Release Notes: Junos® OS Release 14.1R2 for the EX Series, M Series, MX Series, PTX Series, and T Series

11 September 2014

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 14.1R2 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 14.1R2 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 14.1R2 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.

Related Documentation

- Changes in Behavior and Syntax on page 5
- Known Behavior on page 6
- Known Issues on page 6
- *Resolved Issues*
- Documentation Updates on page 7
- Migration, Upgrade, and Downgrade Instructions on page 8
- Product Compatibility on page 9

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 14.1R2 for the EX Series.

- Platform and Infrastructure

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 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 4
- Known Behavior on page 6
- Known Issues on page 6
- *Resolved Issues*
- Documentation Updates on page 7
- Migration, Upgrade, and Downgrade Instructions on page 8
- Product Compatibility on page 9

Known Behavior

There are no changes in known behavior in Junos OS Release 14.1R2 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 4
- Changes in Behavior and Syntax on page 5
- Known Issues on page 6
- *Resolved Issues*
- Documentation Updates on page 7
- Migration, Upgrade, and Downgrade Instructions on page 8
- Product Compatibility on page 9

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R2 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.

- Interfaces and Chassis
- Routing Protocols

Interfaces and Chassis

- 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

- 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

Routing Protocols

- On EX9200 switches, if a session with an unconfigured peer is initiated, 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 4
- Changes in Behavior and Syntax on page 5
- Known Behavior on page 6
- *Resolved Issues*
- Documentation Updates on page 7
- Migration, Upgrade, and Downgrade Instructions on page 8
- Product Compatibility on page 9

Documentation Updates

There are no errata or changes in Junos OS Release 14.1R2 for the EX Series switches documentation.

Related Documentation

- New and Changed Features on page 4
- Changes in Behavior and Syntax on page 5
- Known Behavior on page 6
- Known Issues on page 6
- *Resolved Issues*
- Migration, Upgrade, and Downgrade Instructions on page 8
- Product Compatibility on page 9

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 8

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 4
- Changes in Behavior and Syntax on page 5
- Known Behavior on page 6
- Known Issues on page 6
- *Resolved Issues*
- Documentation Updates on page 7
- Product Compatibility on page 9

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 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 4
- Changes in Behavior and Syntax on page 5
- Known Behavior on page 6
- Known Issues on page 6
- *Resolved Issues*
- Documentation Updates on page 7
- Migration, Upgrade, and Downgrade Instructions on page 8

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 14.1R2 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 14.1R2 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 routers)**—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 routers)**—Junos OS 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 on MX240, MX480, MX960, MX2010, and MX2020 routers**—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 the MX960 router, FPC slot 0 and FPC slot 11 are not NEBS compliant beyond 104°F (40°C). This is a cooling restriction.

For more information about the supported and unsupported Junos OS software features for this MPC, see *Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5Es*.

- **Support for new fixed-configuration queuing MPC on MX240, MX480, MX960, MX2010, and MX2020 routers—**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.

  For more information about the supported and unsupported Junos OS software features for this MPC, see *Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC5Es*.

- **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 on MX2010 and MX2020 3D Universal Edge Routers**—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 routers)**—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**—Starting in Junos OS Release 14.1R1, 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.

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.

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.]

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.

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.

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 14.1R2. 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)**—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 Routers)**—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 via 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] statements.

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.]

Related Documentation

- Changes in Behavior and Syntax on page 41
- Known Behavior on page 50
- Known Issues on page 51
- Documentation Updates on page 69
- Migration, Upgrade, and Downgrade Instructions on page 75
- Product Compatibility on page 85

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 14.1R2 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 to 1400 bytes.

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

- **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.

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.**

Routing Policy and Firewall Filters

- **New firewall filter match condition supported on MPC line cards (MX Series)**—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 and should be removed 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 14.1R1, 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.]

- **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.]

Services Applications

- **Restrictions for maximum blocksize for NAT port block allocation**—Begining with Junos OS Release 14.1, the maximum blocksize for NAT port block allocation (PBA) is 32,000.

Subscriber Management and Services

NOTE: Although present in the code, the subscriber management features are not supported in Junos OS Release 14.1R2. 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 provide 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.

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 10
- Known Behavior on page 50
- Known Issues on page 51
- Documentation Updates on page 69
- Migration, Upgrade, and Downgrade Instructions on page 75
- Product Compatibility on page 85

Known Behavior

This section contains the known behaviors, system maximums, and limitations in hardware and software in known behavior in Junos OS Release 14.1R2 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.

Related Documentation

- New and Changed Features on page 10
- Changes in Behavior and Syntax on page 41
- Known Issues on page 51
- Documentation Updates on page 69
- Migration, Upgrade, and Downgrade Instructions on page 75
- Product Compatibility on page 85

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R2 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)
- Forwarding and Sampling
- General Routing
- Interfaces and Chassis
- J-Web
- Layer 2 Features
- Layer 2 Ethernet Services
- MPLS
- Network Management and Monitoring
- Operation, Administration, and Maintenance (OAM)
- Platform and Infrastructure
- Routing Protocols
- Services Applications
- Subscriber Access Management
- User Interface and Configuration
- VPNs

Class of Service (CoS)

- When the class of service daemon experiences multiple crashes within a short time, it might not be able to restart. PR969900

- SNMP get-request for OID jnxCosIngressQstatTxedBytes(ingress queue) might return the statistics/value of jnxCosQstatTxedBytes(egress queue). PR1011641

Forwarding and Sampling

- Accounting-data log file contains multiple header lines. PR881832
- In the PPPoE environment with the 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 the routing-instance hierarchy, but this knob is present under this hierarchy. This PR is opened to remove next-hop-group knob from the routing-instance hierarchy. PR731264
- 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
- Management interface(fxp) can transmit IPv6 packets to network interfaces. << topology example >> logical logical +---------+ +-----------+ +-----------+ | R1 |fe-1/3/0 fxp0| M10i |ge-0/0/2 ge-0/0/3| R2 | | |------------------| |--------------------| | | | | | | +---------+ +-----------+ +-----------+ .2 .1 .1 .2 2001:1::/64 2001:2::/64 <----------------> <----------------> ospf3 ospf3 ==>IPv6 packtes PR955132
- "show chassis fabric topology" displays error when HSL2 link fault between F13 and F2S. PR962268
- The amount of time that it takes for other Juniper Networks devices that function as hardware virtual tunnel endpoints (VTEPs) to learn a new MAC address after the first packet is sent from this MAC address is a maximum of 4.5 seconds. (The amount of time depends upon the server configuration on which VMware NSX is running.) During this time, traffic destined for this MAC address is flooded into the VXLAN. PR962945
- On T4000 with Type-5 FPC (T4000-FPC5-3D), if a single request timeout or occasional timeouts were seen over long period of time, the timeout error bit is not cleared correctly. This leads destination to be marked dead, and the traffic can't 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
- 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
- Autoheal denied reason may not be shown if CRC errors occur on the same cable from F13 side more than once in an autoheal window, subsequently the error is seen is again from LCC side. PR973783

- A sudden Routing Engine crash can be seen in an MPLS scenario where we attempt to read out of scope memory. PR988418
- An entity with a particular MAC address is moved so that its traffic is handled by a different Juniper Networks device that functions as a hardware virtual tunnel endpoint (VTEP). This MAC address is not learned by entities served by the new hardware VTEP until the hardware VTEP that previously handled its traffic ages out the MAC address. During this transitional period, traffic destined for this MAC address is dropped. PR988270
- On M7i/M10i with enchanced CFEB, M320 with E3-FPC, M120, and MX 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
- 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/MX/T Series routers (platforms) with Network Address Port Translation (NAPT) configuration, when the router receives the packet whose value of protocol filed 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
- RPD would push routes to meta-db on MS-PIC even when no-install knob is configured. This may have inconsistent OIF between actual forwarding next-hop versus next-hop on meta-db on MS-PIC resulting in incorrect Out interface on flow records. PR1002287
- Service PIC on MS-MPC card could generate a core file and restart on receiving a stray SIGQUIT signal due to it not handling the signal. With this fix we ignore SIGQUIT signal and avoid Service PIC restart. 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 work mgd to copy the files. This should not be a problem. During unified ISSU, when backup chassis switchover is done, the original Bm (new Bs) copies 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
- An overflowing number of error messages would be seen after loading a scaled CoS configuration. PR1006380
- After loading from scaled config to baseline configuration, get some error messages on the vty "Err] XQCHIP(46):XQ-chip[0]: RATE ran out of credits on nru interface" "Err] TNPC CM received unknown trigger (type Queue, id 1)". PR1006387
- Ingress queuing is not supported on MPC5 when OTN is enabled. Enabling ingress queuing with OTN would lead to line-card crash. PR1008569
- When the SIB plane state changed to fault state, it should read the FPGA for the power related info instead of reading from the cpld. PR1009402
- On TXP 100g brooklyn pic with gres enabled observed minimal ipv4 traffic loss around(0.04% to 0.05%) on aggregated 100g brooklyn pic during graceful switchover with all-chassis enabled on TXP platform. But the same feature is working fine in 40G brooklyn pic. PR1014420

Interfaces and Chassis

- For Automatic Protection Switching (APS) on SONET/SDH interfaces, there are no operational mode commands that display the presence of APS mode mismatches. An APS mode mismatch occurs when one side is configured to use bidirectional mode, and the other side is configured to use unidirectional mode. PR65800
- When the GE port is configured with WAN PHY mode, a "Zero length TLV" message might be reported from the port. This is a cosmetic issue. PR673937
- Demux Subscriber IFLs might show the interface as 'Hardware-Down' eventhough the underlying ae bundle and its member link shows up. PR971272
- PP0 static chap local-name is not used. PR978154
- After the connections with NSX controllers are disabled on a Juniper Networks device, interfaces that were configured to be managed by OVSDB continue passing traffic. PR980577
- IGMP joins do not work for PPP subscribers that are using MLPPP and LNS. PR1001214
- Fabric Blackholing logic recovery for certain cases will be done with different action (Phase 1/2/3) based on the problem. PR1009502
- In a corner case, it is possible for Packet Forwarding Engine to see IFL delete before NH delete and because of this the Packet Forwarding Engine can crash. This is a day one timing issue, but is being seen only now. There is no work around for this issue. The Packet Forwarding Engine will reboot after crashing and there will be a traffic loss on all interfaces on this Packet Forwarding Engine, till it comes up and syncs all FIB states from Routing Engine. PR1001953
- In a high-scale VPLS configuration, modification of a tunnel interface through a restart or reconfiguration may cause the packet processing engine to access an invalid interface, resulting in minor packet loss and logging of packet processing engine traps. Existing traffic flows on the Packet Forwarding Engine are not affected. The router recovers quickly and normal operation resumes with the new configuration. PR976972

J-Web

- Jweb 10.1 : PPPoE Logical Interfaces configuration page design need to change. PR493451
- JWEB:On https service chassis viewer is not launching on Internet Explorer. PR819717
- Basic value entry format error check is not present in Configure-->Security-->IPv6 Firewall Filters, but the same is present in IPv4 Firewall Filters. 1009173
- On Configure->System Properties->Management Access ->Certificates-> cli generated certificates are not reflecting at J-Web. PR915069

Layer 2 Features

- In BGP signaled VPLS/VPWS scenario, rpd process memory leak might occur when a group with wildcard configuration is applied to the routing instance. PR987727

Layer 2 Ethernet Services

- 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
- In MX 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
- On MX Seriesplatform with DHCP service enabled, issuing CLI command "show dhcp-security binding" might result in jdhcpd process crash. PR1007577

MPLS

- For point-to-multipoint LSPs configured for VPLS, the "ping mpls" command reports 100 percent packet loss even though the VPLS connection is active. PR287990

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
- 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 pfe statistics traffic counter is less then 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 the master Routing Engine to the 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 shoot up after every 20 seconds or so. PR988969

Operation, Administration, and Maintenance (OAM)

- For CFM down Mep CCM packets , Junos OS does not support setting 802.1p when DOWN MEP CFM session is configured on an interface which had interface vlan map operation and a vlan tag is pushed. PR1016071

Platform and Infrastructure

- Backing up the configuration with transfer-on-commit does not work in a MX-VC environment. PR947444
- After rebooting 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
- AS Number will be displayed when using XML RPC Traceroute. PR988727
- When we uninstall an SDK package, the configuration 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 file. Before the un-installation of the SDK package, the config statement [set jnx-ifinfo traceoptions flag all] should also get deleted from the configuration which is relevant to the package being deleted. PR992486
- On MX Series router with MPCs or 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 the "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 the wrong packets, and the packet forwarding would be affected. PR999649
- When Micro-BFD configurations are added after the ae bundle configuration, then micro-bfd session for all the member links remains in "Down" state. Below is the snippet as reference, when ae100 LACP state is "Disturbing", while micro-BFD session remain in "Down" state while on the other end the session would be in "Init" state. user@ndoeA> show lacp interfaces ae100 Aggregated interface: ae100 LACP state:

Role Exp Def Dist Col Syn Aggr Timeout Activity xe-0/3/0 Actor No No Yes Yes Yes Yes Fast Active xe-0/3/0 Partner No No Yes Yes Yes Yes Fast Active xe-0/3/1 Actor No No Yes Yes Yes Yes Fast Active xe-0/3/1 Partner No No Yes Yes Yes Yes Fast Active LACP protocol: Receive State Transmit State Mux State xe-0/3/0 Current Fast periodic Collecting distributing xe-0/3/1 Current Fast periodic Collecting distributing user@ndoeA> show bfd session address 10.10.100.145 Detect Transmit Address State Interface Time Interval Multiplier 10.10.100.145 Down xe-0/3/0 0.000 1.000 3 10.10.100.145 Down xe-0/3/1 0.000 1.000 3 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

Routing Protocols

- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. PR51975
- 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 file. The problem is with route-ordering. And it is auto-corrected after collecting this soft-assert-coredump, without any impact to traffic/service. PR815146
- The routing-protocol process(RPD) may crash with core-dump in 12.2R1 or later releases. This is due a BFD-Triggered Local Repair for Rapid Convergence feature introduced in 12.2R1. The fix for the issue has been added in 14.2 and later the fix will be added in 14.1 as well. PR926188
- 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
- In the P2MP environment with OSPF adjacency established, one router's time is set to an earlier date than another router. OSPF adjacency might not come up when one router goes down and comes up. PR991540

Services Applications

- When you specify a standard application at the [edit security idp idp-policy rulebase-ips rule 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
- 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
- If a destination-prefix or source-prefix is used like the following example, the 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 52 characters in length, it will get truncated, if the truncated jpool name gets overlapped with an other generated jpool name, it will lead to 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 { 10.20.20.1/32; } } 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 > 52 characters. Second jpool = _jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_LONG_NAME_2 > 52 characters. The resulted jpool "_jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_" will be used incorrectly in both terms. PR973465
- In the L2TP scenario, when the username is more than 200 characters long, the L2TP daemon (jl2tpd) on LAC might crash. PR979047
- The L2TP daemon (jl2tpd) might crash when the length field of the L2TP control message is set less than 12. PR998894
- The following messages are being logged at ERR not DEBUG severity: Jun 16 08:22:23.694 ROUTER-RE0 mspd[3618]: mspd: No member config Jun 16 08:22:23.695 ROUTER-RE0 mspd[3618]: mspd: Building package info This PR sets the correct severity. PR1003640

Subscriber Access Management

- MIB entries for jnxUserAAAAccessPoolRoutingInstance might not appear after deleting and re-adding an assignment pool under a routing instance. PR998967

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
- On configure->clitools->point and click->system->advanced->deletion of saved core context on "No" option is not happening at J-Web. PR888714

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

Related Documentation

- New and Changed Features on page 10
- Changes in Behavior and Syntax on page 41

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

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 recover 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. The limitation will be addressed in 14.1R2 release. 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 laser. 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 ISSU upgrading there would be a RPD crash happening. 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 validatio 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 is supported in Junos OS for the MD + MA name. Junos OS needs to support 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 equipment fail to response Topology Change Notification ACK packet when interoperates with Cisco equipment. After the fix, Juniper 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 multi homing (FEC 129) is 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 a 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 roouters 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 routesr 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 MX240/480/960 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 router with trio linecard 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 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 utilisation 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/MX/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 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 10
- Changes in Behavior and Syntax on page 41
- Known Behavior on page 50
- Known Issues on page 51
- Documentation Updates on page 69
- Migration, Upgrade, and Downgrade Instructions on page 75

Documentation Updates

This section lists the errata and changes in Junos OS Release 14.1R2 documentation for the M Series, MX Series, and T Series.

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 decapsulation 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 Concentator.

- 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 Overiew" 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 Overiew" 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.

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 enabled (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.

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*)

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 an DP, 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 10
- Changes in Behavior and Syntax on page 41
- Known Behavior on page 50
- Known Issues on page 51
- Migration, Upgrade, and Downgrade Instructions on page 75
- Product Compatibility on page 85

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 76
- Upgrade and Downgrade Support Policy for Junos OS Releases on page 78
- Upgrading a Router with Redundant Routing Engines on page 78

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-14.1R21-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-14.1R21-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 10
- Changes in Behavior and Syntax on page 41
- Known Behavior on page 50
- Known Issues on page 51

- Documentation Updates on page 69
- Product Compatibility on page 85

Product Compatibility

- Hardware Compatibility on page 85

Hardware Compatibility

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

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

Related Documentation

- New and Changed Features on page 10
- Changes in Behavior and Syntax on page 41
- Known Behavior on page 50
- Known Issues on page 51
- Documentation Updates on page 69
- Migration, Upgrade, and Downgrade Instructions on page 75

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 14.1R2 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 14.1R2 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 L4 optics.
- **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 89 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 91
- Known Behavior on page 92
- Known Issues on page 92
- Documentation Updates on page 94
- Migration, Upgrade, and Downgrade Instructions on page 95
- Product Compatibility on page 98

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 14.1R2 for the PTX Series.

- VPNs on page 92

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 router. 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 86
- Known Behavior on page 92
- Known Issues on page 92
- Documentation Updates on page 94
- Migration, Upgrade, and Downgrade Instructions on page 95
- Product Compatibility on page 98

Known Behavior

There are no changes in known behavior in Junos OS Release 14.1R2 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 86
- Changes in Behavior and Syntax on page 91
- Known Issues on page 92
- Documentation Updates on page 94
- Migration, Upgrade, and Downgrade Instructions on page 95
- Product Compatibility on page 98

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 14.1R2 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.

- MPLS
- VPNs

MPLS

- When issuing "traceroute mpls rsvp lsp-name" from the MPLS LSP ingress node, if there are PTX Series routers on the LSP path, PTX would not list the correct downstream router's IP in the TLV of the response packet. PR966986

VPNs

- The problem is seen in PTX Series routers where the composite nexthops 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 platform. If Resync blob is not set, RPD will create indirect nexthop for transit route on PE-PE connection network on PTX. If Resync blob is set, RPD will create composite nexthop for transit route on PE-PE connection network on PTX. Using composite nexthop (cnh) can help scaled network. However, either indirect (inh) or composite nexthops work properly in control and forwarding planes. PR1007311

Related Documentation

- New and Changed Features on page 86
- Changes in Behavior and Syntax on page 91
- Known Behavior on page 92
- Documentation Updates on page 94
- Migration, Upgrade, and Downgrade Instructions on page 95
- Product Compatibility on page 98

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.

- General Routing
- Platform and Infrastructure

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 does not display Halt message on the CRAFT-Interface confirming that the system has halted. PR971303
- On PTX Series routers with GRES configuration, the chassis daemon might crash when Routing Engine switchover is executed. PR993857

- 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 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
- 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
- The problem is seen in PTX Series routers where the composite nexthops 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 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 86
- Changes in Behavior and Syntax on page 91
- Known Behavior on page 92
- Known Issues on page 92
- Documentation Updates on page 94
- Migration, Upgrade, and Downgrade Instructions on page 95
- Product Compatibility on page 98

Documentation Updates

There are no outstanding issues with the published documentation for Junos OS Release 14.1R2 for the PTX Series.

Related Documentation

- New and Changed Features on page 86
- Changes in Behavior and Syntax on page 91
- Known Behavior on page 92
- Known Issues on page 92
- Migration, Upgrade, and Downgrade Instructions on page 95
- Product Compatibility on page 98

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 14.1R2

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-14.1R21-domestic-signed.tgz
    ```

All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-14.1R21-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 86
- Changes in Behavior and Syntax on page 91
- Known Behavior on page 92
- Known Issues on page 92
- Documentation Updates on page 94
- Product Compatibility on page 98

Product Compatibility

- Hardware Compatibility on page 99

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 86
- Changes in Behavior and Syntax on page 91
- Known Behavior on page 92
- Known Issues on page 92
- Documentation Updates on page 94
- Migration, Upgrade, and Downgrade Instructions on page 95

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

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.