groups are used in the configuration, the expansion of interfaces <*> apply-groups will be done against all interfaces during the configuration validation process, even if the apply-group is configured only under a specific interface stanza. [PR967233](#)
- This is a corner case. On MX Series routers with MPCs/MICs, some stale unilists nexthops are present. If an interface is down for more than ARP timeout interval, the broken

selectors in unilist nexthops, and then traffic will be blackholed when the interface is up again. [PR980052](#)

- Have BFD session between one router supporting inline-BFD (Trio and Junos OS 13.3 or higher) and the other which doesn't support inline-BFD (any version and non-Trio, or Trio and Junos OS less than release 13.3). When the "failure detection time" is less than 50ms, the BFD session might flap. [PR982258](#)
- GRES doesn't clear system login to original master-only fxp0 addresses causing stale login sessions. [PR991029](#)
- On MX2020/MX2010 we might see sporadic FO request time-out error reported under heavy system traffic load. This would mean the request returning into a grant took longer than +/-30usec. The packet will still get forwarded through the fabric hence no operational impact. [May 6 18:56:59.174 LOG: Err] MQCHIP(2) FO Request time-out error [May 6 19:33:47.555 LOG: Info] CMTFPC: Fabric request time out pfe 2 plane 6 pg 0, trying recovery. [PR991274](#)
- When we uninstall an SDK package, the config related to that package is still left out in the config file. After this if commit sync is issued, though commit is successful, it leads to a commitd core. Before the un-installation of SDK package, the config statement [set jnx-ifinfo traceoptions flag all] should also get deleted from the config which is relevant to the package being deleted. [PR992486](#)
- On MX Series routers with MPCs/MICs or T4000 router with type5. When the firewall filter under the [forwarding-options] hierarchy within a bridge domain is removed, it might result in lookup error and frame drop might be observed. [PR999083](#)
- In the IRB interface environment with "destination-class-usage" configuration. If the bridge domain ID is the same as Destination Class Usage (DCU) ID (bridge domain ID and DCU ID are generated by system), the firewall filter might match wrong packets, the packet forwarding would be affected. [PR999649](#)
- On M7i, or M10i equipped with Enhanced Compact Forwarding Engine Board (CFEB-E). When a MPLS LSP flaps, the CFEB-E is unable to recover 8 bytes of JTREE memory per event. [PR1000385](#)
- When receiving traffic coming on MPC and going out on DPC, the MAC entry on a Packet Forwarding Engine might not be up-to-date and the frames targeted to a known MAC address will be flooded across the bridge domain. [PR1003525](#)
- With NSR enabled, when activating a BGP session in a routing instance, and the interface route is imported into the main routing instance, the TCP receive window might decrement until it hits 0 after receiving incoming BGP traffic arrives from the main routing instance. [PR1003576](#)
- On MX Series routers with MPCs/MICs, routers in the same VRRP instance might both claim to be VRRP master after performing unified in-service software upgrade (ISSU) upgrading to specific Junos OS versions. [PR1004471](#)
- On MX-VC platform, if there is a dark window larger than 5s (TCP timeout timer) from Packet Forwarding Engine to Routing Engine during unified ISSU for some reasons, some VC members might get unexpectedly rebooted. [PR1005309](#)

- In PPPoE over ATM subscriber management environment with active subscribers is present, when issue the "show arp" command, an ARP core file is generated. [PR1006306](#)
- The non-first IP fragments containing UDP payload may be mistakenly interpreted as PTP packets if the following conditions are met: - the byte at the offset 9 in the IP packet contains 0x11 (decimal 17) - UDP payload - the two bytes at the offset 22 in the IP packet contain the value 0x01 0x3f (decimal 319; byte 22=0x01 and byte 23=0x3f) - PTP protocol The mis-identification of the packet as PTP will trigger the corruption of the fragment payload. [PR1006718](#)
- Micro BFD sessions are used to monitor the status of individual LAG member links. When micro BFD configurations are added after the LAG bundle configuration in separate commit, the micro BFD sessions for all the member links might remain in "Down" state. [PR1006809](#)
- Memory allocated in reference to the BFD session was not getting freed up. This resulted in memory leak and the memory exhaustion triggered crash. [PR1007432](#)
- For MX Series routers with MPCs/MICs or Chassis-based line cards, if there are more than 8K PPPoE subscribers with SRL (ratelimit) on a same FPC, new subscribers might not be able to connect any more due to filter memory threshold. [PR1009232](#)
- If rate-limit has been configured in scheduler for MX-VC VCP ports, ISSU might fail. [PR1009590](#)
- MPLS traffic going through the ingress pre-classifier logic may not determine mpls payload correctly classifying mpls packet into control queue versus non-control queue and expose possible packet re-order. [PR1010604](#)
- Issue: This change addresses missing NULL check in a trace message which was resulting in Packet Forwarding Engine crash. The error path involves scenario where ifbd is not yet created for an IFL. This is possible under certain IPC reordering scenarios. The Packet Forwarding Engine should not crash by differencing a NULL pointer in this case. [PR1014090](#)
- The fix was committed for this PR# but it also needs DDOS configuration additional to this fix and it is as below: 1) check the "show ddos-protection protocols statistics terse" 2) For each of the Control plane protocols on the system like ospf/vrrp/pvstp, it is recommended to configure 2X of the rate as give below example along with increasing DDOS rate for virtual-chassis control. Example, ##### set system ddos-protection protocols virtual-chassis control-high bandwidth 20000 set system ddos-protection protocols virtual-chassis control-high burst 20000 set system ddos-protection protocols ospf aggregate bandwidth 1000 set system ddos-protection protocols ospf aggregate burst 1000 set system ddos-protection protocols vrrp aggregate bandwidth 100 set system ddos-protection protocols vrrp aggregate burst 100 [PR1017640](#)
- On MX Series routers with MPCs/MICs, when there are next-hop changes, the "heap 0" memory of the FPC may experience memory leakage which will eventually causes memory exhaustion. [PR1019794](#)
- For MX Series routers with inline Network Address Translation (NAT) service, when using "source-prefix" or "destination-prefix" in a NAT translation rule, a pool is implicitly created, appending "_jinpool_" with the rule name and term name with a form :

jinpool{rule_name}_{term_name}. The name might be cropped due to the maximum length limitation (64 characters). If that happens, both pools might get the same name and result in the indeterminate behavior (statistic issue, drop or incorrect translation). [PR1020033](#)

- problem scenario: The error logs "CHASSISD_FCHIP_CONFIG_MD_ERROR" will appear during FPC normal boot up time and also during FPC restart time for each plane and for each LMNR FPC. Problem statement: This Error logs "CHASSISD_FCHIP_CONFIG_MD_ERROR" are observed only in M320 chassis containing FPCs based on LMNR chipsets. Due to this error log, the rate limit for the fabric port connecting the Packet Forwarding Engine 1 will be set to the default values. [PR1020551](#)
- When receiving traffic coming on MPC and going out on DPC, an Ethernet frame with known DMAC will be flooded to the whole bridge domain after flapping the link which the given MAC is learned for more than 32 times. [PR1026879](#)
- When a layer 2 frame entered the VPLS end point on the label-switched interface (LSI) with VLAN tagged, the frame is wrongly interpreted and treated as no VLAN frame. So the VLAN tag will not be popped although the outbound interface has a pop configuration. [PR1027513](#)
- In normal case, network-service enhanced-ip would make BFD over AE distributed to Packet Forwarding Engine (control plane independent). However due to this software issue, it would remain running on Routing Engine. [PR1031916](#)

Routing Protocols

- High CPU utilization is observed by routing process when high number (around 1000) of Rosen based MVRP configuration is committed in one-shot. It will take more than 1 hour for CPU usage by routing process to come to normal condition. [PR947732](#)
- Performing CLI command "clear multicast bandwidth-admission interface <int>" on 64-bit Junos OS results the rpd process crash. The command should be used without the interface qualifier on the impacted releases. [PR949680](#)
- In a scaled setup a restart routing or NSR switchover can result in duplicate msdp entries. [PR977841](#)
- On a platform with an IGMP configuration in which two receivers are joined to the same (S,G) and IGMP immediate-leave is configured, when one of the receivers sends a leave message for the (S,G), the other receiver might not receive traffic for 1-2 minutes. [PR979936](#)
- In the P2MP environment with OSPF adjacency established. One router's time is set to earlier date than another router. OSPF adjacency might not come up when one router goes down and comes up. [PR991540](#)
- Bringing up DFWD based BFD sessions at scale causes a churn in DFW as a result of which the FPC CPU usage remains at 100% for a prolonged timespan. [PR992990](#)
- When all the following conditions are met, if the knob "path-selection always-compare-med" is configured, the rpd process might crash. - routing-instance (VR, VRF) with no BGP configuration - rib-group in default instance with routing-instance.inet.0 as secondary-rib - rib-group applied to BGP in default instance

- BGP routes from master tables (inet.0) leaked to the routing-instance table (routing-instance.inet.0) [PR995586](#)
- When IS-IS is configured for traffic engineer (TE), after remove family mpls from the interface and remove the specific interface from [edit protocols rsvp] and [edit protocols mpls] hierarchy level, corresponding link is not removed from the TED as expected. [PR1003159](#)
- When there are more than 65535 "flow-spec" routes existing in the routing table, the rpd process might crash because it exceeds the current maximum supportable scaling numbers (Current scaling numbers are in the range of 10K~16K). [PR1004575](#)
- When having ECMP routes and multiple levels of route/next-hop recursion, a particular sequence of routes churn may result in rpd process crash and traffic outage. [PR1006523](#)
- Abnormal ip6 route-calculation behavior can be seen when ospf3-te-shortcut is configured. [PR1006951](#)
- When the same PIM RP address is learned in multiple VRFs, with NSR configured, RPD on the backup Routing Engine may crash due to memory corruption by the PIM module. [PR1008578](#)
- When deleting a routing-instance or making changes to the routing-instance, the deletion of the routing-instance to kernel might come before the deletion of the IFLs in the routing-instance, resulting in rpd crash. This is a timing issue, hard to reproduce. [PR1009426](#)
- During unified in-service software upgrade (ISSU), when a Bidirectional Forwarding Detection (BFD) session negotiation is happening, if the session is configured with 10 seconds or higher interval, BFD session would flap. [PR1010161](#)
- Misconfiguring BGP route validation session to the router itself might lead to rpd process crash. [PR1010216](#)
- When inet.3/inet6.3 is not enabled, BGP group uses inet6.0 table to advertise the routes for both inet6 unicast and inet6 labeled-unicast families. When BGP family is changed, BGP sessions re-establish. When BGP starts to advertise routes to the peer, BGP expects to see route label however if the old inet6 unicast routes are still present (not completely cleaned), then rpd process crashes. The fix is to separate bgp group for inet6 unicast with inet6 labeled-unicast with same rib. The old peers are cleaned up in the old group and new peers are established in new group. Thus, new peer establishment is not delayed by the cleanup of the old peer. [PR1011034](#)
- Issue: IsisRouterTable MIB issues, when we do "show snmp mib walk isisRouterHostName/isisRouterTable" we were not getting exact hostname as it is in "show isis hostname" so the actual implementation was not as per RFC-4444, because it was showing only the hostnames of the devices which are immediate neighbors of Dut. Fix: added level info to get sysid_entry per each level correctly and filled data(isisRouterTable) correctly. [PR1011208](#)
- In scaled BFD scenarios, BFD ISSU poll negotiation will fail causing the BFD session to flap during ISSU. [PR1012859](#)
- Under certain sequence of events RPD can assert after a RPD_RV_SESSIONDOWN event. [PR1013583](#)

- With multicast discard route present, if a RP router has no pd- interface, it might not generate (S,G) join to upstream when receiving MSDP source active (SA) message. [PR1014145](#)
- For 64-bit Junos OS, the route protocols process (rpd) might crash and dump core during IBGP route churn when using IBGP multipath and multiple levels of IBGP route/next-hop recursion. [PR1014827](#)
- This PR is implementing traceoptions debug enhancements to detect route-record corruption events. The route-record traceoptions debug will be enabled as follows:
----- user@router> edit Entering configuration mode [edit]
user@router# set routing-options traceoptions flag route-record [edit] user@router#
commit ----- [PR1015820](#)
- The OpenSSL project released a security advisory on 2014-08-06 that contained nine security issues. The following four issues affect Junos OS: CVE-2014-5139: Crash with SRP ciphersuite in Server Hello message CVE-2014-3509: Race condition in ssl_parse_serverhello_tlsext CVE-2014-3511: OpenSSL TLS protocol downgrade attack CVE-2014-3512: SRP buffer overrun See JSA10649 for more information. [PR1016458](#)
- When receiving open message with any capability after the "add-path" capability from BGP peer, the session will be bounced. [PR1016736](#)
- With BGP multipath configured, if a BGP route's multiple protocol nexthops are resolved to different types of IGP routes with a same metric, high rpd process utilization might be observed due to the BGP multipath task. [PR1017372](#)
- The snmp trap generated when an ipv6 BFD session goes up/down does not contain the ipv6 bfd session address. [PR1018122](#)
- The Junos OS implementation of RFC3107 uses unspecified label (0x000000) when sending route with label withdrawn message. This means Junos OS sends 0x000000 instead of 0x800000 for label withdrawn, which is inconsistent with RFC 3107. [PR1018434](#)
- Under following combination of events: * graceful-restart is enabled and * bidirectional PIM is enabled and * rpd is restarted, and * multicast traffic for bidir rp group hits the box. Pim creates the discard route and this traffic is pruned. [PR1019560](#)
- Establish two BFD sessions between two routers, one is single-hop BFD for directly connected interface and the other is multi-hop MPLS OAM BFD. If configuring the MPLS OAM on the same interface with single-hop BFD, when bringing down MPLS OAM from the ingress, it might result in the OAM BFD session deleted on ingress but it still receiving OAM BFD down packet from egress. Since there is no session matching this BFD packet, it does a normal look up and brings down the single-hop BFD session which is on the same interface. [PR1021287](#)
- If auto-export feature is enabled together with rib-groups configuration option, the rpd process might crash. [PR1028522](#)

Services Applications

- If a destination-prefix or source-prefix is used like below example, the Network Address Translation (NAT) rule and term names will be used to generate an internal jpool with a form : `_jpool_{rule_name}_{term_name}`. If the generated jpool name exceeds 64 characters in length, it will get truncated. If the truncated jpool name get overlapped with other generated jpool name it will lead to an inconsistent pool usage. `user@router# show services nat rule A_RULE_NAME_WHICH_IS_LONG_12345 { ... term A_TERM_ALSO_WITH_LONG_NAME_1 { from { source-address { 10.20.20.1/32; } } then { translated { source-prefix 10.10.10.1/32; <--- translation-type { source static; } } } } term A_TERM_ALSO_WITH_LONG_NAME_2 { from { source-address { 10.20.20.22/32; } } then { translated { source-prefix 10.10.10.2/32; <--- translation-type { source static; } } } } } First jpool = _jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_LONG_NAME_1 > 64 characters. Second jpool = _jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_LONG_NAME_2 > 64 characters. The resulted jpool "_jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_" will be used wrongly in both terms. PR973465`
- On MX240/MX480/MX960 routers with MS-DPC with "deterministic-port-block-allocation block-size" configuration. In rare condition, when the "block-size" is set to a larger value (in this case, block-size=16128), the Services PIC might crash. [PR994107](#)
- In the NAT environment, a same pool is used in several terms of a nat-rule. If any pool parameter is modified, the configuration change is ignored. [PR994200](#)
- The redundant services PIC (rsp-) interfaces or redundant Multiservices (rms-) interfaces configured with "hot-standby" mode might flap upon committing any configuration change (will happen for even an unrelated interface description change). [PR1000591](#)
- The following messages are being logged at ERR not DEBUG severity: `mspd[3618]: mspd: No member config mspd[3618]: mspd: Building package info` This PR sets the correct severity. [PR1003640](#)
- When removing a basic-nat44 translation term, there is a chance the prefix that was used for this translation will become wedged. Any attempt to reuse this prefix for dynamic-nat44 or napt-44 will be such that no address/port allocation will succeed. [PR1008214](#)
- Software tunnel count management is inconsistent and incorrect, thus the output of "show service software statistics" might be incorrect. [PR1015365](#)
- Configured Port Control Protocol (PCP) lifetime is ignored and NAT pool mapping-timeout is used instead for pinholes existence on the CGNAT public interface. [PR1017155](#)

- L2TP LNS dropped all tunnels/sessions after a commit. [PR1020420](#)
- With Real Time Streaming Protocol (RTSP) Application Layer Gateway (ALG) enabled, the PIC might crash in case the Transport header in status reply from the media server is bigger than 240 bytes. [PR1027977](#)

Subscriber Access Management

- MIB entries for jnxUserAAAAccessPoolRoutingInstance may not appear after deleting and re-adding an assignment pool under a routing instance. [PR998967](#)

User Interface and Configuration

- CST: chassis core generated while applying group config on chassis > FPC. [PR936150](#)

VPNs

- In the Rosen MVPN environment, the RP-PE is an assert loser, another PE is sending traffic over the data-mdt. If a new receiver PE with higher rate comes up, because internal workflow processes wrong, the receiver PE might reset data-mdt. This leads to traffic loss. [PR999760](#)
- Serving site B is not receiving all the traffic from serving site A when traffic is reduced from the exceeded cmcast limit. [PR1001861](#)
- In the 12.3 release after issuing a "request pim multicast-tunnel rebalance" command the software may place the default encapsulation and decapsulation devices for a Rosen MVPN on different tunnel devices. [PR1011074](#)

Resolved Issues: 14.1R2

- [Class of Service \(CoS\) on page 170](#)
- [Forwarding and Sampling on page 171](#)
- [General Routing on page 171](#)
- [Interfaces and Chassis on page 173](#)
- [Layer 2 Ethernet Services on page 174](#)
- [MPLS on page 175](#)
- [Platform and Infrastructure on page 175](#)
- [Routing Protocols on page 177](#)
- [Services Applications on page 177](#)
- [Software Installation and Upgrade on page 178](#)
- [User Interface and Configuration on page 178](#)
- [VPNs on page 178](#)

Class of Service (CoS)

- Manually setting max-queues-per-interface to 4 on PB-4OC3-IOC12-SON-SFP doesn't work. The ports will still work with 8 queue while displaying 4 queue from CLI output. [PR981253](#)

- On MX Series routers with MPC and MPCE and other type of linecard, DPCE, when the Default Frame Relay DE Loss Priority Map is configured and committed, all FPCs are getting restarted with core files. [PR990911](#)

Forwarding and Sampling

- Less impact on customer environment, it is just a ease of debugging issue. [PR950553](#)
- DPC crashed after deactivate/activate [routing-instances TPIX bridge-domains IX bridge-options. [PR983640](#)

General Routing

- When nonstop active routing (NSR) is configured and the memory utilization of rpd process on the backup Routing Engine is high (1.4G or above), the rpd crash on the backup Routing Engine may bounce the BGP sessions on the master Routing Engine. [PR942981](#)
- There is a regression issue in Release 14.1 and later for single chassis with NSR and the MXVC environment. RPD might crash during GRES or membership switchover due to asynchronized routing table between Routing Engines. [PR950767](#)
- Under particular scenarios, commit action might lead the Context-Identifier to be ignored when OSPF protocol refreshes its database. Then the PE router will stop advertising this Context-Identifier. [PR954033](#)
- "show interfaces et-x/y/z extensive" will display MRU now. MRU can be configured at "set interfaces et-x/y/z gigether-options mru" If MRU is not configured then it is defaulted to MTU + 8. MRU displayed from the CLI does not include the CRC [PR958162](#)
- On MX Series Virtual Chassis (MX-VC), if multiple VCP ports are configured between MPC5E cards, traffic might not be load balanced over the VCP ports. Besides, packets might get lost due to VC ingress and egress next-hop caches getting out of synchronization. [PR960803](#)
- Although receiving the flow specification (flowspec) routes with packet-length, icmp-code, or icmp-type matching rules from a BGP peer properly, the local firewall filter in the Packet Forwarding Engines might not include these matching rules. [PR968125](#)
- On an MX VC-Mm Routing Engine switch, the last flap time and associated error counters for the VCP interfaces sometimes get reset. The last flap time can be incorrectly reported as 'Never', for those VCP that have previously flapped. [PR971995](#)
- tnping member1-RE0 from member0-RE0 fails because of a replication panic at "rnh_index_alloc: nhindex 624 could not be allocated err=12" [PR977445](#)
- Changing service-set configuration continuously during scaled traffic conditions may result in mspmand process crash and a core file generated. [PR978032](#)
- Juniper Distributed Application Framework (JDAF) serviceability feature enables CLI based inspection of various JDAF service counters. [PR978640](#)
- On T Series router with FIB Localization enabled, if reboot the Routing Engine while scaled traffic running, the FIB-remote FPC might crash. [PR979098](#)

- In rare condition, when PPPoE subscribers log in with large amounts of configuration data, the subscriber management infrastructure daemon (smid) and authentication service process (authd) might crash, and no new subscribers could connect to the router. [PR980646](#)
- In scenario of NG-MVPN with P2MP LSP as provider tunnel, Kernel Routing Table (KRT) might get stuck after making changes for MVPN, then traffic loss will be seen. Besides, rpd process might crash while trying to generate a live core file. [PR982959](#)
- With a firewall policer configured on more than 256 IFFs (interface address family) of a PIC, then offline and online the PIC might cause the FPC to crash. [PR983999](#)
- OpenSSL library in Junos OS was patched to resolve CVE-2010-5298. [PR984416](#)
- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX with DPC. In a race condition, the Dense Port Concentrator (DPC) may crash when ifls get added to an ifl-set while that same ifl-set get deactivated/deleted in class-of-service. For example:

```
# set interfaces interface-set interface_set_JTAC_ge-3/0/0 interface ge-3/0/0 unit 100 # deactivate class-of-service interfaces interface-set interface_set_JTAC_ge-3/0/0 # commit or (quick commit of following changes) # set interfaces interface-set interface_set_JTAC_ge-3/0/0 interface ge-3/0/0 # commit # deactivate class-of-service interfaces interface-set interface_set_JTAC_ge-3/0/0 # commit
```

[PR985974](#)
- When the logical interface's (IFL) MTU is changed (set interfaces et-x/y/z unit 0 family inet mtu xx), the static route goes to dead state and never recovers on its own. [PR989021](#)
- During large scale MVPN routes churn events, some core-facing IGP protocols (like OSPF or LDP) might flap or experience a long convergence time. [PR989787](#)
- When the interface-mac-limit on vtep interfaces is reached, any new OVSDB MACs advertised from the same remote VTEP are never getting added to the bridge mac-table. [PR992084](#)
- Group VPN member registration in MX Series router will not succeed if the same interface is used for both data traffic and server-member communication. This limitation will apply if a group VPN service-set is applied on the interface in which the member is communicating with the Group key server. [PR993001](#)
- The fabric performance of MPC1, MPC2, or 16xE MPC in 'increased-bandwidth' mode on an MX960 populated with SCBE's will be less compared to redundant mode due to XF1 ASIC scheduling bugs. [PR993787](#)
- On 10X10GE SFPP, when an interface configured for CCC and asynchronous-notification, and it is told to turn off its later, its laser flaps on and off for some period of time. [PR996277](#)
- The PIC memory gauge counters show up as 0 after a GRES switchover in the "show chassis pic fpc-slot X pic-slot Y" output. [PR1000111](#)

- Because of MCNH change from 13.3 to 14.1 and later, which used new FLOOD_MCNH to replace old MCNH_P2MP, while unified ISSU upgrading there would be a RPD crash. [PR1000494](#)
- When using AMS load-balancing if a PIC in the AMS bundled is offline for any reason and the operator on-lines the PIC, there is slight 30 to 40 second momentary traffic loss. [PR1005665](#)

Interfaces and Chassis

- Queue stats counters for AE interface will become invalid after deactivating ifl on the AE interface. [PR926617](#)
- Strange FRU Insertion trap[RE PCMCIA card 0] is generated when Routing Engine master-switching is done on box with RE-1800. [PR943767](#)
- When an ifl containing some vrrp group configuration is deleted, snmp walk on vrrp MIB may loop continuously. [PR957975](#)
- If there is an IRB interface configured for "family inet6" in a bridge-domain on an MX Series router, the Packet Forwarding Engine might not correctly update the next hop for an IPv6 route when the MAC address associated with the next hop moves from an AE interface to a non-AE interface. [PR958019](#)
- Temperature Top and Bottom are swapped in show chassis environments output for Type3/Type4 FPCs of T Series [PR975758](#)
- In the multilink frame relay (mlfr) environment with "disable-tx" configuration, when the differential delay exceeds the red limit, the transmission is disabled on the bundle link. When it is restored, the link should be added back. But in this case, the link stays in the disable state, and it is not rejoined to the bundle. [PR978855](#)
- With nonstop active routing (NSR) enabled, the VRRP tracking routes state on backup Routing Engine might not get synchronized when adding/deleting the tracking routes. [PR983608](#)
- When upgrading to Release 13.3R2, customer may see the following messages: Chassis control process: rtslib: ERROR kernel does not support all messages: expected 104 got 103,a reboot or software upgrade may be required Chassis control process: Chassis control process: rtslib: WARNING version mismatch for msg macsec (103): expected 99 got 191,a reboot or software upgrade may be required Chassis control process: Chassis control process: rtslib: ERROR kernel does not support all messages: expected 104 got 103,a reboot or software upgrade may be required Chassis control process: Chassis control process: rtslib: WARNING version mismatch for msg macsec (103): expected 99 got 191,a reboot or software upgrade may be required These messages are generated during validation of the new chassis management daemon against the old kernel, and are harmless. [PR983735](#)
- 1GbE SFP(EX-SFP-1FE-LX) output optical power is restored after reseating by manual removal/insert of SFP although the IF is disabled. [PR984192](#)
- SNMP OID VRRP-MIB::vrrpAssolpAddrRowStatus returns only one Ip address when the interface ifl has configured with two virtual-addressees under two vrrp-groups. [PR987992](#)

- Following messages could be seen on the router for the FPC slot which are even empty. These messages are cosmetic and could be ignored. `chassisd[1637]: %DAEMON-6: FPC 0 does not support Pic power off config cmd ignoring the config change`
`chassisd[1637]: %DAEMON-6: FPC 2 does not support Pic power off config cmd ignoring the config change.` [PR988987](#)
- CFMD may crash after configuration change of an interface in a logical system which is under OAM config for a l2vpn instance. [PR991122](#)
- In Ethernet OAM connectivity-fault-management, Junos OS default encodes MAID(MD name and MA name) in character format. Currently only 43 octets are supported in Junos OS for the MD + MA name. Junos OS needs to support a maximum length of 44 octets for MAID per the standards. [PR997834](#)
- On MX Series router with MPCs or MICs or T4000 router with type5 FPC, when the "Hardware-assisted-timestamping" is enabled, the MPC modules might crash with a core file generated. The core files could be seen by executing CLI command "show system core-dumps". [PR999392](#)

Layer 2 Ethernet Services

- In DHCPv6 subscriber environment, changing the c-tags (inner vlan) without clear the DHCPv6 clients first is not recommended, it might cause the subscriber to use the old inner vlan even after DHCPv6 RENEW process. [PR970451](#)
- When Cisco running in an old version of PVST+, it doesn't carry VLAN ID in the end of BPDU. So Juniper Networks equipment fails to respond Topology Change Notification ACK packet when interoperates with Cisco equipment. After the fix, Juniper Networks equipment will read the VLAN ID information from Ethernet header. [PR984563](#)
- Layer 2 Control Protocol process (l2cpd) is used to enable features such as Layer 2 protocol tunneling or nonstop bridging. If a router receives a Link Layer Discovery Protocol (LLDP) packets with multiple management address TLV, memory leak might occur which resulting in l2cpd process crash. [PR986716](#)
- `jnxLacpTimeout` trap may show negative values and incorrect values for `jnxLacpifIndex` and `jnxLacpAggregateifIndex`. [PR994725](#)
- In race condition, when FPC gets rebooted or reset, link(s) from this FPC which are part of aggregate-ethernet bundle would remain in LACP "Detached" state indefinitely.
`user@node> show lacp interfaces ae102` Aggregated interface: ae102 LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity xe-2/0/0 Actor No Yes No No No Yes Fast Active xe-2/0/0 Partner No Yes No No No Yes Fast Passive xe-2/0/1 Actor No No Yes Yes Yes Yes Fast Active xe-2/0/1 Partner No No Yes Yes Yes Yes Fast Active LACP protocol: Receive State Transmit State Mux State xe-2/0/0 Defaulted Fast periodic Detached xe-2/0/1 Current Fast periodic Collecting distributing
`user@node> show interfaces xe-2/0/0 terse` Interface Admin Link Proto Local Remote xe-2/0/0 up up xe-2/0/0.0 up up aenet --> ae102.0 xe-2/0/0.32767 up up aenet --> ae102.32767
This issue would be seen when associated aggregate-ethernet bundle is configured for vlan-tagging. To clear this condition, the affected interface should be deactivated and activated using cli commands. ===== [edit] `user@node# deactivate interfaces xe-2/0/0` [edit] `user@node# commit` [edit] `user@node# activate interfaces xe-2/0/0` [edit] `user@node# commit` ===== [PR998246](#)

MPLS

- snmpwalk/snmpgetnext or "show snmp mib walk" fail when polling MPLSLSP OCTETS, MPLSLSP PACKETS, MPLSLSP INFO OCTETS or MPLSLSP INFO PACKETS. [PR981061](#)
- LSP metric modification leads to Constrained Shortest Path First (CSPF) computation and resignaling. It should update RSVP routes directly. [PR985099](#)
- In the MPLS environment with "egress-protection" configuration, there is a direct LDP session between primary PE and protector. One context-id is configured as primary PE's loopback address or any LDP enabled interface address. When delete the whole apply-group or delete the ldp policy from apply-group, the routing protocol daemon (rpd) might crash. [PR988775](#)
- In the virtual private LAN service (VPLS) environment with multihoming (FEC 129) configured, when the router receives the label request for the Forwarding Equivalency Class (FEC) 129, if there is no route for the specific FEC 129, the routing protocol daemon might crash. [PR992983](#)

Platform and Infrastructure

- When using OSPF/OSPFv3 with interface type point-to-point, it is possible for the OSPF session (using multicast traffic exclusively) to come up before next-hop resolution is done (ARP, or ND). In this case, transit traffic will be discarded, until resolution is done. When you have multiple links available, then the route will be balanced using a "unilist" next-hop. When one of the links in the "unilist" don't have Layer 2 resolution, these next-hops will actually drop traffic. The fix added by this PR will make unilist not contain forwarding and non-forwarding at the same time. When the next hop resolution will be done, then the link will be added to the unilist. [PR832974](#)
- The error message 'unlink(): failed to delete .perm file: No such file or directory' was logged when disconnecting from a Telnet session to the router. [PR876508](#)
- Starting with Junos OS Release 13.3 and later, the range of CLI screen-width is 40 through 1024 (in earlier Junos OS releases, the range is 0 through 1024). This PR restores the option of setting screen-width to 0 resulting in unlimited screen width. [PR936460](#)
- The Routing Engine and FPCs are connected with an internal Ethernet switch. In some rare case, the FPCs might receive a malformed packet from the Routing Engine (for example, packet gets corrupted somewhere on its way from the Routing Engine to FPC), then the toxic traffic might crash the FPC. [PR938578](#)
- MPC Type 2 3D might crash with CPU hog due to excessive link flaps causing the interrupts to go high. [PR938956](#)
- The issue might come when a non-template filter gets deleted (but does not get completely cleaned up) and the same filter index gets reassigned to a template filter. This could be considered as a timing issue given it comes with a very specific sequence of events only. [PR949975](#)
- On MX Series routers with MPCs or MICs, VPLS traffic might get blocked for about 5 minutes (timer of MAC address aged-out) after re-negotiating control-word. [PR973222](#)
- With NG-MVPN, multicast traffic might get duplicated and/or blackholed if a PE router, with active local receivers, is also a transit node and the P2MP LSP is branched down

over an aggregate interface with members on different Packet Forwarding Engines. [PR973938](#)

- On MX Series Virtual Chassis platforms with interface alias configured, this feature might not work as expected and cause interface flapping after commit. [PR981249](#)
- no-propagate-ttl doesn't work for L3VPV when PE is configured with l3vpn-composite-nexthop and its core interfaces are hosted on MPC based FPC. [PR985688](#)
- On MX Series routers with MPCs or MICs, when filter is applied on the interface with the action of "then next-interface", the packets that are forwarded by the firewall filter would be corrupted. [PR986555](#)
- Interface alias was not shown in the show commands when configured. Now interface alias will be shown (IF CONFIGURED) in show commands containing interface names. A |display no-interface-alias command adds the ability to show the actual interface name if it's needed. [PR988245](#)
- When services packet(interface-style) is diverted to different routing-instance using a firewall filter, route lookup of the services packet was matching a reject route which results in PPE thread timeout. [PR988553](#)
- TXP with Release 13.1R4 might not trigger autoheal after 65535 CRC error event on inter-chassis optical hsl2 link. Customer will need to do manual fabric plane reset to recover the faulty SIBs after the 65535 CRC error event. [PR988886](#)
- NPC core ../../src/pfe/ukern/cpu-ppc/ppc603e_panic.c:68 [PR989240](#)
- On logical systems, backup rpd of logical systems is not getting SIGHUP when the "commit fast-synchronize" statement at the [edit system] hierarchy level is enabled. It causes the issue "restarting backup rpd" of logical systems (as part of recovery mechanism). [PR990347](#)
- When two midplane link errors are present between F13 and F2 Sibs then CLOS rerouting logic does not work properly. This can introduce RODR packet drops and result in destination errors in the plane. [PR992677](#)
- "delete" or "deactivate" of apply-group defining the entire TACACS or RADIUS configuration configured under [edit system apply-group <>] does not take effect on commit. This could lead to TACACS or RADIUS based authentication to still continue working despite removal (delete/deactivate) of configuration. [PR992837](#)
- On MX240/MX480/MX960 routers with Multiservices DPC (MS-DPC), the MS-DPC might crash when the MPLS or VPLS with LAG Enhanced is configured. [PR993716](#)
- Packets dropped with IPv6 reject route are currently subjected to loopback IPv6 filter processing on MX Series routers with MPCs or MICs, as a result the packet dropped by a reject route may be seen from the "show firewall log". [PR994363](#)
- On MX Series routers with MPCs/MICs or T4000 router with type5 FPC, if the CoS scheduler is configured without transmit-rate while with buffer-size temporal, the Packet Forwarding Engine might not allocate buffer for the associated queue. The issue might lead to packets loss. [PR999029](#)

- Handle CHASSISD_FRU_UNSUPPORTED event with auto-image-upgrade.slax script. [PR1000476](#)
- MS PIC may reset after GRES in case of excessive resolve traffic. [PR1001620](#)

Routing Protocols

- In PIM-SM network with "bootstrap routing" RP selection mechanism used, it is observed that some bootstrap messages (BSMs) generation and forwarding behavior of Junos OS does not conform to RFC standard, specifically in the section 3.2 (Bootstrap message generation), 3.3 (Sending Candidate-RP-Advertisement Messages) and 3.4 (Creating the RP-Set at the BSR). [PR871678](#)
- In Protocol Independent Multicast (PIM) scenario, if interface get deleted before the (S,G) route is installed in the Routing Information Base (RIB), then this interface index might be re-used by kernel for another interface and thus cause routing protocol process (rpd) core. [PR913706](#)
- In certain rare circumstances, BGP NSR replication to the backup Routing Engine may not make forward progress. This was due to an issue where an internal buffer was not correctly cleared in rare circumstances when the backup Routing Engine was experiencing high CPU. [PR975012](#)
- On EX9200 switches or MX Series platforms with IGMP snooping enabled on an IRB interface, some transit TCP packets may be wrongly considered as IGMP packets, causing packets to be dropped. [PR979671](#)
- Due to some corner cases, certain commits could cause the input and/or output BGP policies to be reexamined causing an increase in rpd CPU utilization [PR979971](#)
- PPMD filter is not programmed properly which is resulting the Routing Engine absorbing BFD packets instead of the Packet Forwarding Engine. [PR985035](#)
- In Junos OS, by default the RIP protocol "send" option is set to Multicast RIPv2. When this "send" option is changed from "multicast" (active) to "none" (passive) or vice-versa, rpd core file might be seen on the router. [PR986444](#)
- in V4 RG, member site receives traffic from both serving sites for few sources upon withdraw/inject routes for 30 seconds. [PR988561](#)

Services Applications

- Clearing the stateful firewall subscriber analysis causes the active subscriber count to display a very huge number. The large number is seen because when a subscriber times out, the number of active subscribers is decremented. If it is set to zero using the clear command, then a decrement would give an incorrect result. There is no impact to the overall functionality. [PR939832](#)
- Jflowd core crashes because of the interface name mismatch between Jflowd config parsing and SRRB. Config parsing treats the interface as ms-*/*/*(without subunit) while SRRB reports ms-*/*/*. The fix is to treat interface name without any subunit as interface with subunit .0. [PR968922](#)
- If a PPPoE/PPP user disconnects in the access network without the LAC/LNS noticing it to tear down the connection (also the PPP keepalive hasn't detected yet), and a

second PPP request comes from the same subscriber on the L2TP tunnel (same or different LAC/tunnel), then a second route is added to the table having the next hop "service to unknown". [PR981488](#)

- The cflow export would cease due to memory exhaustion when flow-monitoring is enabled using Adaptive Services II PIC due to memory leak condition. While in this condition, user would see increments in "Packet dropped (no memory)" as below:
user@node> show services accounting errors Service Accounting interface: sp-3/0/0, Local interface index: 320 Service name: (default sampling) Interface state: Accounting Error information Packets dropped (no memory): 315805425, Packets dropped (not IP): 0 [PR982160](#)
- In H323 ALG with CGNAT scenario, the MS-PIC might crash when the ALG is deleting an H323 conversation due to the deleting port is outside of allocated NAT port-block range. [PR982780](#)
- On M Series, MX Series, and T Series routers (platforms) with Services PIC with dynamic-nat44 translation-type configured, when the flows are cleared the IP addresses in use are never freed. This issue is present in Junos OS Release 11.4R7 and all more recent releases without this fix. [PR986974](#)
- In large scale L2TP LNS environment. When the SNMP MIB JNX-L2TP-MIB is walked continuously, the memory of the L2TP daemon (jl2tpd) increases due to memory leak. [PR987678](#)
- In the Layer-2 Tunneling Protocol (L2TP) environment with "failover-within-preference" configuration. There are two L2TP network servers (LNSs) with different preference, one LNS is primary and another is backup. If the primary LNS is dead, the router doesn't try to create L2TP tunnel to the backup LNS. [PR990042](#)

Software Installation and Upgrade

- By upgrade-with-config, user can specify a configuration to be applied on upgrade, but the configuration file will not be loaded post upgrading. As a result, router will bring up with old configuration. [PR983291](#)

User Interface and Configuration

- When load large scale configuration, due to the ddl object not being freed properly after it's accessed, load configuration failed with error: Out of object identifiers. [PR985324](#)

VPNs

- Upon withdraw/inject bgp routes in the serving PEs for two different route-groups, member/regular sites receive traffic from both serving sites for 60 seconds. [PR973623](#)
- The S-PMSI tunnel might fail to be originated from ingress PE after flapping the routes to customer multicast source. [PR983410](#)
- In MVPN scenario, a multihomed ingress PE might fail to advertise type-4 after losing routes to local sources. [PR984946](#)

- In route-group scenario, source route is flapped on preferred serving site. After that the member site fails to originate type-4 even though it has type-5 and type-3 from non-preferred serving sites. [PR994687](#)
- Make the assert winner send the assert messages in a spaced way just as PIM Hellos and Joins are sent. With fix, the assert winner sends the assert message more often such that helps the other routers on the LAN to maintain state. For now, the robustness count is hard-coded as 3. This will later be enhanced by way of a CLI knob such that the robust count is configurable. [PR999019](#)

Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 14.1R7 documentation for the M Series, MX Series, and T Series.

- [Adaptive Services Interfaces Feature Guide on page 180](#)
- [Chassis-Level Feature Guide on page 181](#)
- [Ethernet Interfaces Feature Guide on page 181](#)
- [Firewall Filters Feature Guide for Routing Devices on page 182](#)
- [High Availability Feature Guide for Routing Devices on page 182](#)
- [Interchassis Redundancy Using Virtual Chassis Feature Guide for MX Series Routers on page 182](#)
- [Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide on page 184](#)
- [Junos OS Administration Library for Routing Devices on page 185](#)
- [Layer 2 Configuration Guide, Bridging, Address Learning, and Forwarding on page 185](#)
- [Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices on page 185](#)
- [MPLS Applications Feature Guide for Routing Devices on page 185](#)
- [Overview for Routing Devices on page 186](#)
- [OSPF Feature Guide on page 186](#)
- [Services Interfaces Configuration Guide on page 186](#)
- [Standards Reference on page 189](#)

- [Subscriber Management Network Access Feature Guide on page 189](#)
- [Subscriber Management Provisioning Guide on page 189](#)
- [System Log Messages Reference on page 189](#)
- [Traffic Sampling, Forwarding, and Monitoring Feature Guide for Routing Devices on page 190](#)
- [User Access and Authorization Feature Guide for Routing Devices on page 190](#)
- [VPLS Feature Guide for Routing Devices on page 191](#)
- [VPNs Library for Routing Devices on page 191](#)

Adaptive Services Interfaces Feature Guide

- The following items describe updates for aggregated Multiservices (AMS) interfaces information:
 - The description for the **rejoin-timeout** statement under the hierarchy **[edit interfaces interface-name load-balancing-options member-failure-options drop-member-traffic]** should be changed to the following:

Configure the time by when failed members (members in the **DISCARD** state) should rejoin the aggregated Multiservices (AMS) interface automatically. All members that do not rejoin by the configured time are moved to the **INACTIVE** state and the traffic meant for each of the members is dropped.

If multiple members fail around the same time, then they are held in the **DISCARD** state using a single timer. When the timer expires, all the failed members move to **INACTIVE** state at the same time.
 - The following information should be added to the “Aggregated Multiservices Interface” section in the “Understanding Aggregated Multiservices Interfaces” topic:

Member interfaces are identified as mams in the configuration. The chassisd process in routers that support AMS configuration creates a mams entry for every multiservices interface on the router.

When you configure **services-options** at the ams interface level, the options apply to all member interfaces (mams) for the ams interface.

The options also apply to service sets configured on ms- interfaces corresponding to the ams interface’s member interfaces. All settings are per PIC. For example, session-limit applies per member and not at an aggregate level.



NOTE: You cannot configure services-options at both the ams (aggregate) and member-interface level. If services-options is configured on ms-x/y/z, it also applies to service sets on mams-x/y/z.

When you want services-options settings to apply uniformly to all members, configure services-options at the ams interface level. If you need different settings for individual members (for example, because of a syslog configuration), configure services-options at the member-interface level.

- The **show interfaces load-balancing** command topic should include the following description for **Last change** in the table:

Time elapsed since the last change to the interface. Changes that affect the elapsed time displayed include internal events that may not have changed the state of any member.

The “Configuring Secured Port Block Allocation,” “port,” and “secured-port-block-allocation” topics should include the following note:

If you make any configuration changes to a NAT pool that has secured port block allocation configured, you must delete the existing NAT address pool, wait at least 5 seconds, and then configure a new NAT address pool. We also strongly recommend that you perform this procedure if you make any changes to the NAT pool configuration, even if you do not have secured port block allocation configured.

- The descriptions in the “Options” section of the IPsec **protocol** statement at the **[edit services ipsec-vpn ipsec proposal *proposal-name*]** and **[edit services ipsec-vpn rule *rule-name* term *term-name* then manual direction *direction*]** hierarchy levels fail to state that the **ah** and **bundle** options are not supported on MS-MPCs and MS-MICs on MX Series routers.

Chassis-Level Feature Guide

- The “Configuring Redundancy Fabric Mode for Active Control Boards on MX Series Routers” topic incorrectly states that on MX Series routers that contain the enhanced SCB with Trio chipset and the MPC3E, redundancy mode is enabled by default. The correct default behavior is that on MX Series routers that contain the enhanced SCB, regardless of the type of DPC or MPC installed on it, the default mode is the redundancy mode.

Ethernet Interfaces Feature Guide

- In the Output Fields section of the **show interfaces (10-Gigabit Ethernet)**, **show interfaces (Gigabit Ethernet)**, and **show interfaces (Fast Ethernet)** command topics of the *Ethernet Interfaces Feature Guide*, the descriptions of the **Bit errors** and **Errored blocks** fields that are displayed under the PCS Statistics section of the output are ambiguous. The following are the revised descriptions for these fields:
 - **Bit errors**—The number of seconds during which at least one bit error rate (BER) occurred while the PCS receiver is operating in normal mode.
 - **Errored blocks**—The number of seconds when at least one errored block occurred while the PCS receiver is operating in normal mode.
- The **[edit protocols lacp]** hierarchy level topic fails to mention that the ppm centralized statement is supported at this level for MX Series routers. This statement has been supported from Junos OS Release 9.4. You can use the **ppm** statement 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 protocols lacp]** hierarchy level. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing

Engine by configuring the **no-delegate-processing** configuration statement at the **[edit routing-options ppm]** statement hierarchy level.

Firewall Filters Feature Guide for Routing Devices

- The following additional information regarding the de-encapsulation of GRE packets as a terminating action for firewall filters applies to the "Firewall Filter Terminating Actions" topic:



NOTE: The *decapsulate* action that you configure at the **[edit firewall family inet filter *filter-name* term *term-name*]** hierarchy level does not process traffic with IPv4 and IPv6 options. As a result, traffic with such options is discarded by the decapsulation of GRE packets functionality.

High Availability Feature Guide for Routing Devices

- The "Nonstop Active Routing System Requirements" topic should include the **inet-mvpn** and **inet6-mvpn** protocol families for BGP in the list of supported family types. The topic previously documented that NSR supports next-generation MVPN starting with Junos OS 14.1R1, but did not include the specific names of the next-generation MVPN protocol families in the list.
- The topic "Improving the Convergence Time for VRRP" failed to include the following information:
 - Disable duplication address detection for IPv6 interfaces—Duplicate address detection is a feature of the Neighbor Discovery Protocol for IPv6. Duplicate address detection is enabled by default and determines whether an address is already in use by another node. When duplicate address detection is enabled, convergence time is high after an IPv6 interface that has been configured for VRRP tracking comes up. To disable duplicate address detection, include the **ipv6-duplicate-addr-transmits 0** statement at the **[edit system internet-options]** hierarchy level. To disable duplicate address detection only for a specific interface, include the **dad-disable** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet6]** hierarchy level.

Interchassis Redundancy Using Virtual Chassis Feature Guide for MX Series Routers

- In the "Junos OS 13.2 Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers", the "Support for MX Series Virtual Chassis (MX Series routers with MPC3E interfaces)" feature description failed to mention that you can configure a two-member MX Series Virtual Chassis on both MPC3E modules and MPC4E modules. The correct description for this feature is as follows:
 - **Support for MX Series Virtual Chassis (MX Series routers with MPC3E and MPC4E interfaces)**—Extends support for configuring a two-member MX Series Virtual Chassis to MX240, MX480, and MX960 routers with any of the following modules installed:

- MPC3E (model number MX-MPC3E-3D)
- 32x10GE MPC4E (Model number: MPC4E-3D-32XGE-SFPP)
- 2x100GE + 8x10GE MPC4E (Model number: MPC4E-3D-2CGE-8XGE)

All MX Series Virtual Chassis features are supported on these modules.

In earlier Junos OS releases, MX Series routers did not support MX Series Virtual Chassis configuration on MPC3E and MPC4E modules.

[See [Junos OS High Availability Library for Routing Devices](#) and [Junos OS for MX Series 3D Universal Edge Routers](#).]

- The following additional information applies to the “Virtual Chassis Components Overview” topic in the *Interchassis Redundancy Using Virtual Chassis Feature Guide for MX Series Routers* for Junos OS Release 11.2 and later releases.

When you configure chassis properties for MPCs installed in a member router in an MX Series Virtual Chassis, keep the following points in mind:

- Statements included at the `[edit chassis member member-id fpc slot slot-number]` hierarchy level apply to the MPC (FPC) in the specified slot number only on the specified member router in the Virtual Chassis.

For example, if you issue the `set chassis member 0 fpc slot 1 power off` statement, only the MPC installed in slot 1 of member ID 0 in the Virtual Chassis is powered off.

- Statements included at the `[edit chassis fpc slot slot-number]` hierarchy level apply to the MPCs (FPCs) in the specified slot number on *each* member router in the Virtual Chassis.

For example, if you issue the `set chassis fpc slot 1 power off` statement in a two-member MX Series Virtual Chassis, both the MPC installed in slot 1 of member ID 0 *and* the MPC installed in slot 1 of member ID 1 are powered off.



BEST PRACTICE: To ensure that the statement you use to configure MPC chassis properties in a Virtual Chassis applies to the intended member router and MPC, we recommend that you always include the `member member-id` option before the `fpc` keyword, where *member-id* is 0 or 1 for a two-member MX Series Virtual Chassis.

Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide

- The **address-allocation** statement topic fails to state the following additional information regarding addresses allocation on MS-MICs and MS-MPCs:

Regardless of whether the round-robin method of allocation is enabled by using the **address-allocation round-robin** statement, round-robin allocation is enabled by default on MS-MICs and MS-MPCs.

- The topic "Configuring Secured Port Block Allocation" contains a note listing configuration changes that require a reboot of the services PIC. The note has been updated to include a change to the NAT pool name.
- Configuration example [Configuring Inline Network Address Translation - Interface-Service Service Set](#) should state that a Modular Port Concentrator (MPC) with a Trio chipset is required, *not* a Multiservices Dense Port Concentrator.
- The following information regarding the guidelines for configuration of IP addresses for NAT processing applies to the "Configuring Source and Destination Addresses Network Address Translation Overview " section of the "Network Address Translation Rules Overview" topic:

The addresses that are specified as valid in the **inet.0** routing table and not supported for NAT translation are **orlonger** match filter types. You cannot specify any regions within such address prefixes in a NAT pool.

- The following information regarding the working of APP with NAT rules applies to the "Network Address Translation Rules Overview" topic:

For MX Series routers with MS-MICs and MS-MPCs, although the address pooling paired (APP) functionality is enabled within a NAT rule (by including the **address-pooling** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level), it is a characteristic of a NAT pool. Such a NAT pool for which APP is enabled cannot be shared with NAT rules that do not have APP configured.

Junos OS Administration Library for Routing Devices

- The **extend-size** statement topic fails to note that when you include this statement to increase the size of the configuration database, you must reboot the router after committing the configuration to make the change effective.

Layer 2 Configuration Guide, Bridging, Address Learning, and Forwarding

- The following additional information applies to the "Configuring VLAN Identifiers for Bridge Domains and VPLS Routing Instances" topic:

The maximum number of Layer 2 interfaces that you can associate with a bridge domain or a VPLS instance on MX Series routers is 4000.

Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices

- The Options section for the **flow-export-rate** statement under the hierarchy **[edit forwarding-options sampling instance *instance-name* family inet output inline-jflow]** did not include the default value. The default value is:

Default: 1 for each PFE on the FPC to which the sampling instance is applied.

- The description for the **max-packets-per-second**, **maximum-packet-length**, and **run-length** statements at the **[edit forwarding-options sampling instance *instance-name* input]** hierarchy level failed to include the following:



NOTE: This statement is not supported when you configure inline flow monitoring (by including the **inline-jflow** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6) output]** hierarchy level).

- The "Configuring RPM Timestamping" topic failed to mention that RPM timestamping is also supported on the MS-MPCs and MS-MICs on MX Series routers.

MPLS Applications Feature Guide for Routing Devices

- The "Configuring Miscellaneous LDP Properties," "Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols," "authentication-key-chain (LDP)," and "authentication-key-chain (BGP and BMP)" topics should include the following information: You must also configure the authentication algorithm using the **authentication-algorithm *algorithm*** statement. This statement must be included at the **[edit protocols (bgp | ldp)]** hierarchy level when you configure the **authentication-key-chain *key-chain*** statement at the **[edit protocols (bgp | ldp)]** hierarchy level.
- The "Path Computation for LSPs on an Overloaded Router" topic should state that when you set the overload bit on a router running IS-IS, only new LSPs are prevented from transiting through the router. Any existing Constrained Path Shortest First (CPSF) LSPs remain active and continue to transit through the router. The documentation

incorrectly states that any existing LSPs transiting through the router are also rerouted when you configure the overload bit on an IS-IS router.

Overview for Routing Devices

- The "Configuring Automatic Mirroring of the CompactFlash Card on the Hard Disk Drive" and the "mirror-flash-on-disk" topics should not include support for MX5, MX10, and MX40 3D Universal Edge Routers. On the MX Series, this feature is supported only on the MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.

OSPF Feature Guide

- **OSPF domain-id interoperability (MX Series)**— Starting in Junos OS Release 14.1R5, to enable interoperability with routers from other vendors, you can set the AS number for domain-id attributes to 0 for the following commands:

```
[edit]
user@host# set routing-instances routing-instance name protocols ospf domain-id
domain-id
user@host# set policy-options community community name members domain-id
attributes
```



CAUTION: Do not downgrade the Junos OS after configuring the AS number for domain-id attributes to 0. Set the AS number to a non-zero value and commit the configuration before downgrading the Junos OS.

Services Interfaces Configuration Guide

- The following additional information applies to the sample configuration described in the "Example: Flow-Tap Configuration" topic of the "Flow Monitoring" chapter.



NOTE: The described example applies only to M Series and T Series routers, except M160 and TX Matrix routers. For MX Series routers, because the flow-tap application resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC), the Packet Forwarding Engine must send the packet to a tunnel logical (vt-) interface to encapsulate the intercepted packet. In such a scenario, you need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use.

- The following additional information applies to the working of basic NAT on AMS interfaces of MS-MPCs and MS-MICs for the "Aggregated Multiservices Interface" section of the "Understanding Aggregated Multiservices Interfaces" topic:



NOTE: With basic NAT44, load balancing on AMS interfaces of MS-MICs and MS-MPCs does not work properly if the ingress hash key is source IP address and the egress hash key is destination IP address.

If the service set is applied on the Gigabit Ethernet or 10-Gigabit Ethernet interface that functions as the NAT inside interface, then the hash keys used for load balancing might be configured in such a way that the ingress key is set as destination IP address and the egress key is set as source IP address. Because the source IP address undergoes NAT processing, it is not available for hashing the traffic in the reverse direction. Therefore, load balancing does not happen on the same IP address, and forward and reverse traffic do not map to the same PIC. With the hash keys reversed, load balancing occurs correctly.

With next-hop services, for forward traffic, the ingress-key on the inside-interface load-balances traffic, and for reverse traffic, the ingress-key on the outside-interface load-balances traffic or per-member-next-hops steer reverse traffic. With interface-style services, the ingress-key load-balances forward traffic, and the egress-key load-balances forward traffic or per-member-next-hops steer reverse traffic. Forward traffic is traffic entering from the inner side of a service-set, and reverse traffic is traffic entering from the outer side of a service-set. The forward key is the hash key used for the forward direction of traffic, and the reverse key is the hash key used for the reverse direction of traffic (depends on whether it relates to interface-services or next-hop-services style.)

With stateful firewalls, you can configure the following combinations of forward and reverse keys for load balancing. In the following combinations presented for hash keys, FOR-KEY refers to the forward key, REV-KEY denotes the reverse key, SIP signifies source IP address, DIP signifies destination IP address, and PROTO refers to protocol such as IP.

- FOR-KEY: SIP, REV-KEY: DIP
- FOR-KEY: SIP,PROTO REV-KEY: DIP, PROTO
- FOR-KEY: DIP, REV-KEY: SIP
- FOR-KEY: DIP,PROTO REV-KEY: SIP, PROTO
- FOR-KEY: SIP,DIP REV-KEY: SIP, DIP
- FOR-KEY: SIP,DIP,PROTO REV-KEY: SIP, DIP,PROTO

With static NAT configured as basic NAT44 or destination NAT44, and with stateful firewall configured or not, if the forward direction of traffic must undergo NAT processing, configure the hash keys as follows:

- FOR-KEY: DIP, REV-KEY: SIP
- FOR-KEY: DIP,PROTO REV-KEY: SIP, PROTO

If the reverse direction of traffic must undergo NAT processing, configure the hash keys as follows:

- FOR-KEY: SIP, REV-KEY: DIP
- FOR-KEY: SIP,PROTO REV-KEY: DIP, PROTO

With dynamic NAT configured, and with stateful firewall configured or not, only the forward direction traffic can undergo NAT. The forward hash key can be any combination of SIP, DIP, and protocol, and the reverse hash key is ignored.

-
- The functionality to log the cflowd records in a log file before they are exported to a cflowd server (by including the **local-dump** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output flow-server *hostname*]** hierarchy level) is not supported when you configure inline flow monitoring (by including the **inline-jflow** statement at the **[edit forwarding-options sampling instance *instance-name* family inet output]** hierarchy level).
 - The following information regarding the interoperation of FTP ALG and address-pooling paired features is missing from the "ALG Descriptions" topic of the "Application Properties" chapter:

On MS-MPCs and MS-MICs, for passive FTP to work properly without FTP application layer gateway (ALG) enabled (by not specifying the **application junos-ftp** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* from]** and the **[edit services nat rule *rule-name* term *term-name* from]** hierarchy levels), you must enable the address pooling paired (APP) functionality (by including the **address-pooling** statement at the **[edit services nat rule *rule-name* term *term-name* then translated]** hierarchy level). Such a configuration causes the data and control FTP sessions to receive the same NAT address.

- The following information regarding the restriction on prefix lengths you can configure in NAT pools on MS-MPCs and MS-MICs applies to the "Configuring Source and Destination Addresses Network Address Translation Overview" section of the "Network Address Translation Rules Overview" topic:

On MX Series routers with MS-MPCs and MS-MICs, if you configure a NAT address pool with a prefix length that is equal to or greater than /16, the PIC does not contain sufficient memory to provision the configured pool. Also, memory utilization problems might occur if you attempt to configure many pools whose combined total IP addresses exceed /16. In such circumstances, a system logging message is generated stating that the NAT pool name failed to be created and that the service set is not activated. On MS-MPCs and MS-MICs, you must not configure NAT pools with prefix lengths greater than /16.

- The "Configuring Unicast Tunnels" topic incorrectly shows the **backup-destination** statement. This statement does not apply to unicast tunnels and should be removed.

Standards Reference

- The *Supported Network Management Standards* topic incorrectly states that Junos OS supports `mplsL3VpnIfConfTable` as part of compliance with RFC 4382, MPLS/BGP Layer 3 Virtual Private Network (VPN) MIB. Junos OS does not support this table.

Subscriber Management Network Access Feature Guide

- The [LAC Tunnel Selection Overview](#), [Configuring Weighted Load Balancing for LAC Tunnel Sessions](#), and [weighted-load-balancing \(L2TP LAC\)](#) topics in the *Junos OS Broadband Subscriber Management and Services Library* incorrectly describe how weighted load balancing works on an L2TP LAC. The topics state that the tunnel with the highest weight (highest session limit) within a preference level is selected until it has reached its maximum sessions limit, and then the tunnel with the next higher weight is selected, and so on.

In fact, when weighted load balancing is configured, tunnels are selected randomly within a preference level, but the distribution of selected tunnels is related to their weight. The LAC generates a random number within a range equal to the aggregate total of all session limits for all tunnels in the preference level. Portions of the range—pools of numbers—are associated with the tunnels according to their weight; a higher weight results in a larger pool. The random number is more likely to be in a larger pool, so a tunnel with a higher weight (larger pool) is more likely to be selected than a tunnel with a lower weight (smaller pool).

For example, consider a level that has only two tunnels, A and B. Tunnel A has a maximum sessions limit of 1000 and tunnel B has a limit of 2000 sessions, resulting in an aggregate total of 3000 sessions. The LAC generates a random number in the range from 0 through 2999. A pool of 1000 numbers, the portion of the range from 0 through 999, is associated with tunnel A. A pool of 2000 numbers, the portion of the range from 1000 through 2999, is associated with tunnel B. If the generated number is less than 1000, then tunnel A is selected, even though it has a lower weight than tunnel B. If the generated number is 1000 or larger, then tunnel B is selected. Because the pool of possible generated numbers for tunnel B (2000) is twice that for tunnel A (1000), tunnel B is, *on average*, selected twice as often as tunnel A.

Subscriber Management Provisioning Guide

- The table in topic “AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS” incorrectly indicates that VSA 26-1 (Virtual-Router) supports CoA Request messages. VSA 26-1 does not support CoA Request messages.

System Log Messages Reference

- The formats of the MSVCS_LOG_SESSION_OPEN and MSVCS_LOG_SESSION_CLOSE system log messages in the "MSVCS System Log Messages" chapter are incorrectly specified. The following is the correct and complete format of the MSVCS_LOG_SESSION_OPEN and MSVCS_LOG_SESSION_CLOSE system log messages:

*App: application, source-interface-name fpc/pic/port\address in hexadecimal format
source-address:source-port source-nat-information ->
destination-address:destination-port destination-nat-information (protocol-name)
hh:mm:ss.milliseconds protocol-name (tos tos-bit-value, ttl ttl-value, id id-number,
offset offset-value, flags [ip-flag-type], proto protocol- name (protocol-id), length
number)*

Traffic Sampling, Forwarding, and Monitoring Feature Guide for Routing Devices

- The **enhanced-hash-key** configuration statement topic fails to mention that the **src-prefix-len** option is available for configuration at the **[edit forwarding-options enhanced-hash-key family inet6 layer-3-services src-prefix-len]** hierarchy level. You can use the **src-prefix-len** option to include the source prefix length in the hash key for enhanced IP forwarding engines.

User Access and Authorization Feature Guide for Routing Devices

- The "Configuring the SSH Protocol Version" topic incorrectly states that both version 1 and version 2 of the SSH protocol are enabled by default. The topic should state that version 2 of the SSH protocol is enabled by default, and you must explicitly configure version 1 if you want to enable it.
- The "Example: DHCP Complete Configuration" and "dhcp" topics should not include support for the MX Series Universal Edge 3D Routers. This feature is supported only on the M Series and the T Series.

VPLS Feature Guide for Routing Devices

- The following information regarding the working of firewall filters and policers with MAC addresses applies to the "Configuring Firewall Filters and Policers for VPLS " topic:

The behavior of firewall filters processing with MAC addresses differs between DPCs and MPCs. On MPCs, interface filters are always applied before MAC learning occurs. The input forwarding table filter is applied after MAC learning is completed. However, on DPCs, MAC learning occurs independently of the application of filters. If the CE-facing interface of the PE where the firewall filter is applied is an MPC, then the MAC entry times out and is never learned again. However, if the CE-facing interface of the PE where the firewall filter is applied is a DPC, then the MAC entry is not timed out and if the MAC address entry is manually cleared, it is relearned.

VPNs Library for Routing Devices

- The "Routing Instances Overview" topic should include the following instance types: Ethernet VPN (EVPN) and Internet Multicast over MPLS. Use the Ethernet VPN instance type, which is supported on the MX Series only, to connect a group of dispersed customer sites using a Layer 2 virtual bridge. Use the Internet Multicast over MPLS instance type to provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.

To configure an EVPN instance type, include the **evpn** statement at the **[edit routing-instances *routing-instance-name* instance-type]** hierarchy level. To configure an Internet Multicast over MPLS instance type, include the **mpls-internet-multicast** statement at the **[edit routing-instances *routing-instance-name* instance-type]** hierarchy level.

Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)
- [Product Compatibility on page 199](#)

Migration, Upgrade, and Downgrade Instructions

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

- [Basic Procedure for Upgrading to Release 14.1 on page 192](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 194](#)
- [Upgrading a Router with Redundant Routing Engines on page 194](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 195](#)
- [Upgrading the Software for a Routing Matrix on page 196](#)
- [Upgrading Using Unified ISSU on page 197](#)
- [Downgrading from Release 14.1 on page 197](#)
- [Changes Planned for Future Releases on page 198](#)

Basic Procedure for Upgrading to Release 14.1

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



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

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

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

The download and installation process for Junos OS Release 14.1 is different from previous Junos OS releases.

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<http://www.juniper.net/support/downloads/>
2. Select the name of the Junos 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-14.1R71-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-14.1R71-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 14.1 **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 11.4, 12.3, and 13.3 are EEOL releases. You can upgrade from Junos OS Release 11.4 to Release 12.3 or even from Junos OS Release 11.4 to Release 13.3. 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 12.1 (a non-EEOL release) to Junos OS Release 13.3 or directly downgrade from Junos OS Release 13.3 to Junos OS Release 12.1.

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>

Upgrading a Router with Redundant Routing Engines

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

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

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

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

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 also 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 [Multicast Protocols Feature Guide for Routing Devices](#).

Upgrading the Software for a Routing Matrix

A routing matrix can be 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 line-card chassis (LCCs) in the routing matrix (specified in the Junos OS CLI by using the **lcc** option). To avoid network disruption during the upgrade, ensure the following conditions 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 of 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 all LCCs connected to the SCC or SFC are all **re0** or are all **re1**.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC 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 the 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 or RE-DUO-C1800-16G Routing Engines.



BEST PRACTICE: 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, perform the following steps:

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. 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.
4. Install the new software on the new backup Routing Engine (re0).

For the detailed procedure, see the [Routing Matrix with a TX Matrix Router Deployment Guide](#) or the [Routing Matrix with a TX Matrix Plus Router Deployment Guide](#).

Upgrading Using Unified 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 [High Availability Feature Guide for Routing Devices](#).

Downgrading from Release 14.1

To downgrade from Release 14.1 to another supported release, follow the procedure for upgrading, but replace the 14.1 **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 [Installation and Upgrade Guide](#).

Changes Planned for Future Releases

- **Introduction of the all keyword to prevent accidental execution of certain clear commands**—The **all** keyword is planned to be introduced in Junos OS Release 14.2 (as an optional keyword) and in Junos OS Release 15.2 (as a mandatory keyword) for certain **clear** commands that are used for clearing protocol and neighbor sessions. This makes users explicitly select the **all** keyword to clear all protocol or session information. Thus, it prevents accidental clearing or resetting of protocols or neighbor sessions, which might disrupt network operations.

The **all** keyword is planned to be introduced for the following **clear** commands:

- **clear arp**
- **clear bgp neighbor**
- **clear bfd adaptation**
- **clear bfd session**
- **clear igmp membership**
- **clear isis adjacency**
- **clear isis database**
- **clear ldp neighbor**
- **clear ldp session**
- **clear mld membership**
- **clear mpls lsp**
- **clear msdp cache**
- **clear multicast forwarding-cache**
- **clear (ospf | ospf3) database**
- **clear (ospf | ospf3) neighbor**
- **clear pim join**
- **clear pim join-distribution**
- **clear pim register**
- **clear rsvp sessions**

In Junos OS Release 14.2 and 15.1—The **all** keyword will be *optional*. Therefore, when you type any of these **clear** commands followed by the **?** in the CLI, the **all** keyword will be listed as an option after the **<[Enter]>** keyword. You can execute the **clear** command directly or with the **all** keyword to clear all information. For example, when you type **clear mpls lsp ?**, you will see:

```
user@host> clear mpls lsp ?
```

Possible completions:

```
<[Enter]>      Execute this command
all            Reset 'all' the nontransit or egress LSPs
               originating on this router      <----->
autobandwidth  Clear LSP autobandwidth counters
logical-system Name of logical system, or 'all'
name           Regular expression for LSP names to match
optimize       Perform nonpreemptive optimization computation now
...
```

Both **clear mpls lsp** or **clear mpls lsp all** will function identically in these releases.

In Junos OS Release 15.2 and later—The **all** keyword will be *mandatory*. Therefore, when you type a **clear** command followed by the **?** in the CLI, the **<[Enter]>** option to execute the command directly (without specifying any options) will not be available.

For example, when you type **clear mpls lsp ?**, you will see **all** listed as an option but not **<[Enter]>** to execute the command directly. Therefore, you will have to type **clear mpls lsp all** and then press **<[Enter]>** if you want to clear information about all the nontransit or egress LSPs originating on the router.

```
user@host> clear mpls lsp ?
```

Possible completions:

```
all            Reset 'all' the nontransit or egress LSPs
               originating on this router      <----->
autobandwidth  Clear LSP autobandwidth counters
logical-system Name of logical system, or 'all'
name           Regular expression for LSP names to match
optimize       Perform nonpreemptive optimization computation now
...
```

Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Product Compatibility on page 199](#)

Product Compatibility

- [Hardware Compatibility on page 200](#)

Hardware Compatibility

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

To determine the features supported on M Series, MX Series, and T Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

<http://pathfinder.juniper.net/feature-explorer/>

Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 57](#)
- [Known Behavior on page 70](#)
- [Known Issues on page 72](#)
- [Resolved Issues on page 93](#)
- [Documentation Updates on page 179](#)
- [Migration, Upgrade, and Downgrade Instructions on page 192](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 14.1R7 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Resolved Issues on page 212](#)
- [Documentation Updates on page 221](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)
- [Product Compatibility on page 225](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 14.1R7 for the PTX Series.

- [Hardware on page 201](#)
- [Interfaces and Chassis on page 203](#)
- [MPLS on page 205](#)
- [Network Management and Monitoring on page 206](#)
- [Routing Protocols on page 206](#)

Hardware

- **New FPC with eight Packet Forwarding Engines (PTX5000)**—Starting in Junos OS Release 14.1, a new FPC (FPC2-PTX-P1A), with eight Packet Forwarding Engines and two PIC slots, is supported on the PTX5000. The FPC is capable of forwarding at 960 Gbps speed, and it supports 300W of PIC power per PIC slot. The new FPC supports the following PICs from Release 14.1:
 - P2-100GE-CFP2 (4x100G CFP2 PIC)
 - P1-PTX-24-10GE-SFPP (24x10G LAN PIC)
 - P1-PTX-24-10G-W-SFPP (24x10G LAN/WAN PIC)
 - P1-PTX-2-100G-C-WDM-C (2x100G LH DWDM PIC)

The following PICs are supported on the FPC from Release 14.1R2:

- P1-PTX-2-40GE-CFP (2x40-Gigabit Ethernet PIC with CFP)

- P1-PTX-2-100GE-CFP (2x100-Gigabit Ethernet PIC with CFP)
- **New 4-port 100-Gigabit Ethernet PIC (PTX5000)**—Beginning with Junos OS Release 14.1, a new 4-port 100-Gigabit Ethernet PIC with CFP2 (P2-100GE-CFP2) is supported on the FPC FPC2-PTX-P1A in a PTX5000. The PIC supports 100GBASE-LR4 and 100GBASE-SR10 transceivers. The CFP2-100G-SR10-D transceiver is not dual-rate, it supports Ethernet only. The CFP2-100G-SR10-D2 transceiver is dual-rate, but only when used in a PIC that supports OTN, such as P2-100GE-OTN.
- **New SIB to support high density FPC (PTX5000)**—Starting in Junos OS Release 14.1, a new high-density SIB (SIB2-I-PTX5000) provides switch fabric capacity of 960 Gbps speed per FPC slot for the FPC FPC2-PTX-P1A in a PTX5000.
- **New high-capacity DC PSM and PDU (PTX5000)**—Starting in Junos OS Release 14.1, the following DC power supply module (PSM) and DC power distribution unit (PDU) are added to provide power to a new, high-density FPC—FPC2-PTX-P1A—and other components in a PTX5000:
 - PTX High Capacity-60A DC PDU (PDU2-PTX-DC)
 - PTX High Capacity-60A DC PSM (PSM2-PTX-DC)
- **Fabric capacity on PTX5000**—Starting with Junos OS Release 14.1, the PTX5000 supports nine Switch Interface Boards (SIBs). The packet transport router with FPC2-PTX-P1A FPCs provides up to 16 terabits per second (Tbps), full duplex (8 Tbps of any-to-any, nonblocking, half-duplex) switching. The chassis with SIB-I-PTX5008 provides an 8+1 active redundancy that supports line rate for all the eight FPC slots.

[See [Fabric Fault Handling Overview on PTX5000 Packet Transport Router](#).]
- **Enhanced midplane (PTX5000)**—Starting in Junos OS Release 14.1, the PTX5000 supports a new enhanced midplane. The PTX5000BASE2 model is a chassis with an enhanced midplane that requires high capacity 60-A DC PDUs and PSMs. The enhanced midplane is identified as **Midplane-8Se** in the output from the **show chassis hardware** operational-mode CLI command.
- **New AC PSM and PDU (PTX5000)**—Starting with Junos OS Release 14.1R2, new AC power supply modules (PSMs) and power distribution units (PDUs) are added to provide power to the FPC2-PTX-P1A FPC and other components in a PTX5000 router. You can install two redundant AC PDUs and each AC PDU supports up to eight PSMs. All PSMs are considered to be a part of single zone to provide power to a common power bus. Run the **show chassis hardware** operational mode command to view the AC PSM and PDU details. The **show chassis environment pdu pdu-number** displays the firmware version for all the microcontrollers on the PDU.
- **Support for 4-port 100-Gigabit Ethernet OTN PIC (PTX5000)**—Starting with Junos OS Release 14.1R2, a 4-port 100-Gigabit Ethernet OTN PIC—P2-100GE-OTN—is supported on the FPC2-PTX-P1A FPC in PTX5000 routers.
- **Support for P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC (PTX5000)**—Starting with Junos OS Release 14.1R2, PTX5000 supports the P2-10G-40G-QSFPP PIC on the FPC2-PTX-P1A FPC. You can configure the P2-10G-40G-QSFPP PIC to operate in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode.

Interfaces and Chassis

- **Support for physical interface damping (PTX Series)**—Beginning with Junos OS Release 14.1, interface damping is supported on *physical interfaces* to address periodic flaps with long up and down durations (in seconds) as opposed to instantaneous multiple flaps with very short up and down durations (in milliseconds) addressed by the Interface hold timers. When the interface is placed in the suppressed state, the interface link state is set to down. Interface event damping uses an exponential back-off algorithm to suppress interface up and down event reporting to the upper-level protocols. To configure interface damping, include the **damping** statement at the **[edit interfaces *interface-name*]** hierarchy level. You use the **show interfaces extensive** command to view the interface damping values and link state.
- **Adaptive load balancing (ALB) for aggregated Ethernet bundles (PTX Series)**—Starting in Junos OS Release 14.1, ALB evenly distributes data flows across aggregated Ethernet member links. Network administrators use this feature to manage uneven or overloaded data flows on member links. ALB supports up to 32 member links and up to 50 aggregated Ethernet bundles. The algorithm determines which link to use by considering the scanned packet or bit rate associated with each hash value in conjunction with the mapping of hash values to a given link. ALB is applied to IPv4, IPv6, and MPLS packet headers. ALB is disabled by default.



NOTE: ALB is not applied to multicast traffic.

To configure ALB, include the **adaptive** statement at the **[edit interfaces *ae-interface* aggregated-ether-options load-balance]** hierarchy level. Under the adaptive statement, you can set the following ALB options: tolerance percentage, scan-interval, and pps.

[See [Configuring Aggregated Ethernet Interfaces on PTX Series Packet Transport Routers](#).]

- **SFPP-10G-CT50-ZR (PTX Series)**—The SPFF-10G-CT50-ZR tunable transceiver provides a duplex LC connector and supports the 10GBASE-Z optical interface specification and monitoring. The transceiver is not specified as part of the 10-Gigabit Ethernet standard and is instead built according to Juniper Networks specifications. Only WAN-PHY and LAN-PHY modes are supported. To configure the wavelength on the transceiver, use the **wavelength** statement at the **[edit interfaces *interface-name* optics-options]** hierarchy level. The following interface module supports the SPFF-10G-CT50-ZR transceiver:

- 10-Gigabit Ethernet LAN/WAN OTN PIC with SFP+ (model number: P1-PTX-24-10G-W-SFPP)—Supported in Junos OS Release 13.2R3, 13.3R2, 14.1, and later

For more information about interface modules, see the “Cables and Connectors” section in the *Interface Module Reference* for your router.

[See [10-Gigabit Ethernet 10GBASE Optical Interface Specifications](#) and [wavelength](#).]

- **SFPP-10G-ZR-OTN-XT (PTX Series)**—The SFPP-10G-ZR-OTN-XT dual-rate extended temperature transceiver provides a duplex LC connector and supports the 10GBASE-Z optical interface specification and monitoring. The transceiver is not specified as part

of the 10-Gigabit Ethernet standard and is instead built according to ITU-T and Juniper Networks specifications. The following interface modules support the SFPP-10G-ZR-OTN-XT transceiver:

- 10-Gigabit Ethernet PIC with SFP+ (model number: P1-PTX-24-10GE-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later
- 10-Gigabit Ethernet LAN/WAN OTN PIC with SFP+ (model number: P1-PTX-24-10G-W-SFPP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3, and later

For more information about interface modules, see the “Cables and Connectors” section in the *Interface Module Reference* for your router.

[See [10-Gigabit Ethernet 10GBASE Optical Interface Specifications](#).]

- **New Flexible PIC Concentrator (FPC) model number FPC-SFF-PTX-T (PTX3000)**—Starting in Junos OS Release 14.1, a new FPC is supported on the PTX3000. The FPC-SFF-PTX-T does not interoperate with other Type 5 FPCs in the same chassis. The FPC-SFF-PTX-T model has a 10ms RTT buffer capacity and does not support IPv6 or IP multicast features.

[See [PTX3000 FPCs Supported](#).]

- **Support for high-density FPC (PTX5000)**—Starting with Junos OS Release 14.1, a new high-density FPC, FPCE (model number: FPC2-PTX-P1A), is supported on the PTX5000. This FPC has eight Packet Forwarding Engines and a forwarding capacity of 9600 million packets per second (Mpps).

[Table 1 on page 204](#) provides information regarding the Type 5 PICs that are supported on the FPC2-PTX-P1A FPC:

Table 1: Type 5 PICs Supported on FPC2-PTX-P1A

Type 5 PIC	PIC Model Number
10-Gigabit Ethernet PIC with SFP+	P1-PTX-24-10GE-SFPP
10-Gigabit Ethernet LAN/WAN OTN PIC with SFP+	P1-PTX-24-10G-W-SFPP
100-Gigabit DWDM OTN PIC	P1-PTX-2-100G-WDM
100-Gigabit Ethernet PIC with CFP2	P2-100GE-CFP2

To meet the increased power requirements of the high-density FPC, the following new power distribution unit (PDU) and power supply module (PSM) are supported on the PTX5000:

- PTX High Capacity 60A DC PDU (PDU2-PTX-DC)
- PTX High Capacity 60A DC PSM (PSM2-PTX-DC)



NOTE: The PTX High Capacity 60A DC PDU can support a maximum of eight PSMs.

[See [PTX5000 FPCs Supported](#).]

- **Support for dual-rate speed (PTX Series)**—Starting in Junos OS Release 14.1R3, for tPTX3000, support for the dual-rate speed for the 24-port 10-Gigabit Ethernet PIC (P1-PTX-24-10GE-SFPP) enables you to switch all port speeds to either 1-Gigabit Ethernet or 10-Gigabit Ethernet. The default is 10 Gbps. All ports are configured to the same speed; there is no mixed-rate-mode capability. You can use either the SFP-1GE-SX or the SFP-1GE-LX transceiver for 1 Gbps. Changing the port speed causes the PIC to reboot.

To configure all ports on the P1-PTX-24-10GE-SFPP to operate at 1 Gbps, use the **speed 1G** statement at the **[edit chassis fpc fpc-number pic pic-number]** hierarchy level. To return all ports to the 10-Gbps speed, use the **delete chassis fpc fpc-number pic pic-number speed 1G** command.

[See [speed \(24-port and 12-port 10 Gigabit Ethernet PIC\)](#) and [10-Gigabit Ethernet PIC with SFP+ \(PTX Series\)](#).]

MPLS

- **Require BFD-triggered Packet Forwarding Engine local repair (PTX Series)**—Starting in Junos OS Release 14.1, this feature enables you to configure BFD and MPLS ping for fast-failure detection without relying on fast physical level detection. With links between routers, when a route goes down, the local Packet Forwarding Engine does a local repair and traffic is quickly re-routed around the broken link. The RPD is then informed of the down link and does a global repair and pushes down the updated route information to all other FPCs.

[See [PTX Series Packet Transport Routers](#).]

- **Link protection for MLDP**—Beginning in Junos OS Release 14.1, link protection for MLDP is supported to enable fast reroute of traffic carried over LDP LSPs in case of a link failure. LDP point-to-multipoint LSPs can be used to send traffic from a single root or ingress node to a number of leaf nodes or egress nodes traversing one or more transit nodes. When one of the links of the point-to-multipoint tree fails, the subtrees may get detached until the IGP reconverges and MLDP initiates label mapping using the best path from the downstream to the new upstream router. To protect the traffic in the event of a link failure, you can configure an explicit tunnel so that traffic can be rerouted using the tunnel. Junos OS supports make-before-break (MBB) capabilities to ensure minimum packet loss when attempting to signal a new LSP path before tearing down the old LSP path. This feature also adds targeted LDP support for MLDP link protection.

[See [Example: Configuring LDP Link Protection](#).]

- **Entropy label and FAT label support (PTX Series)**—Starting in Release 14.1, Junos OS supports entropy labels and Flow Aware Transport for Pseudowires (FAT) labels. Entropy label and FAT label when configured on the label-switching routers (LSRs)

and label edge routers (LERs) perform load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload.

In Junos OS Release 14.1, entropy labels can be used for RSVP-signaled label-switched paths (LSPs) and point-to-point LDP-signaled LSPs. FAT flow labels can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for virtual private LAN service (VPLS) and virtual private wire service (VPWS) networks.

[See [Configuring the Entropy Label for LSPs](#) and [FAT Flow Labels Overview](#).]

Network Management and Monitoring

- **SNMP notifying target for removed notify target configuration (PTX Series)**—Beginning with Junos OS Release 14.1, when a trap target is deleted from Juniper Networks devices, either a syslog event or a syslog trap is generated as per the user configuration. The existing SNMP trap `jnxSyslogTrap` is sent to all target network management systems (NMSs) specified in the SNMP agent including the target NMS, which is being deleted. By default, in the event of target deletion, only a syslog event is generated. To trigger a trap on deletion of a trap target, configure a syslog event policy, which sends the syslog as a trap to the network management systems.

Routing Protocols

- **Selecting backup LFA for IS-IS routing protocol (PTX Series)**—Starting with Junos OS Release 14.1, the default loop-free alternate (LFA) selection algorithm or criteria can be overridden with an LFA policy. These policies are configured for each destination (IPv4 and IPv6) and a primary next-hop interface. These backup policies enforce LFA selection based on admin-group, srlg, neighbor, neighbor-tag, bandwidth, protection-type, and metric attributes of the backup path. During backup shortest-path-first (SPF) computation, each attribute (both node and link) of the backup path, stored per backup-next hop, is accumulated by IGP. For the routes created internally by IGP, the attribute set of every backup path is evaluated against the policy configured per destination per prefix primary next-hop interface. The first or the best backup path is selected and installed as the backup next hop in the routing table. To configure the backup selection policy, include the **backup-selection** configuration statement at the **[edit routing-options]** hierarchy level. The **show backup-selection** command displays the configured policies for a given interface and destination. The display can be filtered against a particular destination, prefix, interface, or logical systems.

Related Documentation

- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Resolved Issues on page 212](#)
- [Documentation Updates on page 221](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)

- [Product Compatibility on page 225](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 14.1R7 for the PTX Series.

- [Interfaces and Chassis on page 208](#)
- [Network Management and Monitoring on page 208](#)
- [Routing Protocols on page 208](#)
- [VPNs on page 208](#)

Interfaces and Chassis

- **Configuring the Ethernet mode for the P2-10G-40G-QSFPP PIC (PTX Series)**—You can configure the P2-10G-40G-QSFPP PIC to operate either in 10-Gigabit Ethernet mode or in 40-Gigabit Ethernet mode. When the PIC is in 40-Gigabit Ethernet mode, you must execute the **show interfaces diagnostics optics et-fpc/pic/port** command. The output of this command displays the diagnostic optics information of the corresponding 40-Gigabit Ethernet port of the PIC. However, when the PIC is in 10-Gigabit Ethernet mode, you must execute the **show interfaces diagnostics optics et-fpc/pic/port:channel** command. The output of this command displays the diagnostic optics information of the corresponding 10-Gigabit Ethernet port of the PIC.

[See [P2-10G-40G-QSFPP PIC Overview](#).]

Network Management and Monitoring

- **New system log message indicating the difference in the Packet Forwarding Engine counter value (PTX Series)**—Effective in Junos OS Release 14.1R3, if the counter value of a Packet Forwarding Engine is reported lesser than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the **MIB2D_COUNTER_DECREASING** system log message for that specific counter.

[See [MIB2D_COUNTER_DECREASING](#).]

Routing Protocols

- **Configuring and establishing targeted sessions with third-party controllers using LDP targeted neighbor (PTX Series)**—Starting with Junos OS Release 14.1R5, you can configure LDP targeted neighbor to third-party controllers for applications such as route recorder that wants to learn label-FEC bindings of an LSR. LDP targeted neighbor helps to establish a targeted session with controllers for a variety of applications.

VPNs

- **Support for chained composite next hops for Layer 3 VPN transit traffic (PTX Series)**—Starting in Junos OS Release 14.1, chained composite next hops for Layer 3 VPN transit traffic are enabled by default on PTX Series routers. You no longer need to configure the **transit l3vpn** statement at the **[edit routing-options forwarding-table chained-composite-next-hop]** hierarchy level. Chained composite next hops facilitate the handling of large volumes of transit traffic in the core of large networks.

[See [Chained Composite Next Hops for Transit Devices](#).]

Related Documentation

- [New and Changed Features on page 201](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Resolved Issues on page 212](#)
- [Documentation Updates on page 221](#)

- [Migration, Upgrade, and Downgrade Instructions on page 222](#)
- [Product Compatibility on page 225](#)

Known Behavior

There are no changes in known behavior in Junos OS Release 14.1R7 for the PTX Series.

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

Related Documentation

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Issues on page 209](#)
- [Documentation Updates on page 221](#)
- [Resolved Issues on page 212](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)
- [Product Compatibility on page 225](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R7 for the PTX Series.

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

- [General Routing on page 209](#)
- [Class of Service \(CoS\) on page 210](#)
- [Interfaces and Chassis on page 210](#)
- [MPLS on page 211](#)
- [Routing Protocols on page 211](#)

General Routing

- The ifl count is incorrect and will not be repaired until a pic restart. [PR882406](#)
- CCG locks to cc-8k even when configured signal type is cc-8k-400, without off-frequency [PR895450](#)
- This is expected behavior because ipv6 transit stat is not supported yet. [PR896253](#)
- The maximum number of Ingress LSP each PFE can support is changed from 16k to 10k. If more ingress LSPs are needed, we need to distribute them to different PFEs (which can be on the same or different FPC). [PR900388](#)
- It is reported that on PTX Series platforms, when the firewall filter is configured on the loopback interface of the device, due to bad error handling or NULL pointer, all the FPC

on device may continuously crash and unstable. Because the issue is not reproducible, the trigger of the issue is not clear. [PR996749](#)

- In case ATS interface remain down and Control plane link to CSE "CB Link" is up. In this scenario memory leak in ESP interfaces tx ring on CSE causes continues message as below or it can crash jtapperd as well. >jtapperd: Err: Unable to get mbuff from the tx memory pool [PR1080550](#)
- On PTX Series platform, the FPC may crash when the interface goes down (e.g. disabling the interface or interface flap occurs) due to failure during CNH NH change (composite next hop, e.g. created in P2MP LSP scenario). Based on the current observation, the probability of the issue might be low. [PR1082429](#)
- 14.1R6 - P2MP + NSR may observe ~100ms loss during RE switchover in steady state with link protection enabled. This however is not observed from releases 14.2 and higher. The problem is undergoing further triage and is targeted to enter 14.1R7 release window. [PR1095488](#)
- The Laser output power alarm was incorrect prior to 14.2 and it was modified through the [PR1084963](#). "Laser output power" was not displayed correctly. It was displaying the transmitter output power which is not related to the thresholds and would not generate alarms if it was not in the threshold range. [PR1098417](#)
- When PTX series router is used as a PE router for L2circuit connections, there are certain Layer 2 related protocols like LACP, LLDP, STP that will not get tunneled in to the L2circuit path by the PTX. Instead of tunneling the packets into the L2circuit path, the PTX will punt it to the RE causing the packets to not reach the other end of the L2circuit. Packets like ARP work fine. The PTX chipset performs a classification on the affected L2 packets and sends it to the RE instead of pushing in through the MPLS tunnel in L2circuit scenarios. [PR1131797](#)

Class of Service (CoS)

- In case of member links of an AE (aggregated ethernet) interface scatter over multiple PFEs, if the FPC where member links of the AE interface reside gets reset or interface disable, there may be a dip in the output of SNMP walk on AE related queue MIB (such as jnxCosQstatTxedPkts). The behavior is intermittent and not seen every time. [PR1122343](#)

Interfaces and Chassis

- On dual Routing Engine platforms, when adding the logical interfaces (IFLs) and committing, due to the device control process (dcd) on backup Routing Engine may fail to process the configuration and keep it in the memory, in some cases (not happening all the time), it might be observed that the memory of the dcd keeps increasing on backup Routing Engine. [PR1014098](#)
- In the dual Routing Engines scenario with GRES and ae0 interfaces configuration, if GRES is disabled on system, the backup Routing Engine should remove the ae0 bundle, however it doesn't go clean and ae0 remains in backup Routing Engine. After switching Routing Engine mastership to make other Routing Engine as master, the new master

Routing Engine (which was backup earlier) continues to use invalid MAC address "00:00:00:00:00:00". [PR1089946](#)

MPLS

- Currently we allow configuration of both fast-reroute and link-protection/node-link-protection on a single LSP. But when customer configure both types of protection on the LSPs, it might cause scaling issues in customer network. Now, there is a change requirement, restrict the configuration to either fast-reroute or link/node-link protection on per LSP basis. [PR860960](#)
- In RSVP-based P2MP scenario, if a sub-LSP switchover to bypass LSP due to PIC offline, then when a new sub-LSP is established via setup-protection, then the deletion of the old sub-LSP might result in deletion of both the sub-LSPs. [PR1132050](#)

Routing Protocols

- When we have two paths for the same route, the route gets pointed to Unilist NH which in turn gets pointed to two separate Unicast NHs. The route is determined by OSPF and we have BFD enabled on one of the paths which runs through an I2circuit path. When the link on the I2circuit gets cut, the link flap is informed by BFD as well as through OSPF LSAs. Ideally the BFD should inform the link down event before the OSPF LSA. But at the current situation, the OSPF LSAs update the current event a second before BFD. Due to this reason, we do get the route to be pointing to a new UNILIST NH with the weights swapped. But the Unicast NH for which the L3 link is down, gets added to the Unilist NH, the BFD assumes the link to be up, and hence updates the weights inappropriately and hence we do see traffic loss. Once the BFD link down event is processed at OSPF protocol level, now the route points to only UNICAST NH and hence we do see traffic flowing through the currently active link. The traffic outage would be hardly for less than a second during FRR. Also, this can be avoided if the BFD keepalive intervals are maintained around 50 ms with multiplier as 3 as opposed to 100ms with a multiplier of 3. [PR1119253](#)

Related Documentation

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Resolved Issues on page 212](#)
- [Documentation Updates on page 221](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)
- [Product Compatibility on page 225](#)

Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

- [Resolved Issues: 14.1R7 on page 212](#)
- [Resolved Issues: 14.1R6 on page 213](#)
- [Resolved Issues: 14.1R5 on page 215](#)
- [Resolved Issues: 14.1R4 on page 217](#)
- [Resolved Issues: 14.1R3 on page 218](#)
- [Resolved Issues: 14.1R2 on page 220](#)

Resolved Issues: 14.1R7

General Routing

- When a labeled BGP route resolves over a route with MPLS label (e.g. LDP/RSVP routes), after clearing the LDP/RSVP routes, in the short window before the LDP/RSVP routes restore, if the BGP routes resolves over a direct route (e.g. a one-hop LSP), the rpd process might crash. [PR1063796](#)

High Availability (HA) and Resiliency

- With NSR enabled on multiple RE system, when dynamic GRE tunnel is configured, performing RE switchover might causing rpd crash repeatedly on backup RE. [PR1130203](#)

Interfaces and Chassis

- During subscriber login/logout the below error log might occur on the device configured with GRES/NSR. /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV (pp0.1073751222) [PR1058958](#)

Multiprotocol Label Switching (MPLS)

- MPLS auto-bandwidth does not reset MAX Avg Bandwidth when overflow or underflow threshold limit is configured. It may lead to wrong bandwidth reservations occasionally. [PR954663](#)
- When an LSP is link-protected and has no-local-reversion configured, if the primary link (link1) is down and LSP on bypass (link2), then another link (link3) is brought up, before the LSP switch to link3, if link1 is enabled and link3 is disabled, the LSP will stuck in bypass LSP forever. This is a timing issue. [PR1091774](#)
- When multipoint LDP (M-LDP) in-band signaling is enabled to carry multicast traffic across an existing IP/MPLS backbone and routing process is enabled to use 64bit mode, the rpd might crash due to accessing an uninitialized local variables. [PR1118459](#)

Platform and Infrastructure

- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)

Routing Protocols

- From Junos release 14.1R1 or above, the rpd process might crash while executing CLI command "show isis backup spf results". [PR1037114](#)
- In multicast environment, when the RP is FHR (first hop router) and it has MSDP peers, when the rpf interface on RP changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)

Software Installation and Upgrade

- In certain conditions, when /var is not mounted from a persistent filesystem, executing a Junos upgrade will have unexpected results. This is caused by an inexact check of whether it is running from an Emergency VAR. [PR1112334](#)

VPNs

- For Layer 2 circuit, PTX3000 uses different VCCV (Virtual Circuit Connectivity Verification) BFD control packet format from that of MX and the other PTX platforms. PTX3000 negotiates Router-alert control channel type, and uses PW Associated Channel Header of Channel Type : 0x0021. However, MX and the other PTX platforms use the Channel Type is 0x0007 without IP/UDP headers. JUNOS takes the Channel-type 0x0007 as default. MX and the other PTX platforms work as expected. This is PTX3000 specific issue. [PR1116356](#)

Resolved Issues: 14.1R6

- [Forwarding and Sampling on page 214](#)
- [General Routing on page 214](#)
- [Interfaces and Chassis on page 215](#)
- [MPLS on page 215](#)
- [Network Management and Monitoring on page 215](#)

Forwarding and Sampling

- In PTX Series Carrier-Grade Service Engine (CSE) jflow solution environment, because the sampling process (sampled) may get into a continuous loop when handling asynchronous event (for example, aggregated tethered services interface flapping, or route update, or IFL/IFD update), the sampled may never come out of that loop which may result in high CPU usage (up to 90% sometimes). Because, sampled is not able to consume any states (such as route updates, interface updates) generated by kernel and this results in memory exhaustion and finally resulting in the router not making any updates and forcing a router reboot. [PR1092684](#)

General Routing

- On the PTX3000, if the distributed Bidirectional Forwarding Detection (BFD) session is configured over aggregated Ethernet interface with member links hosted on different FPCs, the BFD session might not come up due to the hello packets can not find the proper outgoing interface. [PR992652](#)
- The FPC on PTX Series router might crash and reboot when the Packet Forwarding Engine is handling a fatal error. When the error happened, "TQCHIP0: Fatal error pq_t_min_free_cnt is zero" log message will be seen. [PR1084259](#)
- On PTX Series platforms, some non-fatal interrupts (for example, CM cache or AQD interrupts) are logged as fatal interrupts. The following log messages will be shown on CM parity interrupt: fpc0 TQCHIP 0: CM parity Fatal interrupt, Interrupt status: 0x10 fpc0 CMSNG: Fatal ASIC error, chip TQ fpc0 TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180010 msec TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180005 msec [PR1089955](#)
- Entropy Label Capability is enabled by-default. On PTX Series transit LSRs that carry LSPs with Entropy Label Capability, packet loss can be observed due to data errors when one or more labeled route entries are not properly removed from the hash table (ie. following LSP optimization or MBB event) because the 'stale' entries are pointing to corrupted route memory. As a result, when the MPLS label that is associated with the 'stale' entry is re-used, data errors are seen for packets using the corresponding label. [PR1100637](#)
- On PTX Series platform, when pulling out FPC or SIB ungracefully (for example, pulling the line card out of the chassis unintentionally when the line card is carrying the traffic), there might be small probability that it can impact any of the FPCs with Grant Scheduler (GS) and Request Table (RT) fatal interrupt occurred. [PR1105079](#)

Interfaces and Chassis

- After removing a child link from aggregated Ethernet bundle, in the output of "show interface <AE> detail", the packet count on the remaining child link spikes, then if add back the previous child link, the count recover to normal. [PR1091425](#)

MPLS

- In the output of the cli command "traceroute mpls ldp", the addresses of the interfaces on transit PTX routers might be shown as "127.0.0.1". [PR1081274](#)

Network Management and Monitoring

- Due to inappropriate cleanup in async library, disabling multiple interfaces while SNMP is polling interface oids might cause mid2d process to crash. [PR1097165](#)

Resolved Issues: 14.1R5

- [General Routing on page 215](#)
- [Interfaces and Chassis on page 216](#)
- [MPLS on page 216](#)
- [Platform and Infrastructure on page 217](#)
- [Routing Protocols on page 217](#)

General Routing

- Prior to this fix "show interface diagnostics optics" command shows output for all four lanes for 10G ports of 48x10GE 12x40GE QSFP+ PIC. Normal behavior would be to display output for only the lane that the port belongs to. [PR959514](#)
- On PTX Series routers with MPLS environment (30k transit LSP), large number of MPLS interfaces (in this case, 200 interfaces) are configured with 0 or 1 MPLS labels. When these interfaces flap, the FPC kernel memory usage might leak. [PR995893](#)
- A reboot may be required when chassis is powered up first time. [PR1034662](#)
- When there is link/node protection/ECMP for RSVP/LDP transit or egress LSPs with huge scaling and continuous flapping of LSPs like auto-bandwidth case, traffic might get black-holed upon LSP re-optimizations. The issue would get triggered if the same unilist list-id (unilist list-id is a unique id for unilist nexthop) is allocated for two different unilist forwarding topologies. This situation arises when the unilist list-id wraps around after max value of 65535. After the wraparound, if there is long living list-id (which can be due to some node/link protected LSP that has not been re-optimized for a long time), the Packet Forwarding Engine assigns the same list-id during allocation (upon other LSP re-optimizations) and this will trigger the issue as the new unilist will be directed to incorrect interface. [PR1043747](#)
- On PTX Series or T Series platform running Junos OS release 12.1 or later, for interface connected via optical system like DWDM, when the interface is admin disabled, there might be a delay (300-400msec) for system to detect the event and during which time, traffic blackhole might be seen. Please note if disable the interface by breaking the Rx or Tx link, issue will not happen. [PR1043762](#)

- On PTX Series platform with one of the following protocols configured, flapping the protocols will trigger the Composite Next-hop change operation. In rare condition, since it is not properly programmed, the FPC might crash. This is a day-1 issue. - LDP - MPLS - Point-to-multipoint LSP - RSVP - Static LSPs [PR1045794](#)
- For PTX Series router, the unilist next-hop member will become 'replaced' status on Packet Forwarding Engine after interface flapping with ARP timeout. While the problem is happening, routing-table will display all right next-hop status but can not forward traffic since forwarding next-hop in Packet Forwarding Engine is in 'replaced' status and no longer active. [PR1046778](#)
- On PTX Series platform, non-revertive feature for clock synchronous sources does not work correctly. After deleting the primary clock and then adding it back, it will fallback to the primary clock but not stay in secondary. [PR1052549](#)
- When the port on 24x 10GE(LWO) SFP+ (which never went link up since the PIC is onlined) is configured as CLI loopback, the ports will receive framing error until the interface gets physically linked up (that is, with real fiber instead of CLI loop). There would be no problem in normal use. This is only seen in self-loopback testing with CLI loopback. [PR1057364](#)
- In LDP tunneling over single hop RSVP based LSP environment, after enabling "chained-composite-next-hop", the router may fail to create the chained composite next hops if the label value of VPN is equal with the label value of LDP. [PR1058146](#)
- On PTX Series routers, the interrupt-driven basis link down detection (an interrupt-driven link-down notification is generated to trigger locally attached systems to declare the interface down within a few milliseconds of failure) may fail after performing a unified in-service software upgrade (ISSU). The interrupt might got prevented after performing unified ISSU due to disable the interrupt registers before unified ISSU but never restored after. [PR1059098](#)
- On PTX Series routers, the interrupt-driven link down detection may stop working. When the line card is receiving multiple back-to-back fault in very short duration (no matter from remote or local), it may fail to detect all the received interrupts, and this failure may cause delay of the link down detection (for example, it may take PTX Series router ~300ms to make interface down). [PR1060279](#)
- The configured buffer-size will not take effect until either "transmit-rate" or "excess-rate" is configured. [PR1072179](#)

Interfaces and Chassis

- Configuring ODU FRR under otn-options for 2x100G DWDM PIC is an unsupported command on a PTX Series router, and incorrectly adding such a configuration could result in an FPC crash and restart. [PR1038551](#)

MPLS

- This is a regression issue related to a timing factor. When LDP session flaps, over which entropy label TLV or any unknown TLV is received, the LDP speaker might not send label withdraw for some prefixes to some neighbors. As a result, these neighbors will still use stale labels for the affected prefixes. [PR1062727](#)

- In Junos OS Release 14.1 and later, the "load-balance-label-capability" knob is introduced to enable the router to push and pop the load balancing label, which causes LDP and RSVP to advertise the entropy label TLV to neighboring routers. PTX Series router have the capability and it is reflected in their default forwarding-options configuration. However, there is a software defect in the way Entropy Label Capability (ELC) TLV is encoded in the LDP label mapping message. It might cause the LDP session between routers to go down. [PR1065338](#)

Platform and Infrastructure

- If the system has service-related configurations, error message generated by mountd might be seen: "can't delete exports for /packages/mnt/jbase: Bad address" [PR991814](#)
- In some rare conditions, setting up configuration access privileges, using the "allow-configuration-regexps" or "deny-configuration-regexps" statements will crash the management daemon (mgd), which serves a central role in the user-interface component of Junos OS. [PR1029384](#)

Routing Protocols

- In MPLS TE scenario, if IS-IS shortcuts for family inet6 are enabled, the LSP flapping might cause memory leak, which could result in traffic blackhole or FPC crash. [PR1049675](#)
- When running Simple Network Management Protocol (SNMP) polling to specific IS-IS Management Information Base (MIB) with invalid variable, it will cause routing protocol process (rpd) crash. [PR1060485](#)
- On PTX Series platform with transit BGP-LU chained composite next-hop configured, when advertising LDP routes via BGP labeled unicast (BGP-LU), if the LDP LSP itself is tunneled over an RSVP LSP, the rpd process might crash. Notes: The "set routing-options forwarding-table chained-composite-next-hop transit labeled-bgp" is enabled by default on PTX Series routers. [PR1065107](#)

Resolved Issues: 14.1R4

- [General Routing on page 217](#)
- [MPLS on page 218](#)
- [Network Management and Monitoring on page 218](#)
- [Routing Protocols on page 218](#)

General Routing

- On PTX5000, the packet drop is observed along with the parity error read from l3bnd_ht entry corresponding to certain addresses. With this SRAM parity error, ASIC will unconditionally drop the packet even if PTX Series does not use l3bnd_ht during lookup. The parity check for l3bnd_ht lookup for PTX5000 will be disabled to avoid the SRAM parity error and packet drop as a workaround. We also add new log message to report the counter value change for slu.hw_err trap count - TL[num]: SLU hw error count xxx (prev count yyy). [PR1012513](#)

- LACP on aggregated Ethernet interfaces does not currently support unified ISSU on PTX Series. A warning message is present before performing unified ISSU if LACP is configured. So the user can discontinue the unified ISSU process. [PR1018233](#)
- On PTX Series with equal-cost multipath (ECMP) route, bouncing the route next-hop interface hosted PIC, the Packet Forwarding Engine might get the route next-hop change message before the interface up message when the PIC is coming up, which results in the next-hop not installed in Packet Forwarding Engine leading to traffic black-holing. [PR1035893](#)

MPLS

- On P2MP MPLS LSP transit router with NSR enabled, when RSVP refresh reduction feature is enabled and LSP link protection is configured on all interfaces, slight P2MP traffic loss might be seen after the graceful Routing Engine switchover is done. [PR1023393](#)
- In MPLS traffic engineering with link or node protection enabled, after adding Shared Risk Link Group (SRLG) configuration, the bypass LSP might ignore the constraint and use an unexpected path. [PR1034636](#)

Network Management and Monitoring

- jnxcpic 380 and jnxcpic 381 definitions have been added in the "mib-jnx-chas-define" file from 14.1R4 release. [PR1036706](#)

Routing Protocols

- Do not use "ifconfig *interface name* down" within shell to bring an interface down. This command may cause unexpected behavior. Use the legitimate CLI configuration command "set interfaces *interface-name* disable" [PR1015736](#)
- With any single hop BFD session and MPLS OAM BFD session configured over same interface, when the interface is disabled and enabled back immediately (e.g. a delay of 10 sec between the two commit check ins), the single hop BFD session might get stuck into Init-Init state due to Down packet is received from other end for MPLS BFD session on the same interface might get demultiplexed to single hop BFD session wrongly. [PR1039149](#)

Resolved Issues: 14.1R3

- [General Routing on page 219](#)
- [Infrastructure on page 219](#)
- [Interfaces and Chassis on page 219](#)
- [MPLS on page 219](#)
- [Routing Protocols on page 220](#)
- [User Interface and Configuration on page 220](#)

General Routing

- On PTX Series platform, when receiving high rate ipv4/ipv6/mpls packets with TTL equals 1, the ICMP TTL expired messages are sent back to the sender not according with the ICMP rate limit settings. [PR893129](#)
- This PR fixes the issue where output ifIndex was being exported as 0. [PR964745](#)
- When "request system halt" is executed on the PTX Series router, the Routing Engine is halted, but the PTX Series router does not display Halt message on the CRAFT-Interface confirming that the system has halted. [PR971303](#)
- If Routing Engine based link protection is enabled on P2MP ingress LSPs in PTX Series and exit interfaces for P2MP LSP branches via ae bundles, packet might duplicate. [PR987005](#)
- On PTX Series routers with GRES configuration, the chassis daemon might crash when Routing Engine switchover is executed. [PR993857](#)
- Because of MCNH change from Release 13.3 to 14.1 and later, which used new FLOOD_MCNH to replace old MCNH_P2MP, while unified ISSU was upgrading, rpd would crash. [PR1000494](#)
- On PTX Series platform working as LSP ingress router, the MPLS auto-bandwidth feature might cause FPC to wedge condition with all interfaces down. [PR1005339](#)
- When large number of IGMP join packets are trying to reach the router, some IGMP packets might get dropped. [PR1007057](#)
- The problem is seen in PTX Series routers where the composite next hops are not observed, for a given VPN mpls route and hence the show route output command gives a truncated value which results in script failure. This may be due to default disabled l3vpn-cn timer in case of transit l3vpn router on PTX Series platform. [PR1007311](#)

Infrastructure

- "delete" or "deactivate" of apply-group defining the entire TACACS or RADIUS configuration configured under [edit system apply-group <>] does not take effect on commit. This could lead to TACACS or RADIUS-based authentication to still continue working despite removal (delete/deactivate) of configuration. [PR992837](#)

Interfaces and Chassis

- On PTX Series platform, CFP-100G-LR4 and CFP2-100G-LR4 optics report incorrect "Laser output power" values on all 4 lanes in **show interface diagnostics optics <intf>**. [PR1021541](#)

MPLS

- When issue "traceroute mpls rsvp lsp-name" from the MPLS LSP ingress node, if there are PTX Series routers on the LSP path, PTX Series would not list correct downstream router's IP in the TLV of the response packet. [PR966986](#)

- When a PTX Series router is at the merge-point (MP) of a bypass LSP, if MPLS explicit-null has been enabled on the router, and the loopback interface has not been configured under protocol RSVP, the bypass LSP might not work correctly. [PR1012221](#)

Routing Protocols

- Do not use "ifconfig <interface name> down" within shell to bring an interface down because it might cause unexpected behavior. Use the legitimate cli configuration command **set interfaces <interface name> disable**. [PR1015736](#)
- Establish two BFD sessions between two routers, one is single-hop BFD for directly connected interface and the other is multi-hop MPLS OAM BFD. If configuring the MPLS OAM on the same interface with single-hop BFD, when bringing down MPLS OAM from the ingress, it might result in the OAM BFD session deleted on ingress but it still receiving OAM BFD down packet from egress. Since there is no session matching this BFD packet, it does a normal look up and brings down the single-hop BFD session which is on the same interface. [PR1021287](#)

User Interface and Configuration

- Commit Error happens with load patch or load replace, which is while applying commit difference on backup Routing Engine as part of fast commit process. [PR1029474](#)

Resolved Issues: 14.1R2

- [General Routing on page 220](#)
- [Platform and Infrastructure on page 221](#)

General Routing

- On PTX Series platform, when receiving high rate ipv4/ipv6/mpls packets with TTL equals 1, the ICMP TTL expired messages are sent back to the sender not according with the ICMP rate limit settings. [PR893129](#)
- This PR fixes the issue where output ifIndex was being exported as 0. [PR964745](#)
- When "request system halt" is executed on the PTX Series router, the Routing Engine is halted, but the PTX Series router does not display Halt message on the CRAFT-Interface confirming that the system has halted. [PR971303](#)
- If Routing Engine based link protection is enabled on P2MP ingress LSPs in PTX Series and exit interfaces for P2MP LSP branches via ae bundles, packet might duplicate. [PR987005](#)
- On PTX Series routers with GRES configuration, the chassis daemon might crash when Routing Engine switchover is executed. [PR993857](#)
- Because of MCNH change from Release 13.3 to 14.1 and later, which used new FLOOD_MCNH to replace old MCNH_P2MP, while unified ISSU was upgrading, rpd would crash. [PR1000494](#)
- On PTX Series platform working as LSP ingress router, the MPLS auto-bandwidth feature might cause FPC to wedge condition with all interfaces down. [PR1005339](#)

- When large number of IGMP join packets are trying to reach the router, some IGMP packets might get dropped. [PR1007057](#)
- The problem is seen in PTX Series routers where the composite next hops are not observed, for a given VPN mpls route and hence the show route output command gives a truncated value which results in script failure. This may be due to default disabled l3vpn-cnh in case of transit l3vpn router on PTX Series platform. [PR1007311](#)

Platform and Infrastructure

- "delete" or "deactivate" of apply-group defining the entire TACACS or RADIUS configuration configured under [edit system apply-group <>] does not take effect on commit. This could lead to TACACS or RADIUS-based authentication to still continue working despite removal (delete/deactivate) of configuration. [PR992837](#)

Related Documentation

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Documentation Updates on page 221](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)
- [Product Compatibility on page 225](#)

Documentation Updates

There are no outstanding issues with the published documentation for Junos OS Release 14.1R7 for the PTX Series.

Related Documentation

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Resolved Issues on page 212](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)
- [Product Compatibility on page 225](#)

Migration, Upgrade, and Downgrade Instructions

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

- [Upgrading Using Unified ISSU on page 222](#)
- [Upgrading a Router with Redundant Routing Engines on page 222](#)
- [Basic Procedure for Upgrading to Release 14.1R6 on page 222](#)

Upgrading Using Unified 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 [High Availability Feature Guide for Routing Devices](#).



NOTE: Unified ISSU on the PTX5000 does not support upgrades from Junos OS Release 13.3 to Junos OS Release 14.1. Upgrading from Junos OS Release 13.3 to Junos OS Release 14.1 breaks the unified ISSU process.

Upgrading a Router with Redundant Routing Engines

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

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

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

Basic Procedure for Upgrading to Release 14.1R6

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



NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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



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



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.

The download and installation process for Junos OS Release 14.1 is different from previous Junos OS releases.

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:
<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 router.



NOTE: After you install a Junos OS Release 14.1 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.

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.

Customers in the United States and Canada, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-14.1R71-domestic-signed.tgz
```


All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-14.1R71-export-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**
 - **scp://hostname/pathname** (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 14.1 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.

Related Documentation

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Resolved Issues on page 212](#)
- [Documentation Updates on page 221](#)
- [Product Compatibility on page 225](#)

Product Compatibility

- [Hardware Compatibility on page 226](#)

Hardware Compatibility

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

To determine the features supported on PTX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

<http://pathfinder.juniper.net/feature-explorer/>

Related Documentation

- [New and Changed Features on page 201](#)
- [Changes in Behavior and Syntax on page 207](#)
- [Known Behavior on page 209](#)
- [Known Issues on page 209](#)
- [Resolved Issues on page 212](#)
- [Documentation Updates on page 221](#)
- [Migration, Upgrade, and Downgrade Instructions on page 222](#)

Finding More Information

For the latest, most complete information about known and resolved issues with Junos OS, see the Juniper Networks Problem Report Search application at:

<http://prsearch.juniper.net>.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at:

<http://pathfinder.juniper.net/feature-explorer/>.

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF format. Find Content Explorer at:

<http://www.juniper.net/techpubs/content-applications/content-explorer/>.

Documentation Feedback

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

- **Online feedback rating system**—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>.
- **E-mail**—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active 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: <http://kb.juniper.net/InfoCenter/>
- 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

16 March 2016—Revision 3, Junos OS Release 14.1R7— EX Series, M Series, MX Series, PTX Series, and T Series.

9 March 2016—Revision 2, Junos OS Release 14.1R7— EX Series, M Series, MX Series, PTX Series, and T Series.

2 March 2016—Revision 1, Junos OS Release 14.1R7— EX Series, M Series, MX Series, PTX Series, and T Series.

12 November 2015—Revision 3, Junos OS Release 14.1R6— EX Series, M Series, MX Series, PTX Series, and T Series.

5 November 2015—Revision 2, Junos OS Release 14.1R6— EX Series, M Series, MX Series, PTX Series, and T Series.

29 October 2015—Revision 1, Junos OS Release 14.1R6— EX Series, M Series, MX Series, PTX Series, and T Series.

26 August 2015—Revision 5, Junos OS Release 14.1R5— EX Series, M Series, MX Series, PTX Series, and T Series.

12 August 2015—Revision 4, Junos OS Release 14.1R5— EX Series, M Series, MX Series, PTX Series, and T Series.

24 July 2015—Revision 3, Junos OS Release 14.1R5— EX Series, M Series, MX Series, PTX Series, and T Series.

16 July 2015—Revision 2, Junos OS Release 14.1R5— EX Series, M Series, MX Series, PTX Series, and T Series.

9 July 2015—Revision 1, Junos OS Release 14.1R5— EX Series, M Series, MX Series, PTX Series, and T Series.

9 June 2015—Revision 8, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

26 May 2015—Revision 7, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

5 May 2015—Revision 6, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

9 April 2015—Revision 5, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

17 March 2015—Revision 4, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

24 February 2015—Revision 3, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

17 February 2015—Revision 2, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

10 February 2015—Revision 1, Junos OS Release 14.1R4— EX Series, M Series, MX Series, PTX Series, and T Series.

18 December 2014—Revision 4, Junos OS Release 14.1R3— EX Series, M Series, MX Series, PTX Series, and T Series.

11 December 2014—Revision 3, Junos OS Release 14.1R3— EX Series, M Series, MX Series, PTX Series, and T Series.

4 December 2014—Revision 2, Junos OS Release 14.1R3— EX Series, M Series, MX Series, PTX Series, and T Series.

20 November 2014—Revision 1, Junos OS Release 14.1R3— EX Series, M Series, MX Series, PTX Series, and T Series.

11 September 2014—Revision 5, Junos OS Release 14.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

25 August 2014—Revision 4, Junos OS Release 14.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

21 August 2014—Revision 3, Junos OS Release 14.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

14 August 2014—Revision 2, Junos OS Release 14.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

8 August 2014—Revision 1, Junos OS Release 14.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

24 July 2014—Revision 6, Junos OS Release 14.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

18 July 2014—Revision 5, Junos OS Release 14.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

15 July 2014—Revision 4, Junos OS Release 14.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

26 June 2014—Revision 3, Junos OS Release 14.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

19 June 2014—Revision 2, Junos OS Release 14.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

12 June 2014—Revision 1, Junos OS Release 14.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

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.