Release Notes: Junos® OS Release 14.1R4 for the EX Series, M Series, MX Series, PTX Series, and T Series

26 May 2015

Contents

Introduction

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

These release notes accompany Junos OS Release 14.1R4 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 14.1R4 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

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

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 one EX9200-SF module to an EX9200-SF2 module, the SF module types can co-exist in the switch *during* the upgrade. You must replace the second EX9200-SF module with another EX9200-SF2 module for normal switch operation.

Flow Monitoring

- **EX9200 Virtual Chassis support for inline flow monitoring (EX9200 Virtual Chassis)**---Starting in 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 the following support:
  - Active sampling and exporting of both IPv4 and IPv6 traffic flows
  - Sampling traffic flows in both 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)

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 this Junos OS release, 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 may 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 7
- Known Behavior on page 8
- Known Issues on page 8
- Resolved Issues on page 10
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 14
- Product Compatibility on page 15

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 14.1R4 for the EX Series.

- Dynamic Host Configuration Protocol
- 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.

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, as highlighted in the following sample:

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
```

```
Chassis                                 JN1221A03RFC      EX9204
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   CABC6474          EX9200-SF2
CB 1              REV 10   750-048307   CABH8948          EX9200-SF2
...
```

Related Documentation

- New and Changed Features on page 5
- Known Behavior on page 8
- Known Issues on page 8
- Resolved Issues on page 10
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 14
- Product Compatibility on page 15

Known Behavior

There are no changes in known behavior in Junos OS Release 14.1R4 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 5
- Changes in Behavior and Syntax on page 7
- Known Issues on page 8
- Resolved Issues on page 10
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 14
- Product Compatibility on page 15

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R5 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

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 jdhcpd 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 software 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

Network Management and Monitoring

- On EX9200 switches, the **ptopoConnLastVerifyTime** MIB might return the wrong 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 ASes. PR963565

Related Documentation

- New and Changed Features on page 5
- Changes in Behavior and Syntax on page 7
- Known Behavior on page 8
- Resolved Issues on page 10
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 14

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 14.1R5

Authentication and Access Control

- On EX9200 switches, the output for the ptopoConnRemotePort MIB might display the wrong 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

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
- If the **disable-logging** option is the only configured option under the **[edit system ddos-protection global]** hierarchy level, and this option is deleted, the kernel might generate a core file. PR1014219
- On EX9200 switches, recurring LMEM data errors might cause a chip wedge. PR1033660
- On EX9200 switches, a process that cores multiple times in a short period of time might not generate a core file. PR1058192

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 for all aggregated Ethernet interfaces and multi-chassis aggregated Ethernet interfaces. PR1034917

Layer 2 Features

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

Routing Protocols

- On EX9200 switches on which virtual private LAN service (VPLS) is enabled, if 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

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, service-provider configurations of Q-in-Q tunneling might not work on multichassis aggregated Ethernet (mc-ae) interfaces. PR973188
- 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 5
- Changes in Behavior and Syntax on page 7
- Known Behavior on page 8
- Known Issues on page 8
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 14
- Product Compatibility on page 15

Documentation Updates

There are no errata or changes in Junos OS Release 14.1R4 for the EX Series switches documentation.

Related Documentation

- New and Changed Features on page 5
- Changes in Behavior and Syntax on page 7
- Known Behavior on page 8
- Known Issues on page 8
- Resolved Issues on page 10
- Migration, Upgrade, and Downgrade Instructions on page 14
- Product Compatibility on page 15

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 14

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 5
- Changes in Behavior and Syntax on page 7
- Known Behavior on page 8
- Known Issues on page 8
- Resolved Issues on page 10
- Documentation Updates on page 14
- Product Compatibility on page 15

Product Compatibility

- Hardware Compatibility on page 15

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 5
- Changes in Behavior and Syntax on page 7

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 14.1 R4 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

This section describes the new features and enhancements to existing features in Junos OS Release 14.1 R4 for the M Series, MX Series, and T Series.

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, **PFEO** 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:

  - Forwarding capability of up to 130 Gbps per Packet Forwarding Engine
  - 100-Gigabit Ethernet interfaces
  - 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
  - Two separate slots for MICs (MIC6-10G and MIC6-100G-CXP)
  - Two Packet Forwarding Engines for each MIC slot
  - Intelligent oversubscription services

  [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)
  - Layer 3 routing protocols
  - MPLS
  - Multicast forwarding
  - Firewall filters and policers
  - Class of service (CoS)
  - Tunnel service
  - Interoperability with existing DPCs and MPCs
  - Internet Group Management Protocol (IGMP) snooping with bridging, integrated routing and bridging (IRB), or VPLS
  - Intelligent hierarchical policers

- Layer 2 trunk port
- MPLS-fast reroute (FRR) VPLS instance prioritization
- Precision Time Protocol (PTP) (IEEE 1588)
- Synchronous Ethernet

The following features are not supported on MPC6E:

- Fine-grained queuing and input queuing
- Unified in-service software upgrade (ISSU)
- Active flow monitoring and services
- 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 ppmd process automatically runs on the Routing Engine and the Packet Forwarding Engine. To disable ppmd on the Packet Forwarding Engine only, include the **no-delegate-processing** statement at the **[edit routing-options ppm]** hierarchy level. The ppmd 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 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. Counters can be configured 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)**—Class-of-service (CoS) hierarchical schedulers can be configured 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, a backup link can be configured to either 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 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 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)**—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)**—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 MPC5Eon 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, the base learning interface approach and the statistical approach can be used 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 such MACs can be configured to ignore the moves. The cooloff time and statistical approach wait time are used internally to find out the looped interface. The interface recovery time can be configured 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. The entropy label can be configured 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 (AE) 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 AE 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.]

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, TFPT, 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 softwire 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 softwire 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*** **softwire-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 softwire flows** output shows active and inactive softwire 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 interoperation using Cisco NAS Port Info AVP (MX Series)**—Starting in Junos OS Release 14.1, you can globally configure LAC interoperation 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 minimalized 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

- **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

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 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 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 13.3R3, 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, 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 softwire statistics** and **clear services inline softwire 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 sents 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.

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 Junos router 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.]

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 14.1R4. 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, 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 subsciber 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.

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
  - 8,000 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 (MX Series)**—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 ovsdb controller *ip-address*]** statement, and OVSDB-managed Virtual Extensible LANs (VXLANs) by using the **[edit bridge-domains *bridge-domain-name* vxlan ovsdb-managed]** or **[edit routing-instances *routing-instance-name* bridge-domain *bridge-domain-name* vxlan ovsdb-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 with MPC line card)**—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 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 in Junos OS Release 14.1 R4, 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 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 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 52
- Known Behavior on page 62
- Known Issues on page 63
- Resolved Issues on page 74
- Documentation Updates on page 109
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 127

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 14.1R4 for the M Series, MX Series, and T Series.

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 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.

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 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 13.2, 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 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.]

MPLS

- **Enhanced support for GRE interfaces for GMPLS (MX Series)**—Starting in Junos OS Release 12.3R7, 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 prior to 14.1, you can configure both fast reroute and node and link protection on the same LSP. 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.

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**—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.
- **Support for BFD for IS-IS IPv6 interfaces**—Starting in Junos OS Release 14.1R2, bidirectional forwarding detection (BFD) is supported for IS-IS IPv6 interfaces. Include the **bidirectional-forwarding-detection** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level. By default, multiple BFD sessions over a single adjacency for IPv4 and IPv6 interfaces that belong to the same IS-IS instance are not automatically created. To enable BFD on IPv4 and IPv6 interfaces configured on the same IS-IS

instance, you must also include the new **bfd-per-address-family** statement at the **[edit protocols isis interface *interface-name*]** hierarchy level. When BFD is enabled for both IPv4 and IPv6 interfaces in a single IS-IS instance, a BFD session is created for each protocol family interface. If either the IPv4 or IPv6 session fails, the adjacency is torn down.

[See Example: Configuring BFD for IS-IS.]

- **BGP hides a route received with a label block size greater than 256 (M Series, MX Series, and T Series)**— 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** 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.

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.

Subscriber Management and Services

NOTE: **Although present in the code, the subscriber management features are not supported in Junos OS Release 14.1R4. 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.

User Interface and Configuration

- **Configuring regular expressions (M Series, MX Series, and T Series)**—In all supported Junos OS releases, regular expressions can no longer be configured 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                    dscd      519     1
lsi.1048832        intf     0                    indr  1048574     4
                                4.4.3.2          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                    dscd      519     1
lsi.1048834        intf     0                    indr  1048574     4
                              4.4.3.2           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.

Related Documentation

- New and Changed Features on page 17
- Known Behavior on page 62
- Known Issues on page 63
- Resolved Issues on page 74
- Documentation Updates on page 109
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 127

Known Behavior

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 14.1R4 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.

- High Availability (HA) and Resiliency

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.

Related Documentation

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 52
- Known Issues on page 63
- Resolved Issues on page 74
- Documentation Updates on page 109
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 127

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R4 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.

Class of Service (CoS)

- COSD errors COSD_GENCFG_WRITE_FAILED: GENCFG write failed (op, minor_type) = (add, fixed classification) for tbl 3 if 1460 sp-3/2/0.16383 Reason: File exists are seen when MS PIC comes online. PR854047

- CoS relevant misconfiguration (e.g. configure classifier exp for LT interfaces implicitly using "interface all" way) might cause cosd crash, if cosd experiences multiple crashes within a short time, it might not be able to restart. PR969900
- Actual-data-rate-downstream for PPPoE subscriber sent in PADI is interpreted in bits instead of kbits. PR1026940

Forwarding and Sampling

- Accounting-data log file contains multiple header lines. PR881832
- In the PPPoE environment with Idle-Timeout attribute is configured. The PPPoE subscribers are terminated early before the Idle-Timeout expires. PR991251

General Routing

- **next-hop-group** knob is not supported under routing-instance hierarchy, but this knob is present under this hierarchy. This PR is opened to remove next-hop-group knob from routing-instance hierarchy. PR731264
- Changing the static route config from next-hop to qualified-next-hop might result in static route getting missed from routing table. Restarting routing process can bring back the routes but with rpd core. PR827727
- DPD may not work with link-type IPSec tunnels when NAT is present between the IPSec peers. PR895719
- Periodic "show subscribers" CLI requests during the GRES recovery (on a scaled system) might lead to spawning of too many subinfo processes. As a side effect, CoA requests might not be serviced while system is kept busy by subinfo processes due to authd might take long time to be recovered (it was observed that authd is not recovered after 1+ hours). PR915677
- Destination ERR alarm is not getting cleared even after FPC offlined. PR937862
- When BCM0 interface goes down, Routing Engine should switch over on M320. PR949517
- The SNMP Get, GetBulk, or GetNext request response for lldpPortConfigTable was not filtering-out the information of interfaces that are configured in the filter-interfaces statement at the [edit snmp] hierarchy level. The issue is resolved now. PR946975
- **show chassis fabric topology** display error when HSL2 link fault between F13 and F2S. PR962268
- On T4000 with Type-5 FPC (T4000-FPC5-3D), if a single request time out or occasional timeouts were seen over long period of time, the timeout error bit is not cleared correctly. This leads to destination be marked dead, and the traffic cannot flow from source Packet Forwarding Engine to destination Packet Forwarding Engine. PR963467
- When the size of apply-macro generated by op-script is equal to 1022 characters, the extensible subscriber services management daemon (essmd) subscribers might get stuck in "terminating" state. PR966764

- In scale DHCP subscribers scenario (e.g. 54K dual-stack DHCPv4/DHCPv6), graceful Routing Engine switchover is configured. If Routing Engine switchover occurs, after that execute the command "root@user> show dynamic-configuration" many times, large scale DHCP or DHCPv6 subscribers might be terminated. PR968021
- Autoheal denied reason may not be shown if CRC errors occurs on the same cable from F13 side more than once in an autoheal window and subsequently error is seen is again from LCC side PR973783
- In the dual Routing Engines scenario with 8K PPP dual stack subscribers. In rare condition, after Routing Engine switchover, some subscribers are stuck in terminating state forever. PR974300
- A PE device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE. The IGP instance running in the VRF on the PE may be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE. PR977945
- 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
- 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
- If a user configures an MX-VC member with member ID 2, the Virtual Chassis' master Routing Engine may eventually experience a kernel panic. PR989291
- 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 Engine. PR1002107
- When a static discard route is configured with no-install option but actual forwarding using different next hop, if egress sampling is enabled on the forwarding outgoing interface (OIF), traffic leaving that interface would have incorrect OIF on the flow records, resulting in unreliable flow records and incorrect billing. There is no traffic impact though. PR1002287
- A raw IP packet with invalid Memory Buffer (mbuf) length may trigger a kernel crash. The invalid mbuf length might be set by other daemons wrongly. PR1006320
- When the destinations are pointing to protocol next-hops as unilist type or IP forwarding next-hops as unilist, which in scenarios like using Loop-Free Alternate Routes for OSPF (LFA-OSPF) with link protection or MPLS FRR is enabled. If flapping the active interface very fast, especially interface comes back up before Kernel gets a chance to delete all the unilist next-hops, those unilist next-hops which have not been deleted yet would be re-used. As a result, the corresponding destinations are pointing to discard next-hop(s) or replaced next-hop(s) in Packet Forwarding Engine Jtree. The "discard" next-hop(s) causes traffic blackhole while the "replaced" next-hop(s) diverts traffic to other active next-hop(s) in the unlist. Those unilist next-hops which have been already deleted are safe and get updated accordingly. This is a day one timing issue. PR1016649

- On M Series, MX Series and T series platforms with DHCP relay configured, the router might keep filling a specific partition "/var/mfs/sdb" with files named log.XXXX and this would eventually cause DHCP relay fail. PR1017642
- During ISSU on MX-VC platform, there is a chance that the management process (mgd) on both VC-M and VC-B consider itself as protocol backup, resulting in a reboot of the entire VC. PR1020606
- 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
- MS-MPC crashed while the FTP ALG mode enabled. PR1029222
- 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
- CLNS ping fails for l3vpn over ethernet scenario. ISIS routes are being considered as ARP routes which leads to this problem. PR1041251
- For MLPPP interface on MX Series routers with MPCs or MICs, in some very rare conditions, the received fragmented packets might be dropped. PR1041412
- Queue stats on LSQ interfaces are not properly cleaned up when queuing enabled on the IFD and the queues hosted at IFD level. This happens when a subsequent delete & create of LSQ interface (not always though) - 14.1R4.10. PR1044340
- Once default generate route 0.0.0.0/0 is added, deleted or changed, the PFEMAN thread running on the MPC needs more than 600msec to program such changes. This is long enough to trigger LFM or BFD flap. Junos OS 13.3R2 or later is exposed to this symptom. PR1045828
- 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 out of order, packets re-ordering would occur. It reduces the application level throughput for any protocols running over TCP. PR1049613
- Output MTU counter shows incorrect data in the show pfe statistics traffic command output. PR1061111
- Ethernet frame loss measurement (ETH-LM) does not function on XL-based cards. PR1064994
- 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
- Traffic throughput when flowing from NEO to ALOHA is lesser when compared to Aloha to NEO This is observed in 14.1R4. PR1076009

- 'show interfaces queue **<ifl>**' stats are not correct with RLSQ warm-standby mode. Issue seen on NEO as well in 14.1R4.10. PR1082417
- Sparks-IPSec Performance drop while testing with 1 Next-hop style service-set and 10 sessions for PPS 600000, Framesize 1518 . PR1084376

Infrastructure

- Mirroring to next-hop-subgroup stops working, when there is a change in the next-hop-subgroup configuration. PR1049631

Interfaces and Chassis

- Packet Forwarding Engine continues to forward traffic to DHCP client on a demux interface when ae0 interface is down. In this scenario the AE interface bundle has 5 members and configured with minimum link value of 4. When 2 members are down, the ae0 interface also goes down, but Packet Forwarding Engine continues to forward traffic on other members for the demux interface. PR836846
- Demux Subscriber IFLs might show the interface as 'Hardware-Down' even though the underlying ae bundle and its member link show up. PR971272
- Configuring the PAP/CHAP local-name for a static PPPoE logical interface is ignored and the default "JUNOS" is used. This will cause the PAP/CHAP Authentication failure. PR978154
- 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 would not be changed. PR990018
- On MX Series platform, when an aggregate-ethernet bundle participating as L2 interface within bridge-domain goes down, below 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 2:0:0:0:0:8:a:0:0:0:0:0:8:4:0: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. PR990023

- In scenario that MX Series routers serve as a L2TP network server (LNS) and Point-to-Point Protocol (PPP) sessions through established L2TP tunnels, if malformed or incomplete PPPoE packet is received, the jpppd process (daemon which is responsible for PPP based protocol) might crash while stripping HDLC header. PR1002164
- On MX Series platform with large scale of PPPoE subscribers (more than 60k) connected, PPP client process (jpppd) might crash and generate core files when performing Routing Engine switchover. PR1018313
- Configuring ODU FRR on 2x100G DWDM under otn-options could result in a FPC core. PR1038551
- During the test of login/logout of couple of thousands of PPPoE subscibers, pppoed may crash by generating a core dump file. PR1041367
- Disabling demux IFL does not take it down. PR1043264
- clear interfaces interface-set statistics all fails due to memory limitation. PR1045683
- jpppd core dump at ../../../../src/bsd/lib/libthr/thread/thr_kern.c due to heap allocation failure. PR1047660
- On MX Series routers (platforms) with enhanced Switch Control Board (SCB), when the Fan tray is inserted or pulled out, the chassisd process might crash. PR1048021
- 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 was not done correctly due to which the Inherit group was stuck with priority of 0. PR1051135
- LT interface becomes unreachable because the next-hop type remains in hold state. PR105966
- 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. This issue is targeted to be fixed in Junos 14.1R5.. PR1060659
- RPD cores with 20 CFM session for interface in 10ms with AE configuration on the router. PR1051669
- When configuringsampling on an MPC2E-NG, or MPC3E-NG flows are not getting exported,nor are they expiring and remain active after traffic is stopped. PR1062399

J-Web

- PPPoE Logical Interfaces configuration page design need to change. PR493451
- 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 J-Web. 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

Layer 2 Features

- When toggling VLAN tagging type from "flexible-vlan-tagging" to "vlan-tagging" or vice versa, the integrated bridging and routing (IRB) MTU should be changed accordingly. However the IRB MTU is not re-computed in this case, which might lead to connectivity outage. PR928746
- On all Junos OS based platforms running as a ring protection link (RPL) owner, the convergence time may exceed 50ms when link failure occurs. And Ethernet Ring Protection (ERP) may fail if the control VLAN is replaced with a different VLAN at runtime. PR970265
- When the "filter-interfaces" is configured, if SNMP walk is executed, the l2cpd process might in a dead loop. At that moment, if the router receives ARP request packet, it might drop the packet due to high CPU utilization. So the router fails to learn MAC address, the traffic forwarding will be affected. PR982447
- 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
- "account-layer2-overhead" is not supported on 10x10GE MIC. PR1057463

MPLS

- RSVP graceful restart does not function for LSPs that have a forwarding adjacency (FA) label-switched path (LSP) as a next hop. PR60256
- For point-to-multipoint LSPs configured for VPLS, the "ping mpls" command reports 100 percent packet loss even though the VPLS connection is active. PR287990
- When RSVP label-switched-path (LSP) optimize is enabled, RSVP LSP might stay down after a graceful Routing Engine mastership switchover. To resolve the problem, the corresponding label-switched-path configuration needs to be deactivated, then, be activated again. PR1025413

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

Platform and Infrastructure

- 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 was synchronized instead of the configuration on the other Routing Engine. PR841087
- Wrong source IP is used when responding to traceroute in L3VPN setup. This is not an indication of traffic taking wrong link. PR883701
- Backing up the configuration with transfer-on-commit does not work in a MX-VC environment. PR947444
- On MX Series router with trio linecard or T4000 router with type5 FPC, the AE interface is lag-enhanced mode. When the unicast-nexthop via IRB is over an AE interface, the packets might not transported on IRB. PR961685
- After reboot the device, the interface rejects all packets. PR962782
- Certain combinations of Junos OS CLI commands and arguments have been found to be exploitable in a way that can allow root access to the operating system. This may allow any user with permissions to run these CLI commands the ability to achieve elevated privileges and gain complete control of the device. Refer to JSA10634 for more information. PR964860
- In multi-chassis platform, one of LCC's mastership change causes other LCC's SPARE-SIB's Active-LED to be set abnormally instead of "actual active plane's LED". There is no impact on operation, it is a cosmetic issue. * only if spare-SIB is SIB#0. For example, - SCC-RE0(M),RE1(B) | LCC0-RE0(M),RE1(B) | LCC1-RE0(M),RE1(B) - all-chassis SIB0 is spare status. - LCC0's mastership change makes the issue on LCC1. - LCC1's spare-SIB0's active LED to be set abnormally. PR972457
- 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
- XML traceroute does not display as-numbers. PR988727
- The overhead values need to be represented with 8 bits to cover the range "-120..124", but the microcode is only using the last 7 bits. PR1020446
- On MX Series 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. Note it happens only when the router is in "network-services enhanced-ip" mode. PR1024613

- Recurring LMEM data errors may cause Lu chip wedge. 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 behavior is resetting the PCP and DEI bit when the packet is routed. PR1036756
- When "forwarding-class-accounting" is enabled on the interface which is situated on MPC, issue CLI command "show interfaces forwarding-class-counters" will cause MPC crashing due to packet corruption. PR1042461
- Values for the input-traffic-control-profile statement get reset after enabling or disabling the traffic-manager statement with ingress-and-egress mode at the [edit chassis] hierarchy level. PR1052785
- 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

Routing Protocols

- It is necessary that the MSDP peer local-address matches the PIM RP address on routers that are RP. MSDP RPF check might fail in rare cases when both these addresses are not equal. PR35806
- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. PR51975
- 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
- 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 rare cases, rpd may write a core file with signature "rt_notbest_sanity: Path selection failure on ..." The core is 'soft', which means there should be no impact to traffic or routing protocols. PR946415
- In the BGP environment, a group has both IPv4 and IPv6 BGP peers. When all of the IPv4 peers flaps while none of the IPv6 peers flaps, there is a timing issue that if one

of the IPv4 peers comes up before inet.0 RIB is cleaned up, the routing protocol daemon (rpd) might crash. PR986272

- 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
- BFD session may reset on commit if version is configured. The adaptive RX interval gets set to 0 which results in the reset. protocols { bgp { bfd-liveness-detection { version 1; minimum-interval 1000; transmit-interval { minimum-interval 1000; } } } junos@junos> show bfd session extensive| no-more Detect Transmit Address State Interface Time Interval Multiplier 1.1.1.1 Up ae1.0 3.000 4.000 3 Client BGP, TX interval 4.000, RX interval 4.000 Session up time 00:00:06, previous down time 00:00:06 Local diagnostic None, remote diagnostic None Remote state Up, version 1 Session type: Single hop BFD Min async interval 4.000, min slow interval 4.000 Adaptive async TX interval 4.000, RX interval 0.000 Local min TX interval 4.000, minimum RX interval 4.000, multiplier 3 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3 Local discriminator 23, remote discriminator 1911224218 Echo mode disabled/inactive Session ID: 0x2a, update-adj 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
- Routing Protocol Daemon resets when static reverse-path forwarding (RPF) selection is configured and upstream interface in VRF routing instance disabled. PR1054913

Services Applications

- When you specify a standard application at the [edit security idp idp-policy *policy-name* rulebase-ips rule *rule-name* match application] hierarchy level, IDP does not detect the attack on the nonstandard port (for example, junos:ftp on port 85). Whether it is a custom or predefined application, the application name does not matter. IDP simply looks at the protocol and port from the application definition. Only when traffic matches the protocol and port does IDP try to match or detect against the associated attack. PR477748
- Same traffic stream clocks 2 counters in different service-filters. There is no workaround. PR706963
- During a specific scenario, and when AVP hiding is configured on L2TP Network Server (LNS), jl2tpd segmentation fault crash could happen on L2TP Access Concentrator (LAC). PR960107
- Message type for if_msg_ifl_channel_delete should be lower severity and not an error. PR965298
- In the L2TP scenario, when the username is more than 200 characters long, the L2TP daemon (jl2tpd) on LAC might crash. PR979047

- In L2TP scenario, when the LNS is flooded by high rate L2TP messages from LAC, the CPU on Routing Engine might keep too busy to bring up new sessions. PR990081
- The L2TP daemon (jl2tpd) might crash when the length field of the L2TP control message is set less than 12. PR998894
- 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 generate a core file. PR1021881
- 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 unable to login anymore. PR1025235

Subscriber Management and Services

- Access-request packet is not transmitted to RADIUS when Current field shows 4294967295. This happens when the Radius server service is stopped or when the route to Radius server is not available. PR1040240

User Interface and Configuration

- Selecting the Monitor port for any port in the Chassis Viewer page takes the user to the common Port Monitoring page instead of the corresponding Monitoring page of the selected port. PR446890
- User needs to wait until the page is completely loaded before navigating away from the current page. PR567756
- The J-Web interface allows the creation of duplicate term names in the Configure > Security > Filters > IPV4 Firewall Filters page. But the duplicate entry is not shown in the grid. There is no functionality impact on the J-Web interface. PR574525
- Using the Internet Explorer 7 browser, while deleting a user from the Configure > System Properties > User Management > Users page on the J-Web interface, the system is not showing warning message, whereas in the Firefox browser error messages are shown. PR595932
- If you access the J-Web interface using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. PR603669
- On the J-Web interface, next hop column in Monitor > Routing > Route Information displays only the interface address and the corresponding IP address is missing. The title of the first column displays "static route address" instead of "Destination Address". PR684552

VPNs

- When you modify the frame-relay-tcc statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level of a Layer 2 VPN, the connection for the second logical interface might not come up. As a workaround, restart the chassis process (chassisd) or reboot the router. PR32763
- BGP community 0xFF04 (65284) is a well-known community (NOPEER), but it is incorrectly displayed as "mvpn-mcast-rpt" in the CLI command "show route". This is a show command issue only. No operational mis-behavior will be observed on the router/network. PR479156
- In NG MVPN, after IPv6 VRF RP configuration change, we may hit IPv6 data loss for a short period of time. PR1049294
- RPD cores when deactivate logical router with NSR enabled & MVPN Routing instance configured. PR1059057

Related Documentation

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 52
- Known Behavior on page 62
- Resolved Issues on page 74
- Documentation Updates on page 109
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 127

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: 14.1R4 on page 74
- Resolved Issues: 14.1R3 on page 84
- Resolved Issues: 14.1R2 on page 100

Resolved Issues: 14.1R4

- Class of Service (CoS) on page 75
- General Routing on page 75
- High Availability (HA) and Resiliency on page 79
- Interfaces and Chassis on page 79
- Layer 2 Features on page 80
- MPLS on page 80

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
- 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 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
- 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 continously 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 RE 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
- 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 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
- 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
- Somtimes 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

- jnxpic 380 and jnxpic 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
- Trio-based line card 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 DSlite server does not take Softwire 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

- User Interface and Configuration on page 100
- VPNs on page 100

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 enchanced 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 origianl 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
- 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
- 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
- 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, l2 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 then +/-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 MVRF 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 labeld-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 immideate neighbors of Dut. Fix: added level info to get sysis_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
- Softwire tunnel count management is inconsistent and incorrect, thus the output of "show service softwire 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 assignement 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)

- Manually setting max-queues-per-interface to 4 on PB-4OC3-1OC12-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 commited, 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 enchanced 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 16xXE 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::vrrpAssoIpAddrRowStatus 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 MPLSLSPOCTETS, MPLSLSPPACKETS, MPLSLSPINFOOCTETS or MPLSLSPINFOPACKETS. 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 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 gets 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 17
- Changes in Behavior and Syntax on page 52
- Known Behavior on page 62
- Known Issues on page 63
- Resolved Issues on page 74
- Documentation Updates on page 109
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 127

Documentation Updates

This section lists the errata and changes in Junos OS Release 14.1R4 documentation for the M Series, MX Series, and T Series.

- Chassis-Level Feature Guide on page 110
- Ethernet Interfaces Feature Guide on page 110
- Firewall Filters Feature Guide for Routing Devices on page 110
- Interchassis Redundancy Using Virtual Chassis Feature Guide for MX Series Routers on page 110
- Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide on page 112
- Junos OS Administration Library for Routing Devices on page 113
- Layer 2 Configuration Guide, Bridging, Address Learning, and Forwarding on page 113
- Services Interfaces Configuration Guide on page 113
- Subscriber Management Network Access Feature Guide on page 116
- Subscriber Management Provisioning Guide on page 116
- System Log Messages Reference on page 116
- Traffic Sampling, Forwarding, and Monitoring Feature Guide for Routing Devices on page 116
- User Access and Authorization Feature Guide for Routing Devices on page 117
- VPLS Feature Guide for Routing Devices on page 117

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.

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.

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 that can be configured 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.

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.

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.

Related Documentation

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 52
- Known Behavior on page 62
- Known Issues on page 63
- Resolved Issues on page 74
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 127

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 118
- Upgrade and Downgrade Support Policy for Junos OS Releases on page 120
- Upgrading a Router with Redundant Routing Engines on page 120

Basic Procedure for Upgrading to Release 14.1

In order to upgrade to Junos OS 10.0 or later, you must be running Junos OS 9.0S2, 9.1S1, 9.2R4, 9.3R3, 9.4R3, 9.5R1, or later minor versions, or you must specify the **no-validate** option on the **request system software install** command.

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

NOTE: With Junos OS Release 9.0 and later, the compact flash disk memory requirement for Junos OS is 1 GB. For M7i and M10i routers with only 256 MB memory, see the Customer Support Center JTAC Technical Bulletin PSN-2007-10-001 at **https://www.juniper.net/alerts/viewalert.jsp?txtAlertNumber=PSN-2007-10-001&actionBtn=Search**

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

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

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the **juniper.conf** and **ssh** files) might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the Junos OS 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-14.1R41-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-14.1R41-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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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

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

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 adddress **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*.

Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR

Junos OS Release 9.3 introduced NSR support for PIM for IPv4 traffic. However, the following PIM features are not currently supported with NSR. The commit operation fails if the configuration includes both NSR and one or more of these features:

- Anycast RP
- Draft-Rosen multicast VPNs (MVPNs)
- Local RP
- Next-generation MVPNs with PIM provider tunnels
- PIM join load balancing

Junos OS Release 9.3 introduced a new configuration statement that disables NSR for PIM only, so that you can activate incompatible PIM features and continue to use NSR for the other protocols on the router: the **nonstop-routing disable** statement at the **[edit protocols pim]** hierarchy level. (Note that this statement disables NSR for all PIM features, not only incompatible features.)

If neither NSR nor PIM is enabled on the router to be upgraded or if one of the unsupported PIM features is enabled but NSR is not enabled, no additional steps are necessary and you can use the standard upgrade procedure described in other sections of these instructions. If NSR is enabled and no NSR-incompatible PIM features are enabled, use the standard reboot or ISSU procedures described in the other sections of these instructions.

Because the **nonstop-routing disable** statement was not available in Junos OS Release 9.2 and earlier, if both NSR and an incompatible PIM feature are enabled on a router to be upgraded from Junos OS Release 9.2 or earlier to a later release, you must disable PIM before the upgrade and reenable it after the router is running the upgraded Junos OS and you have entered the **nonstop-routing disable** statement. If your router is running Junos OS Release 9.3 or later, you can upgrade to a later release without disabling NSR or PIM—simply use the standard reboot or ISSU procedures described in the other sections of these instructions.

To disable and reenable PIM:

1. On the router running Junos OS Release 9.2 or earlier, enter configuration mode and disable PIM:

   ```
   [edit]

   user@host# deactivate protocols pim
   user@host# commit
   ```

2. Upgrade to Junos OS Release 9.3 or later software using the instructions appropriate for the router type. You can either use the standard procedure with reboot or use ISSU.
3. After the router reboots and is running the upgraded Junos OS, enter configuration mode, disable PIM NSR with the **nonstop-routing disable** statement, and then reenable PIM:

```
[edit]

user@host# set protocols pim nonstop-routing disable
user@host# activate protocols pim
user@host# commit
```

Downgrading from Release 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 17
- Changes in Behavior and Syntax on page 52
- Known Behavior on page 62
- Known Issues on page 63

Product Compatibility

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

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 14.1 R4 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

This section describes the new features and enhancements to existing features in Junos OS Release 14.1R4 for the PTX Series.

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 linerate 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)**—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 131 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.]

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 133
- Known Behavior on page 134
- Known Issues on page 135
- Resolved Issues on page 136
- Documentation Updates on page 140
- Migration, Upgrade, and Downgrade Instructions on page 140
- Product Compatibility on page 143

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 14.1R4 for the PTX Series.

- Interfaces and Chassis on page 134
- VPNs on page 134

Interfaces and Chassis

- 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 corresponding of the corresponding 10-Gigabit Ethernet port of the PIC. For information about the P2-10G-40G-QSFPP PIC, see P2-10G-40G-QSFPP PIC Overview.

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 l3pvn** 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 128
- Known Behavior on page 134
- Known Issues on page 135
- Resolved Issues on page 136
- Documentation Updates on page 140
- Migration, Upgrade, and Downgrade Instructions on page 140
- Product Compatibility on page 143

Known Behavior

There are no changes in known behavior in Junos OS Release 14.1R4 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 128
- Changes in Behavior and Syntax on page 133
- Known Issues on page 135
- Documentation Updates on page 140
- Resolved Issues on page 136

- Migration, Upgrade, and Downgrade Instructions on page 140
- Product Compatibility on page 143

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R4 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 135
- Platform and Infrastructure on page 135

General Routing

- PTX Series does not support the queuing PICs, but by default Junos OS will program chassis scheduler map which will generate the following logs: **"fpc2 COS(cos_chassis_scheduler_pre_add_action:2140): chassis schduler ipc received for non qpic ifd et-2/1/3 with index 131 /kernel: GENCFG: op 8 (COS BLOB) failed; err 5 (Invalid)Fix: Adding check to stop sending chassis scheduler map on PTX platform."** PR910985

Platform and Infrastructure

- In some rare conditions, setting up configuration access privileges using the "allow-configuration-regexps" or "deny-configuration-regexps statements" might crash the management daemon (mgd), which serves a central role in the user-interface component of Junos OS. PR1029384

Related Documentation

- New and Changed Features on page 128
- Changes in Behavior and Syntax on page 133
- Known Behavior on page 134
- Resolved Issues on page 136
- Documentation Updates on page 140
- Migration, Upgrade, and Downgrade Instructions on page 140
- Product Compatibility on page 143

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: 14.1R4

General Routing

- LACP on AE interfaces is currently ISSU unsupported on PTX Series. A warning message is present before performing ISSU if LACP configured. So the user can discontinue the ISSU process. PR1018233
- 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
- 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 a unexpected path. PR1034636

Network Management and Monitoring

- jnxpic 380 and jnxpic 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 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

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 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 128
- Changes in Behavior and Syntax on page 133

Documentation Updates

There are no outstanding issues with the published documentation for Junos OS Release 14.1R4 for the PTX Series.

Related Documentation

- New and Changed Features on page 128
- Changes in Behavior and Syntax on page 133
- Known Behavior on page 134
- Known Issues on page 135
- Resolved Issues on page 136
- Migration, Upgrade, and Downgrade Instructions on page 140
- Product Compatibility on page 143

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

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 14.1R4

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-14.1
    R41-domestic-signed.tgz
    ```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-14.1
R41-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 128
- Changes in Behavior and Syntax on page 133
- Known Behavior on page 134
- Known Issues on page 135
- Resolved Issues on page 136
- Documentation Updates on page 140
- Product Compatibility on page 143

Product Compatibility

- Hardware Compatibility on page 144

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 128
- Changes in Behavior and Syntax on page 133
- Known Behavior on page 134
- Known Issues on page 135
- Resolved Issues on page 136
- Documentation Updates on page 140
- Migration, Upgrade, and Downgrade Instructions on page 140

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

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.