

Junos[®] OS 12.3 Release Notes

Release 12.3R11
15 October 2015
Revision 3

These release notes accompany Release 12.3R11 of the Junos operating system (Junos OS). They describe device documentation and known problems with the software. Junos OS runs on all Juniper Networks ACX Series Universal Access Routers, EX Series Ethernet Switches, M Series, MX Series, and T Series routing platforms, and PTX Series Packet Transport Routers.

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

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

Contents

Junos OS Release Notes for ACX Series Routers	6
New Features in Junos OS Release 12.3 for ACX Series Routers	6
Hardware	7
Firewall Filters	7
Interfaces and Chassis	8
Layer 2 and Layer 3 Protocols	9
Routing Protocols	9
Time Division Multiplexing (TDM)	11
Timing and Synchronization	12
Changes in Default Behavior and Syntax in Junos OS Release 12.3 for ACX	
Series Routers	14
Interfaces and Chassis	14
IPv6	14
Junos OS XML API and Scripting	14
Network Management and Monitoring	15
Security	15
Errata and Changes in Documentation for Junos OS Release 12.3 for ACX	
Series Routers	15
Errata	15
Known Limitations in Junos OS Release 12.3 for ACX Series Routers	18
Class of Service	18
Firewall Filters	18

Interfaces and Chassis	19
MPLS	20
Outstanding Issues in Junos OS Release 12.3 for ACX Series Routers	20
General Routing	20
Resolved Issues in Junos OS Release 12.3 for ACX Series Routers	20
Resolved Issues	21
Previous Releases	21
Upgrade and Downgrade Instructions for Junos OS Release 12.3 for ACX	
Series Routers	23
Basic Procedure for Upgrading to Release 12.3	23
Upgrade and Downgrade Support Policy for Junos OS Releases	26
Downgrading from Release 12.3	26
Junos OS Release Notes for EX Series Switches	28
New Features in Junos OS Release 12.3 for EX Series Switches	28
Hardware	29
Access Control and Port Security	30
Class of Service (CoS)	30
Converged Networks (LAN and SAN)	31
Enhanced Layer 2 Software (ELS) on EX9200 Switches	32
Ethernet Switching and Spanning Trees	32
Firewall Filters and Routing Policy	33
High Availability	34
Infrastructure	34
Interfaces and Chassis	35
IPv6	36
J-Web Interface	37
Layer 2 and Layer 3 Protocols	37
Management and RMON	37
MPLS	38
Virtual Chassis	38
Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX	
Series Switches	39
Authentication and Access Control	40
Dynamic Host Configuration Protocol	40
Layer 2 Features	40
Hardware	40
High Availability	40
Infrastructure	40
Interfaces	41
Network Management and Monitoring	41
Limitations in Junos OS Release 12.3 for EX Series Switches	42
Ethernet Switching and Spanning Trees	42
Firewall Filters	42
Hardware	42
High Availability	43
Infrastructure	44
Interfaces	45
J-Web Interface	46
Layer 2 and Layer 3 Protocols	48

Management and RMON	48
Multicast Protocols	49
Software Installation and Upgrade	49
Virtual Chassis	49
Outstanding Issues in Junos OS Release 12.3 for EX Series Switches	50
Access Control and Port Security	51
Ethernet Switching and Spanning Trees	51
High Availability	51
Infrastructure	51
Interfaces	52
J-Web Interface	52
Layer 2 and Layer 3 Protocols	53
Multicast Protocols	53
Network Management and Monitoring	54
Routing Policy and Firewall Filters	54
Software Upgrade and Installation	54
Virtual Chassis	54
Resolved Issues in Junos OS Release 12.3 for EX Series Switches	54
Issues Resolved in Release 12.3R1	55
Issues Resolved in Release 12.3R2	65
Issues Resolved in Release 12.3R3	67
Issues Resolved in Release 12.3R4	72
Issues Resolved in Release 12.3R5	74
Issues Resolved in Release 12.3R6	77
Issues Resolved in Release 12.3R7	81
Issues Resolved in Release 12.3R8	84
Issues Resolved in Release 12.3R9	87
Issues Resolved in Release 12.3R10	92
Issues Resolved in Release 12.3R11	93
Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches	95
Changes to Junos OS for EX Series Switches Documentation	95
Errata	96
Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches	97
Upgrade and Downgrade Support Policy for Junos OS Releases	97
Upgrading EX Series Switches Using NSSU	98
Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations	99
Upgrading from Junos OS Release 10.4R3 or Later	99
Upgrading from Junos OS Release 10.4R2 or Earlier	101
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers	102
New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers	102
Hardware	103
Class of Service	104
Forwarding and Sampling	110
High Availability (HA) and Resiliency	110

Interfaces and Chassis	111
Junos OS XML API and Scripting	130
Layer 2 Features	131
Layer 2 Tunneling Protocol	133
MPLS	134
Multicast	136
Power Management	139
Routing Policy and Firewall Filters	139
Routing Protocols	141
Security	142
Subscriber Access Management	142
System Logging	153
User Interface and Configuration	154
VPLS	157
VPNs	160
Changes in Default Behavior and Syntax, and for Future Releases in Junos	
OS Release 12.3 for M Series, MX Series, and T Series Routers	162
Changes in Default Behavior and Syntax	162
Changes Planned for Future Releases	180
Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T	
Series Routers	181
Class of Service (CoS)	182
MPLS	183
Routing Policy and Firewall Filters	183
Software Installation and Upgrade	183
Errata and Changes in Documentation for Junos OS Release 12.3 for M Series,	
MX Series, and T Series Routers	184
Errata	184
Changes to the Junos OS Documentation Set	216
Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T	
Series Routers	216
Authentication and Access Control	217
Class of Service (CoS)	217
Forwarding and Sampling	218
General Routing	218
High Availability (HA) and Resiliency	225
Infrastructure	226
Interfaces and Chassis	226
Layer 2 Features	230
MPLS	231
Multicast	233
Network Management and Monitoring	233
Platform and Infrastructure	234
Routing Policy and Firewall Filters	238
Routing Protocols	238
Services Applications	241
Software Installation and Upgrade	245
Subscriber Access Management	245
User Interface and Configuration	245

VPNs	246
Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T	
Series Routers	248
Resolved Issues	249
Previous Releases	251
Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series,	
MX Series, and T Series Routers	357
Basic Procedure for Upgrading to Release 12.3	357
Upgrade and Downgrade Support Policy for Junos OS Releases	359
Upgrading a Router with Redundant Routing Engines	359
Upgrading Juniper Network Routers Running Draft-Rosen Multicast	
VPN to Junos OS Release 10.1	360
Upgrading the Software for a Routing Matrix	361
Upgrading Using Unified ISSU	362
Downgrading from Release 12.3	362
Junos OS Release Notes for PTX Series Packet Transport Routers	364
New Features in Junos OS Release 12.3 for PTX Series Packet Transport	
Routers	364
Hardware	364
Firewall Filters	366
Interfaces and Chassis	366
Network Management and Monitoring	367
User Interface and Configuration	368
Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX	
Series Packet Transport Routers	368
Changes in Default Behavior and Syntax	368
Errata and Changes in Documentation for Junos OS Release 12.3 for PTX	
Series Packet Transport Routers	370
Errata	370
Network Management and Monitoring	370
Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport	
Routers	371
Class of Service (CoS)	371
General Routing	371
General Routing	372
Interfaces and Chassis	373
MPLS	373
Network Management and Monitoring	373
Routing Protocols	373
Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport	
Routers	374
Current Release	374
Previous Releases	374
Junos OS Documentation and Release Notes	386
Documentation Feedback	386
Requesting Technical Support	386
Revision History	388

Junos OS Release Notes for ACX Series Routers

- [New Features in Junos OS Release 12.3 for ACX Series Routers on page 6](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for ACX Series Routers on page 14](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for ACX Series Routers on page 15](#)
- [Known Limitations in Junos OS Release 12.3 for ACX Series Routers on page 18](#)
- [Outstanding Issues in Junos OS Release 12.3 for ACX Series Routers on page 20](#)
- [Resolved Issues in Junos OS Release 12.3 for ACX Series Routers on page 20](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for ACX Series Routers on page 23](#)

New Features in Junos OS Release 12.3 for ACX Series Routers

Powered by Junos OS, the ACX Series Universal Access Routers provide superior management for rapid provisioning to the access network. They are designed to support residential, mobile, and business access. The ACX Series routers include the ACX1000, the ACX1100, the ACX2000, the ACX2200, and the ACX4000 routers.

The following are key features of the ACX Series routers:

- High performance up to 10 Gigabit Ethernet capable
- Seamless MPLS traffic engineering for optimal paths and per-customer quality of service in the access layer
- Built-in Precision Timing Protocol (PTP) and Synchronized Ethernet (SyncE) to eliminate dropped calls and data retransmissions
- Environmentally hardened with 65 W Power over Ethernet (PoE)

The following features have been added to Junos OS Release 12.3 for the ACX Series Universal Access Routers. Following the description is the title of the manual or manuals to consult for further information:

- [Hardware on page 7](#)
- [Firewall Filters on page 7](#)
- [Interfaces and Chassis on page 8](#)
- [Layer 2 and Layer 3 Protocols on page 9](#)
- [Routing Protocols on page 9](#)
- [Time Division Multiplexing \(TDM\) on page 11](#)
- [Timing and Synchronization on page 12](#)

Hardware

- **New ACX4000 Universal Access Router**—Starting in Release 12.3, Junos OS supports the ACX4000 router. This router enables a wide range of business and residential applications and services, including microwave cell site aggregation, MSO mobile backhaul service cell site deployment, and service provider or operator cell site deployment.

The ACX4000 supports use of either four RJ-45 ports or four Gigabit Ethernet SFP transceivers. The ACX4000 also contains an additional two PoE ports, two Gigabit Ethernet SFPs, and two 10-Gigabit Ethernet SFP+ transceivers. The router has two dedicated slots for MICs. For a list of the supported MICs, see the *ACX4000 Universal Access Router MIC Guide*.

[ACX4000 Hardware Guide]

Firewall Filters

- **Filter-based forwarding for routing instances**—For IPv4 traffic only, you can use stateless firewall filters in routing instances to control how packets travel in a network. This is called filter-based forwarding.

You can define a firewall filtering term that directs matching packets to a specified routing instance. This type of filtering can be configured to route specific types of traffic through a firewall or other security device before the traffic continues on its path. To configure a stateless firewall filter to direct traffic to a routing instance, configure a term with the **routing-instance *routing-instance-name*** terminating action at the **[edit firewall family inet filter *filter-name* term *term-name* then]** hierarchy level to specify the routing instance to which matching packets will be forwarded. To configure the filter to direct traffic to the master routing instance, use the **routing-instance default** statement at the **[edit firewall family inet filter *filter-name* term *term-name* then]** hierarchy level.

[ACX Series Universal Access Router Configuration Guide]

- **Forwarding table filters for routing instances**—Forwarding table filter is a mechanism by which all the packets forwarded by a certain forwarding table are subjected to filtering and if a packet matches the filter condition, the configured action is applied on the packet. You can use the forwarding table filter mechanism to apply a filter on all interfaces associated with a single routing instance with a simple configuration. You can apply a forwarding table filter to a routing instance of type forwarding and also to the default routing instance **inet.0**. To configure a forwarding table filter, include the **filter *filter-name*** statement at the **[edit firewall family inet]** hierarchy level.

[ACX Series Universal Access Router Configuration Guide]

Interfaces and Chassis

- **New Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (ACX-MIC-4COC3-1COC12CE) on ACX Series Universal Access Routers**—Starting with Junos OS Release 12.3, a new MIC, Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (ACX-MIC-4COC3-1COC12CE), is supported on ACX Series Universal Access Routers.
- **Support for 6xGE MIC on ACX4000 Universal Access Router**—The ACX4000 now supports 6xGE MICs. The 6xGE MIC features six tri-speed (10/100/1000Gbps) Ethernet ports. Each port can be configured to operate in either RJ45 or SFP mode.
- **Junos OS support for chassis management (ACX4000)**—The following CLI operational mode commands are supported on the ACX4000:

Show commands:

- `show chassis alarms`
- `show chassis craft-interface`
- `show chassis environment`
- `show chassis environment pem`
- `show chassis fan`
- `show chassis firmware`
- `show chassis fpc pic-status`
- `show chassis hardware (clei-models | detail | extensive | models)`
- `show chassis mac-addresses`
- `show chassis pic fpc-slot fpc-slot pic-slot pic slot`
- `show chassis routing-engine`

Restart command:

- `restart chassis-control (gracefully | immediately | soft)`

Request commands:

- `request chassis feb restart slot slot-number`
- `request chassis mic mic-slot mic-slot fpc-slot fpc-slot (offline | online)`
- `request chassis pic offline fpc-slot fpc-slot pic-slot pic-slot`

[See the *ACX Series Universal Access Router Configuration Guide* and the *System Basics: Chassis-Level Features Configuration Guide*.]

- **User-defined alarms**—On an ACX Series router, the alarm contact port (labeled ALARM) provides four user-defined input ports and two user-defined output ports. Whenever a system condition occurs—such as a rise in temperature, and depending on the configuration, the input or output port is activated. The following configuration is supported for user-defined alarms:


```
[edit chassis alarm relay]
input {
  port port-number {
    mode (close | open);
    trigger (ignore | red | yellow;
  }
}
output {
  port port-number {
    input-relay input-relay {
      port port-number;
    }
    mode (close | open);
    temperature;
  }
}
```

To view the alarm relay information, issue the **show chassis craft-interface** command from the Junos OS command-line interface.

[See the *ACX Series Universal Access Router Configuration Guide* and the *System Basics: Chassis-Level Features Configuration Guide*. For a detailed description of the alarm contact port, see the relevant hardware guide for your router.]

Layer 2 and Layer 3 Protocols

- **IPv6 Support**—IPv6 builds upon the functionality of IPv4, providing improvements to addressing, configuration and maintenance, and security. The following IPv6 features are supported on ACX Series routers:
 - Dual stacking (IPv4 and IPv6)
 - Dynamic routes distribution through IS-IS and OSPF for IPv6
 - Internet Control Message Protocol (ICMP) v6
 - IPv6 forwarding
 - IPv6 over MPLS (6PE)
 - IPv6 path maximum transmission unit (MTU) discovery
 - Neighbor discovery
 - Static routes for IPv6

[See the *ACX Series Universal Access Router Configuration Guide* and the *Junos OS Routing Protocols Configuration Guide*.]

Routing Protocols

- **Support for Layer 3 VPNs for IPv4 and IPv6 address families**—You can configure Layer 3 virtual private network (VPN) routing instances on ACX Series routers at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level for unicast IPv4, multicast IPv4, unicast IPv6, and multicast IPv6 address families. If you do not explicitly specify the address family in an IPv4 or an IPv6 environment, the router is configured to exchange unicast IPv4 or unicast IPv6 addresses by default. You can also configure

the router to exchange unicast IPv4 and unicast IPv6 routes in a specified VPN routing and forwarding (VRF) routing instance. If you specify the multicast IPv4 or multicast IPv6 address family in the configuration, you can use BGP to exchange routing information about how packets reach a multicast source, instead of a unicast destination, for transmission to endpoints.

Only the forwarding and virtual router routing instances support unicast IPv6 and multicast IPv6 address families. Unicast IPv6 and multicast IPv6 address families are not supported for VRF routing instances.

A VRF routing instance is a BGP and MPLS VPN environment in which BGP is used to exchange IP VPN routes and discover the remote site, and VPN traffic traverses an MPLS tunnel in an IP and MPLS backbone. You can enable an ACX Series router to function as a provider edge (PE) router by configuring VRF routing instances.

You can configure the following types of Layer 3 routing instances:

- **Forwarding**—Use this routing instance type for filter-based forwarding applications.
- **Virtual router**—A virtual router routing instance is similar to a VRF instance type, but is used for non-VPN-related applications.
- **VRF**—Use the VRF routing instance type for Layer 3 VPN implementations. This routing instance type has a VPN routing table as well as a corresponding VPN forwarding table. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF routing instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. This routing instance type is used to implement BGP or MPLS VPNs in service provider networks or in big enterprise topologies.

[*ACX Series Universal Access Router Configuration Guide*]

- **Support for Multiprotocol BGP**—Multiprotocol BGP (MP-BGP) is an extension to BGP that enables BGP to carry routing information for multiple network layers and address families. MP-BGP can carry the unicast routes used for multicast routing separately from the routes used for unicast IP forwarding.

You can configure MP-BGP on ACX Series routers for IPv4 and IPv6 address families in the following ways:

- To enable MP-BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, include the **family inet** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MP-BGP to carry NLRI for the IPv6 address family, include the **family inet6** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MP-BGP to carry Layer 3 virtual private network (VPN) NLRI for the IPv4 address family, include the **family inet-vpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.

- To enable MP-BGP to carry Layer 3 VPN NLRI for the IPv6 address family, include the **family inet6-vpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MP-BGP to carry multicast VPN NLRI for the IPv4 address family and to enable VPN signaling, include the **family inet-mvpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MP-BGP to carry multicast VPN NLRI for the IPv6 address family and to enable VPN signaling, include the **family inet6-mvpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.

[ACX Series Universal Access Router Configuration Guide]

Time Division Multiplexing (TDM)

- **TDM CESoPSN (ACX1000 and ACX2000 routers)**—Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN) is a method of encapsulating TDM signals into CESoPSN packets, and in the reverse direction, decapsulating CESoPSN packets back into TDM signals—also, referred to as Interworking Function (IWF). The following CESoPSN features are supported:
 - **Channelization up to the ds0 level**—The following numbers of NxDS0 pseudowires are supported for 16 T1 and E1 built-in ports and 8 T1 and E1 built-in ports.
 16 T1 and E1 built-in ports support the following number of pseudowires:
 - Each T1 port can have up to 24 NxDS0 pseudowires, which add up to a total of up to 384 NxDS0 pseudowires.
 - Each E1 port can have up to 31 NxDS0 pseudowires, which add up to a total of up to 496 NxDS0 pseudowires.
 8 T1 and E1 built-in ports support the following number of pseudowires:
 - Each T1 port can have up to 24 NxDS0 pseudowires, which add up to a total of up to 192 NxDS0 pseudowires.
 - Each E1 port can have up to 31 NxDS0 pseudowires, which add up to a total of up to 248 NxDS0 pseudowires.
 - **Protocol support**—All protocols that support Structure Agnostic TDM over Packet (SAToP) support CESoPSN NxDS0 interfaces.
 - **Packet latency**—The time required to create packets (from 1000 through 8000 microseconds).
 - **CESoPSN encapsulation**—The following statements are supported at the **[edit interfaces interface-name]** hierarchy level:
 - **ct1-x/y/z partition partition-number timeslots timeslots interface-type ds**
 - **ds-x/y/z:n encapsulation cesopsn**

- **CESoPSN options**—The following statements are supported at the `[edit interfaces interface-name cesopsn-options]` hierarchy level
 - `excessive-packet-loss-rate` (sample-period *milliseconds*)
 - `idle-pattern` *pattern*
 - `jitter-buffer-latency` *milliseconds*
 - `jitter-buffer-packets` *packets*
 - `packetization-latency` *microseconds*
- **Interfaces show commands**—The `show interfaces interface-name extensive` command is supported for `t1`, `e1`, and `at` interfaces.
- **CESoPSN pseudowires**—CESoPSN pseudowires are configured on the logical interface, not on the physical interface. So the `unit logical-unit-number` statement must be included in the configuration at the `[edit interfaces interface-name]` hierarchy level. When you include the `unit logical-unit-number` statement, Circuit Cross Connect (CCC) for the logical interface is created automatically.

[See the *ACX Series Universal Access Router Configuration Guide*.]

Timing and Synchronization

- **IEEE 1588v2 boundary clock**—The boundary clock has multiple network connections and can act as a source (master) or destination (backup) for synchronization messages. The boundary clock intercepts and processes all Precision Time Protocol (PTP) messages and passes all other traffic. The best master clock algorithm (BMCA) is used by the boundary clock to select the best clock from configured acceptable masters. On ACX Series routers, you can configure a port as a boundary backup or as a boundary master. To configure a boundary clock, include the `boundary` statement at the `[edit protocols ptp clock-mode]` hierarchy level.

[See the *ACX Series Universal Access Router Configuration Guide*.]

- **PTP master boundary clock**—On an ACX Series router, the Precision Time Protocol (PTP) master clock sends unicast packets over UDP to the clients (ordinary and boundary) so they can establish their relative time offset from this master clock. To configure a master clock, include the `master` statement and options at the `[edit protocols ptp]` hierarchy level. On an ACX Series router, you can configure up to 512 remote clock clients. The following configuration is supported for the master boundary clock:

```
[edit protocols ptp master]
announce-interval announce-interval-value;
interface interface-name {
  unicast-mode {
    clock-client ip-address local-ip-address local-ip-address {
      manual;
    }
  }
  transport ipv4;
}
max-announce-interval max-announce-interval;
```

```

max-delay-response-interval max-delay-response-interval;
max-sync-interval max-sync-interval;
min-announce-interval min-announce-interval;
min-delay-response-interval min-delay-response-interval;
min-sync-interval min-sync-interval;
sync-interval sync-interval;

```



NOTE: You must include the boundary statement at the [edit protocols ptp clock-mode] hierarchy level and at least one slave with the slave statement at the [edit protocols ptp] hierarchy level for the remote master configuration to work

[See the *ACX Series Universal Access Router Configuration Guide*.]

- **Clock clients**—A clock client is the remote PTP host, which receives time from the PTP master and is in a slave relationship to the master. The maximum number of configured clock clients is 512. The clock client is included in the configuration of the master clock. Three different types of downstream clients are supported. You can configure any combination of these three types of clients for a given master.
 - Automatic client—For an automatic client, you do not need to configure the exact IP address of the host. Instead, configure a subnet mask for the automatic client, and any host belonging to that subnet can join the master clock through a unicast negotiation—which is a method by which the announce, synchronization, and delay response packet rates are negotiated between the master and the slave before a Precision Time Protocol (PTP) session is established. To configure an automatic client, include the **clock-client ip-address local-ip-address local-ip-address** statement at the [edit protocols ptp master interface *interface-name* unicast-mode] hierarchy level. Include the subnet mask of the remote PTP host in the **clock-client ip-address** statement and the boundary master clock IP address in the **local-ip-address local-ip-address** statement.
 - Manual client—When you configure a manual client, the client immediately receives announce and synchronization packets. To configure a manual client, include the **manual** statement at the [edit protocols ptp master interface *interface-name* unicast-mode clock-client ip-address local-ip-address local-ip-address] hierarchy level.
 - Secure client—For a secure client, you must configure a full and exact IP address, after which it joins the master clock through unicast negotiation. To configure a secure client, include the **clock-client ip-address** statement with the exact IP address of the PTP host at the [edit protocols ptp master interface *interface-name* unicast-mode] hierarchy level.



NOTE: You can configure the maximum number of clients (512) in the following combination:

- Automatic clients 256.
- Manual and secure clients 256—Any combination of manual and secure clients is allowed as long as the combined total amounts to 256.

[See the *ACX Series Universal Access Router Configuration Guide*.]

Changes in Default Behavior and Syntax in Junos OS Release 12.3 for ACX Series Routers

- [Interfaces and Chassis on page 14](#)
- [IPv6 on page 14](#)
- [Junos OS XML API and Scripting on page 14](#)
- [Network Management and Monitoring on page 15](#)
- [Security on page 15](#)

Interfaces and Chassis

- **Connectivity Fault Management MEPs on Layer 2 Circuits and Layer 2 VPNs** — On interfaces configured on ACX Series routers, you no longer need to configure the **no-control-word** statement at either the **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** or the **[edit routing-instances *routing-instance-name* protocols l2vpn]** hierarchy level for Layer 2 circuits and Layer 2 VPNs over which you are running CFM maintenance endpoints (MEPs). This configuration is not needed because ACX Series routers support the control word for CFM MEPs. The control word is enabled by default.

[*ACX Series Universal Access Router Configuration Guide*]

IPv6

- **Change in automatically generated virtual-link-local-address for VRRP over IPv6**—The seventh byte in the automatically generated virtual-link-local-address for VRRP over IPv6 is 0x02. This change makes the VRRP over IPv6 feature in the Junos OS 12.2R5, 12.3R3, 13.1R3, and later releases, inoperable with Junos OS 12.2R1, 12.2 R2, 12.2 R3, 12.2R4, 12.3R1, 12.3R2, 13.1R1, and 13.3R2 releases if automatically generated virtual-link-local-address ID used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

Junos OS XML API and Scripting

- **IPv6 address text representation is stored internally and displayed in command output using lowercase**—Starting with Junos OS Release 11.1R1, IPv6 addresses are stored internally and displayed in the command output using lowercase. Scripts that match on an uppercase text representation of IPv6 addresses should be adjusted to either match on lowercase or perform case-insensitive matches.

- **<get-configuration> RPC with inherit="inherit" attribute returns correct time attributes for committed configuration**—In Junos OS Release 12.3R1, when you configured some interfaces using the interface-range configuration statement, if you later requested the committed configuration using the <get-configuration> RPC with the inherit="inherit" and database="committed" attributes, the device returned `junos:changed-localtime` and `junos:changed-seconds` in the RPC reply instead of `junos:commit-localtime` and `junos:commit-seconds`. This issue is fixed in Junos OS Release 12.3R2 and later releases so that the device returns the expected attributes in the RPC reply.

Network Management and Monitoring

- **New system log message indicating the difference in the Packet Forwarding Engine counter value (ACX Series)**—Effective in Junos OS Release 12.3R9, if the counter value of a Packet Forwarding Engine is reported to be less than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the `MIB2D_COUNTER_DECREASING` system log message for that specific counter.

Security

- In all supported Junos OS releases, regular expressions can no longer be configured if they require more than 64MB of memory or more than 256 recursions for parsing.

This change in the behavior of Junos OS is in line with the Free BSD 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 process (rpd).

Errata and Changes in Documentation for Junos OS Release 12.3 for ACX Series Routers

Errata

- [Class of Service \(CoS\) on page 16](#)
- [Firewall Filters on page 16](#)
- [Hardware on page 17](#)
- [Routing Protocols on page 17](#)

Class of Service (CoS)

- The “Understanding CoS CLI Configuration Statements on ACX Series Universal Access Routers” topic in the *ACX Series Universal Access Router Configuration Guide* incorrectly states that the **[edit class-of-service routing-instances routing-instance-name]** hierarchy level is not applicable for ACX Series routers when in fact it is available.

Firewall Filters

- Support for multifield classifiers is incorrectly omitted from the ACX Series documentation. Multifield classifiers allow fine grained classification by examination of multiple fields in the packet header—for example, the source and destination address of the packet, and the source and destination port numbers of the packet. A multifield classifier typically matches one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. Multifield classifiers are used when a simple BA classifier is insufficient to classify a packet.

In Junos OS, you configure a multifield classifier with a firewall filter and its associated match conditions. This enables you to use any filter match criteria to locate packets that require classification. From a CoS perspective, multifield classifiers (or firewall filter rules) provide the following services:

- Classify packets to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by schedulers in conjunction with the random early discard (RED) algorithm to control packet discard during periods of congestion.
- Police traffic to a specific bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, to a different loss priority, or to both.



NOTE: You police traffic on input to conform to established CoS parameters, setting loss handling and forwarding class assignments as needed. You shape traffic on output to make sure that router resources, especially bandwidth, are distributed fairly. However, input policing and output shaping are two different CoS processes, each with their own configuration statements.

To configure multifield classifiers, include the following statements at the **[edit firewall]** hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        dscp 0;
        forwarding-class class-name;
      }
    }
  }
}
```



```

        loss-priority (high | low);
    }
}
}
simple-filter filter-name {
    term term-name {
        from {
            match-conditions;
        }
        then {
            forwarding-class class-name;
            loss-priority (high | low | medium);
        }
    }
}
}
}

```

The minimum configuration required to define a multifield classifier is the following:

```

[edit firewall]
family family-name {
    simple-filter filter-name {
        term term-name {
            then {
                forwarding-class class-name;
                loss-priority (high | low | medium);
            }
        }
    }
}
}

```

After defining the multifield classifier, you can apply the multifield classifier to an individual interface with the following configuration:

```

[edit interfaces]
interface-name {
    unit logical-unit-number {
        family family {
            filter {
                input filter-name;
            }
        }
    }
}
}

```

[ACX Series Universal Access Router Configuration Guide]

Hardware

- In the “ACX2000 and ACX2100 DC Power Specifications” topic of the *ACX2000 and ACX2100 Router Hardware Guide*, the DC input voltages row in the table presented in the topic incorrectly mentions that the range is 18 to 30 VDC. The correct DC input voltage range for a nominal 24 volt operation is 20 to 30 VDC.

Routing Protocols

- The following additional information about the support for unicast IPv6 and multicast IPv6 address families for routing instances on ACX Series routers applies to the *Routing*

Protocols subsection in the *New Features in Junos OS Release 12.3 for ACX Series Routers* section of the Junos OS 12.3R1 Release Notes and the *Layer 3 VPNs for IPv4 and IPv6 Overview* topic of the *ACX Series Universal Access Router Configuration Guide*:

Only the forwarding and virtual router routing instances support unicast IPv6 and multicast IPv6 address families. Unicast IPv6 and multicast IPv6 address families are not supported for VRF routing instances.

[*Release Notes, ACX Series Universal Access Router Configuration Guide*]

- The *OSPF Configuration Guide* incorrectly includes the **transmit-interval** statement at the **[edit protocols ospf area area interface interface-name]** hierarchy level. The **transmit-interval** statement at this hierarchy level is deprecated in the Junos OS command-line interface.

[*OSPF Configuration Guide*]

Known Limitations in Junos OS Release 12.3 for ACX Series Routers

The following software limitations currently exist in Juniper Networks ACX Series Universal Access Routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

Class of Service

- When the **rewrite-rules** statement is configured with the **dscp** or the **inet-precedence** options at the **[edit class-of-service interfaces]** hierarchy level, the expectation is that the DiffServ code point (DSCP) or IPv4 precedence rewrite rules take effect only on IP packets. However, in addition to the IP packets, the DSCP or IPv4 rewrite takes effect on the IP header inside the Ethernet pseudowire payload as well. [PR664062: This is a known limitation.]

Firewall Filters

- On ACX Series routers, packet drops in the egress interface queue are also counted as *input packet rejects* under the **Filter statistics** section in the output of the **show interface extensive** command when it is run on the ingress interface. [PR612441: This is a known software limitation.]
- When the **statistics** statement is configured on a logical interface, for example **[edit interface name-X unit unit-Y]**, when the (**policer** | **count** | **three-color-policer**) statements are configured in a firewall filter for the **family any**, for example the **[edit firewall family any filter filter-XYZ term term-T then]** hierarchy level, and the configured **filter-XYZ** is specified in the **output** statement of the above logical interface at the **[edit interface name-X unit unit-Y filter]** hierarchy level, the counters from the configuration of another firewall family filter on the logical interface do not work. [PR678847: This is a known limitation.]
- The policing rate can be incorrect if the following configurations are applied together:
 - The **policer** or **three-color-policer** statement configured in a firewall filter, for example **filter-XYZ** at the **[edit firewall family any filter filter-XYZ term term-T then]** hierarchy level, and **filter-XYZ** is specified as an ingress or egress firewall filter on a logical

interface, for example **interface-X unit-Y** at the **[edit interface interface-X unit unit-Y filter (input|output) filter-XYZ]** hierarchy level.

- The **policer** or **three-color-policer** statement configured in a firewall filter, for example **filter-ABC** at the **[edit firewall family name-XX filter filter-ABC term term-T then]** hierarchy level, and **filter-ABC** is configured as an ingress or egress firewall filter on a family of the same logical interface **interface-X unit-Y** at the **[edit interface interface-X unit unit-Y family name-XX filter (input|output) filter-ABC]** hierarchy level.



NOTE: If one of these configurations is applied independently, then the correct policer rate can be observed.

[[PR678950](#): This is a known limitation.]

Interfaces and Chassis

- The ACX Series routers support logical interface statistics, but do not support the address family statistics. [[PR725809](#): This is a known limitation.]
- When the **differential-delay number** option is configured in the **ima-group-option** statement at the **[edit interfaces at-fpc/pic/ima-group-no]** hierarchy level, with a value less than 10, some of the member links might not come up and the group might remain down resulting in traffic loss. A workaround is to keep the differential delay value above 10 for all IMA bundles. [[PR726279](#): This is a known limitation.]
- BERT error insertion and bit counters are not supported by the IDT82P2288 framer. [[PR726894](#): This is a known limitation.]
- The discard error is displayed when a Layer 2 pseudowire is configured with VLAN ID 0 on an NNI interface. The ACX4000 does not support VLAN ID 0 on the NNI interface. [[PR727276](#): This is a known limitation.]
- All 4x supported TPIDs cannot be configured on different logical interfaces of a physical interface. Only one TPID can be configured on all logical interfaces, that is, sub-interfaces of a physical interface. But different physical interfaces can have different TPIDs. As a workaround, use TPID-rewrite. [[PR738890](#): This is a known limitation.]
- The ACX Series routers do not support logical interface statistics for those logical interfaces with **vlan-list/vlan-range**. [[PR810973](#): This is a known limitation.]
- CFM up-mep session (to monitor PW service) does not come up when output **vlan-map** is configured as **push** on AC ifl. This is due to a hardware limitation in Enduro-2. [[PR832503](#): This is a known limitation.]

MPLS

- The scaling numbers for pseudowires and MPLS label routes published for the ACX Series routers are valid only when the protocols adopt graceful restart. In case of nongraceful restart, the scaling numbers would become half of the published numbers. [[PR683581](#)]: This is a known limitation.]

Outstanding Issues in Junos OS Release 12.3 for ACX Series Routers

The following issues currently exist in Juniper Networks ACX Series Universal Access Routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

General Routing

- Customer may see Power Supply Failure/Removed SNMP messages. As long as you do not see any failures and alarms on chassis, these are cosmetic messages. [PR929629](#)
- Upon unplanned power loss, ACX1100 may display this message "/var/log: write failed, filesystem is full." upon login. [PR1027641](#)
- TWAMP packets reflected out on ACX Series routers go out on forwarding class specified under : "set class-of-service host-outbound-traffic forwarding-class <>". Also, the 802.1p/CFI bits of the incoming packet are not preserved after reflection. [PR986997](#)
- Multiple copies of packets are seen during flooding in the VPLS domain. This occurs in a point-to-multipoint topology when traffic is received on the attachment circuit. [PR1069986](#)
- On ACX5000 Series routers, when the transparent clock functionality is enabled, you may notice errors in the correction field of the PTP packet. Ports on ACX5000 series router can be phy or phy-less. The phy ports shows more error in the correction field of the PTP packet. [PR1076050](#)
- On ACX5000 Series routers with the transparent clock functionality enabled, the correction field error on 40G ports is higher, and it is not recommended for the 40G ports to be used for transparent clock operation. [PR1076518](#)

Resolved Issues in Junos OS Release 12.3 for ACX Series Routers

The following issues have been resolved in Junos OS Release 12.3 for Juniper Networks ACX Series routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

Resolved Issues

- Note: There are no resolved issues in Junos OS Release 12.3R11.

Previous Releases

Release 12.3R10

General Routing

- In case of any problem during boot-up, box will automatically reboot to Alternate slice and give customer chance to recover the primary slice by snapshot command. If autosnapshot is enabled, snapshot will be taken automatically on successful boot to Alternate slice. [PR839286](#)
- CFM MEP over L2VPN/L2 circuit on ACX Series no longer requires no-control-word to be configured under L2VPN or l2-circuit hierarchy. [PR864317](#)
- When 16xCHE1/T1 MIC is online Routing Engine driven BFD sessions may flap. [PR892011](#)
- AI-Scripts bundle install fails on ACX Series platform running Junos OS release 12.3X52. [PR925102](#)
- Image download via FTP on to ACX Series through inband takes a long time when compared to FXP. Changing the receive rate towards the CPU has helped in improving the transfer rate. [PR979858](#)

Release 12.3R9**Layer 2 Features**

- If "maintain-subscriber" knob is enabled on the router, DHCPv6 server/relay might be unable to process any packet if deactivate and then activate the routing instance, which means the subscribers can not get the IPv6 addresses. Note, even with the fix, the results of this scenario are also expected if with "maintain-subscriber" knob enabled, Consider using the workaround to avoid this issue. [PR1018131](#): This issue has been resolved.

MPLS

- When the size of a Routing Engine generated packet going over an MPLS LSP is larger than MTU (i.e. MTU minus its header size) of an underlying interface, and the extra bytes leading to IP-fragmentation are as small as <8 bytes, then that small-fragment will be dropped by kernel and lead to packet drop with kernel message "tag_attach_labels(): m_pullup() failed". For example - If SNMP Response with specific size fall into above mentioned condition, then small fragment will be dropped by kernel and eventually the SNMP response will fail. [PR1011548](#): This issue has been resolved.

Release 12.3R8**General Routing**

- jdhcpd crash triggered by DHCP client re-discovering on different vlan. [PR913823](#): This issue has been resolved.

Release 12.3R7

- Note: There are no resolved issues in Junos OS Release 12.3R7.

Release 12.3R6**Interfaces and Chassis**

- On ACX Series-based FPC, one-way packet loss might occur when bridge-domain is stitched to l2circuit via logical tunnel (LT) interface. [PR922375](#): This issue has been resolved.

Release 12.3R5**Junos Scope**

- Heater status shows to be on under normal operating temperatures. This is a cosmetic issue. [PR898079](#): This issue has been resolved.
- This is just an SNMP Trap generated due to inconsistency in temperature parameters and thresholds. [PR859507](#): This issue has been resolved.

Release 12.3R4

Interfaces and Chassis

- When you issue the **show system memory** command on ACX routers, the "unable to load pmap_helper module: No such file or directory" error message is displayed in the output of the command. [PR737616](#): This issue has been resolved.
- CFM MEP over L2VPN/L2 circuit on ACX no longer requires no-control-word to be configured under L2VPN or l2-circuit hierarchy. [PR864317](#) and [PR801746](#): This issue has been resolved.
- When the data-tlv size is set to ≤ 54 then DM frames are considered as invalid. [PR800605](#): This issue has been resolved.

MPLS

- When an ACX Series router is working as a transit router in an MPLS network, then the traceroute response from router provides "nexthop" information as "Unhelpful". [PR669005](#): This issue has been resolved.

Release 12.3R2

Interfaces and Chassis

- A **commit** for configuration change that simultaneously disables RSVP and a point-to-point interface, such as so, t1, and atm, might generate a core file on the routing protocol process. To solve this issue, do not **commit** a configuration change that simultaneously disables RSVP and a point-to-point interface. Instead, disable RSVP and point-to-point interfaces in separate config commits. [PR782174](#): This issue has been resolved.
- On an L2circuit, when any one of the logical interfaces on the PE routers is disabled after changing VLAN-tagging to flexible VLAN tagging or vice versa on a CE router, the pseudowires return an mtu mismatch error. As a workaround, disable and then enable all the logical interfaces on the PE routers. [PR834466](#): This issue has been resolved.

Upgrade and Downgrade Instructions for Junos OS Release 12.3 for ACX Series Routers

This section discusses the following topics:

- [Basic Procedure for Upgrading to Release 12.3 on page 23](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 26](#)
- [Downgrading from Release 12.3 on page 26](#)

Basic Procedure for Upgrading to Release 12.3

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



.....

NOTE: 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 *Understanding System Snapshot on an ACX Series Router*.

.....

The download and installation process for Junos OS Release 12.3 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 web page:
<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-12.3R11  
1-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot source/jinstall-12.3R11  
1-export-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 12.3 **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>.

Downgrading from Release 12.3

To downgrade from Release 12.3 to another supported release, follow the procedure for upgrading, but replace the 12.3 **jinstall** package with one that corresponds to the appropriate release.



.....

NOTE: You cannot downgrade more than three releases. For example, if your routing platform is running Junos OS Release 11.4, you can downgrade the software to Release 10.4 directly, but not to Release 10.3 or earlier; as a workaround, you can first downgrade to Release 10.4 and then downgrade to Release 10.3.

.....

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

Junos OS Release Notes for EX Series Switches

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

New Features in Junos OS Release 12.3 for EX Series Switches

This section describes new features in Release 12.3 of the Junos operating system (Junos OS) for EX Series switches.

Not all EX Series software features are supported on all EX Series switches in the current release. For a list of all EX Series software features and their platform support, see [Feature Explorer](#).

New features are described on the following pages:

- [Hardware on page 29](#)
- [Access Control and Port Security on page 30](#)
- [Class of Service \(CoS\) on page 30](#)
- [Converged Networks \(LAN and SAN\) on page 31](#)
- [Enhanced Layer 2 Software \(ELS\) on EX9200 Switches on page 32](#)
- [Ethernet Switching and Spanning Trees on page 32](#)
- [Firewall Filters and Routing Policy on page 33](#)
- [High Availability on page 34](#)
- [Infrastructure on page 34](#)
- [Interfaces and Chassis on page 35](#)
- [IPv6 on page 36](#)
- [J-Web Interface on page 37](#)
- [Layer 2 and Layer 3 Protocols on page 37](#)
- [Management and RMON on page 37](#)
- [MPLS on page 38](#)
- [Virtual Chassis on page 38](#)

Hardware

Juniper Networks EX9200 Ethernet Switches—The EX9200 Programmable Ethernet Switches support current and planned SDN interfaces and protocols, offering the flexibility and scalability to increase business agility by simplifying the deployment and operation of cloud applications, by offering server virtualization, and by supporting rich media collaboration tools across campuses and data centers. The EX9200 switches provide high performance, scalable connectivity, and carrier-class reliability for high-density environments such as campus aggregation and data center networks.

The first supported Junos OS release for the EX9200 switches is Release 12.3R2.

The EX9200 switches are modular systems that provide high availability and redundancy for all major hardware components, including Routing Engine modules, Switch Fabric (SF) modules, fan trays (with redundant fans), and power supplies. Four line cards are available for the EX9200 switches.

The three EX9200 switches are:

- **EX9204 Ethernet switch**—The EX9204 switch has a capacity of up to 1.6 terabits per second (Tbps), full duplex.

The EX9204 switch has a 6-U chassis. It has two dedicated slots for line cards and a multifunction slot that can be used for either a line card or a host subsystem, all horizontal and all on the front of the switch chassis.

- **EX9208 Ethernet switch**—The EX9208 switch has a capacity of up to 4.8 Tbps, full duplex.

The EX9208 switch has an 8-U chassis and six horizontal line card slots on the front of the switch chassis.

- **EX9214 Ethernet switch**—The EX9214 switch has a capacity of up to 13.2 Tbps, full duplex.

The EX9214 switch has a 16-U chassis and has 12 vertical line card slots on the front of the switch chassis.

The line cards combine a Packet Forwarding Engine and Ethernet interfaces in a single assembly. Line cards are field-replaceable units (FRUs), and they are hot-insertable and hot-removable.

The four line cards available for EX9200 switches are:

- **EX9200-32XS**—32-port SFP+ line card
- **EX9200-40T**—40-port 10/100/1000BASE-T RJ-45 line card
- **EX9200-40F**—40-port 100FX/1000BASE-X SFP line card
- **EX9200-4QS**—4-port 40-Gigabit Ethernet QSFP+ line card

[See [EX9204 Hardware Documentation](#), [EX9208 Hardware Documentation](#), and [EX9214 Hardware Documentation](#).]

Access Control and Port Security

- **MAC limiting enhancements**—The MAC limiting feature for access port security has been enhanced to provide additional flexibility and granularity. The new feature, VLAN membership MAC limit, enables you to configure a MAC limit for a specific interface based on its membership in a particular VLAN (VLAN membership MAC limit). A single interface that belongs to multiple VLANs can thus have more than one MAC limit. [See [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches](#).]
- **VR-aware DHCP server and relay with option 82 on EX8200 switches and EX8200 Virtual Chassis**—VR-aware DHCP (extended DHCP) server with option 82 is now supported on EX8200 standalone switches and EX8200 Virtual Chassis. [See [Understanding DHCP Services for Switches](#) and [Understanding the Extended DHCP Relay Agent for EX Series Switches](#).]
- **VR-aware DHCPv6 server and relay support**—Virtual router-aware (VR-aware) DHCPv6 server and VR-aware DHCPv6 relay are now supported on these switch platforms:
 - EX4500, EX4550, and EX6210 standalone switches
 - EX4200, EX4500, EX4550, mixed EX4200, EX4500, and EX4550, and EX8200 Virtual Chassis[See [dhcpv6 \(DHCP Relay Agent\)](#) and [dhcpv6 \(DHCP Local Server\)](#).]
- **Bypassing 802.1X authentication when adding multiple LLDP-MED end devices**—If you have a large-scale installation of LLDP-MED end devices, you can save configuration time by specifying the `lldp-med-bypass` statement at the `[edit protocols dot1x authenticator interface (all | interface-name)]` hierarchy level. By configuring the `lldp-med-bypass` statement on an interface, you enable the interface to bypass the 802.1X authentication procedure for connecting multiple LLDP-MED end devices. This configuration automatically adds the learned MAC addresses of the LLDP-MED end devices to the switch's static MAC bypass list, so that you do not have to individually add the MAC address of each device. You can issue the `lldp-med-bypass` statement only when the interface is also configured for 802.1X authentication of *multiple* supplicants. [See [lldp-med-bypass](#).]
- **Access control and port security features support added on EX3300 switches**—EX3300 switches now support:
 - Captive portal authentication on Layer 2 interfaces
 - Persistent MAC learning (sticky MAC)[See [Understanding Authentication on EX Series Switches](#) and [Understanding Persistent MAC Learning \(Sticky MAC\)](#).]

Class of Service (CoS)

- **Class-of-service feature support added on EX3300 switches**—EX3300 switches now support:

- IPv6 CoS (multifield classification and rewrite)
- Flexible CoS outer 802.1p marking

[See [Junos OS CoS for EX Series Switches Overview](#).]

Converged Networks (LAN and SAN)

- **Enhanced transmission selection (IEEE 802.1Qaz) support on EX4500 switches**—The EX4500 switch models that support Converged Enhanced Ethernet (CEE) now provide limited support for enhanced transmission selection (ETS) (IEEE 802.1Qaz). ETS is a bandwidth management mechanism that supports dynamic allocation of bandwidth for Data Center Bridging Capability Exchange protocol (DCBX) traffic classes.

EX Series switches do not support the use of ETS to dynamically allocate bandwidth to traffic classes. Instead, the switches handle all DCBX traffic as a single default traffic class, group 7.

However, the switches do support the ETS Recommendation TLV. The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected DCBX peer interface to use.

If the peer interface is willing to learn the ETS state of the switch, it changes its configuration to match the configuration in the ETS Recommendation TLV sent by the EX Series switch (that is, the traffic class group 7).

The switch advertises that it is not willing to change its ETS settings.

The advertisement of the ETS TLV is enabled by default for DCBX interfaces, but you can disable it.

[See [Disabling the ETS Recommendation TLV](#).]

- **Support for IEEE DCBX**—The EX4500 switch models that support Converged Enhanced Ethernet (CEE) now also support the IEEE Data Center Bridging Capability Exchange protocol (IEEE DCBX). Earlier, these switches supported only DCBX version 1.01.

DCBX version 1.01 and IEEE DCBX differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.

DCBX is enabled by default on all 10-Gigabit Ethernet interfaces, and the default setting for the DCBX version on those interfaces is **auto-negotiation**.

When the interface DCBX version is set for **auto-negotiation** (the default):

- The switch sends IEEE DCBX TLVs. If the DCBX peer advertises the IEEE DCBX TLV three times, the switch changes the local DCBX interface to IEEE DCBX.
- If the DCBX peer advertises DCBX version 1.01 TLVs three times, the switch changes the local DCBX interface to **dcbx-version-1.01**.

When the interface DCBX version is set for **dcbx-version-1.01**:

- The switch sends DCBX version 1.01 TLVs and ignores any IEEE DCBX TLVs from the peer.

When the interface DCBX version is set for **ieee-dcbx**:

- The switch sends IEEE DCBX–based TLVs and ignores any DCBX version 1.01 TLVs from the peer.

To configure the DCBX version, use the **set dcbx-version** command at the **[edit protocols dcbx interface (all | interface-name)]** hierarchy level.

The **show dcbx neighbors** command has been updated with additional fields that support the IEEE DCBX feature; these fields include Interface Protocol-Mode and TLV Type.

[See “Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches” on page 95.]

- **VN_Port to VN_Port FIP snooping on EX4500 switches**—You can configure VN_Port to VN_Port (VN2VN_Port) FIP snooping if the hosts are directly connected to the same EX4500 switch. VN2VN_Port FIP snooping on an FCoE transit switch provides security to help prevent unauthorized access and data transmission on a bridge that connects ENodes in the Ethernet network. VN2VN_Port FIP snooping provides security for virtual links by creating filters based on information gathered (snooped) about FCoE devices during FIP transactions. [See [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\)](#); see also “Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches” on page 95.]

Enhanced Layer 2 Software (ELS) on EX9200 Switches

- **Uniform Enhanced Layer 2 Software (ELS) CLI configuration statements and operational commands**—Enhanced Layer 2 Software (ELS) provides a uniform CLI for configuring and monitoring Layer 2 features on EX9200 switches and on MX Series routers in LAN mode (MX-ELM). With ELS, for example, you can configure a VLAN and other Layer 2 features on an EX9200 switch and an MX-ELM router by using the same configuration commands. [See the ELS CLI documentation for EX9200 switches: [Junos OS for EX9200 Switches, Release 12.3](#).]
- The web-based ELS Translator tool is available for registered customers to help them become familiar with the ELS CLI and to quickly translate existing EX Series switch–based CLI configurations into ELS CLI configurations. [See [ELS Translator](#).]

Ethernet Switching and Spanning Trees

- **Ethernet ring protection switching**—Ethernet ring protection switching has been extended to the following switches:
 - EX3300 switches
 - EX4500 switches
 - EX4550 switches
 - EX4550 Virtual Chassis

- Mixed EX4200 and EX4500 Virtual Chassis
- EX8200 switches
- EX8200 Virtual Chassis

Support for all these switches is in addition to that for EX2200, EX3200, and EX4200 switches provided in earlier releases. Ethernet ring protection switching, defined in the ITU-T G.8032 recommendation, provides a means to reliably achieve carrier-class network requirements for Ethernet topologies forming a closed loop. [See [Ethernet Ring Protection Switching Overview](#).]

- **Disable MAC notifications on an interface**—On EX Series switches, when you enable media access control (MAC) notifications, learned and unlearned MAC address and aging SNMP notifications are unicast on all switch interfaces. In a large Layer 2 domain, unicasting might be undesirable because it can lead to significantly excess traffic. You can now disable such notifications on individual interfaces. For example, you might need notifications only for devices that are locally attached to the switch; you might not need notifications that arrive through uplinks. To disable notifications on an interface, issue the **set ethernet-switching-options interfaces *interface-name* no-mac-notification** command. [See [Understanding MAC Notification on EX Series Switches](#).]
- **VLAN pruning within an EX Series Virtual Chassis**—VLAN pruning is now supported within an EX Series Virtual Chassis. When VLAN pruning is enabled within an EX Series Virtual Chassis, all broadcast, multicast, and unknown unicast traffic in a VLAN uses the shortest path possible across the Virtual Chassis to the egress VLAN interface. VLAN pruning within an EX Series Virtual Chassis enables you to conserve Virtual Chassis bandwidth by restricting broadcast, multicast, and unknown unicast traffic in a VLAN to the shortest possible path across the Virtual Chassis instead of broadcasting this traffic to all Virtual Chassis member switches. [See [Enabling VLAN Pruning for Broadcast, Multicast, and Unknown Unicast Traffic in an EX Series Virtual Chassis \(CLI Procedure\)](#).]
- **Spanning-tree protocol concurrent configuration support added on EX3300 switches**—EX3300 switches now support concurrent configuration of RSTP and VSTP. [See [Understanding RSTP for EX Series Switches](#).]
- **Extended Q-in-Q VLAN support for multiple S-VLANs per access interface on EX3300 switches**—EX3300 switches now support filter-based S-VLAN tagging. [See [Understanding Q-in-Q Tunneling on EX Series Switches](#).]

[Firewall Filters and Routing Policy](#)

- **Support for firewall filters with IPv6 EX3300 switches**—EX3300 switches now support IPv6 firewall filters. [See [Firewall Filters for EX Series Switches Overview](#).]
- **Layer 3 unicast routing policy on EX3300 switches**—EX3300 switches now support Layer 3 unicast routing policy.
- **Support for match conditions, actions, and action modifiers for IPv6 firewall filters on EX2200 and EX3300 switches**—Starting with Junos OS Release 12.3R6, you can configure new match conditions, actions, and action modifiers for IPv6 firewall filters

on EX2200 and EX3300 switches. [See [Platform Support for Firewall Filter Match Conditions, Actions, and Action Modifiers on EX Series Switches](#).]

High Availability

- **Nonstop bridging for the Ethernet switching process (eswd), LLDP, LLDP-MED, and spanning-tree protocols on EX3300 Virtual Chassis**—Nonstop bridging (NSB) for the Ethernet switching process (eswd), LLDP, LLDP-MED, and spanning-tree protocols is now supported on EX3300 Virtual Chassis. You can now configure NSB to enable a transparent switchover between the master and backup Routing Engines without having to restart any of these processes or protocols. [See [Understanding Nonstop Bridging on EX Series Switches](#).]
- **Nonstop active routing, graceful protocol restart, and graceful Routing Engine switchover enhancements for standalone EX8200 switches and EX8200 Virtual Chassis**—Nonstop active routing (NSR), which enables a transparent switchover of Routing Engines without requiring restart of supported routing protocols, now supports RSVP and LDP on EX8200 standalone switches and EX8200 Virtual Chassis. Graceful protocol restart, a feature that enables a switch undergoing a restart to inform its adjacent neighbors and peers of the restart, is now supported for RSVP and LDP on standalone EX8200 switches and EX8200 Virtual Chassis. Graceful Routing Engine switchover (GRES) for Layer 2 and Layer 3 VPN LSPs is now supported on standalone EX8200 switches and EX8200 Virtual Chassis. [See [Understanding Nonstop Active Routing on EX Series Switches](#) or [High Availability Features for EX Series Switches Overview](#).]
- **Virtual Router Redundancy Protocol (VRRP) for IPv6 on EX3300 switches**—VRRP for IPv6 is now supported on EX3300 switches. [See [Understanding VRRP on EX Series Switches](#).]
- **Unified ISSU for EX9200 switches**—A unified in-service software upgrade (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. Unified ISSU is supported starting with Junos OS Release 12.3R3 on EX9200 switches with EX9200-40T or EX9200-40F line cards. [See [Unified ISSU Concepts](#).]

Infrastructure

- **Automatic repair of corrupted partition when booting from alternate partition**—Resilient dual-root partitioning has been enhanced to include an automatic snapshot feature. If the automatic snapshot feature is enabled and the system reboots from the alternate root partition, the switch automatically takes a snapshot of the Junos OS root file system in the alternate root partition and copies it to the primary root partition. This automatic snapshot procedure takes place whenever the system reboots from the alternate root partition, regardless of whether the reboot is due to a command or due to corruption of the primary root partition. [See [Understanding Resilient Dual-Root Partitions on Switches](#).]
- **BFD performance improvements**—BFD performance improvements have been made on EX4200 Virtual Chassis, EX4500 Virtual Chassis, and EX8200 switches.

- **IPv4 and IPv6 over GRE tunneling support on EX8200 standalone switches and EX8200 Virtual Chassis**—Generic routing encapsulation (GRE) is an IP encapsulation protocol that is used to transport packets over a network. Information is sent from one network to the other through a GRE tunnel. EX8200 standalone switches and EX8200 Virtual Chassis now support both encapsulation and de-encapsulation. Also, the configuration procedures for standalone EX8200 switches and EX8200 Virtual Chassis are now the same as for EX3200 and EX4200 switches. [See [Understanding Generic Routing Encapsulation](#).]
- **IPv6 for virtual router-aware DHCP**—EX Series switches support IPv6 for virtual router-aware DHCP, that is, for the extended DHCP server and extended DHCP relay. The specific CLI statements supported for EX Series switches are:
 - For extended DHCP server:
 - At the `[edit system services dhcp-local server dhcpv6]` hierarchy level:
 - `group`
 - `overrides`
 - `reconfigure`
 - At the `[edit access address-assignment pool pool-name]` hierarchy level:
 - `family inet6`
 - `dhcp-attributes`
 - `prefix`
 - `range`
 - For extended DHCP relay:
 - At the `[edit forwarding-options dhcp-relay dhcpv6]` hierarchy level:
 - `group`
 - `overrides`
 - `relay-agent-interface-id`
 - `relay-option`
 - `server-group`

[See [Understanding DHCP Services for Switches](#) and [Understanding the Extended DHCP Relay Agent for EX Series Switches](#).]

Interfaces and Chassis

- **LACP standards-based link protection for aggregated Ethernet interfaces**—LACP standards-based link protection can be enabled on a global level (for all aggregated Ethernet interfaces on the switch) or for a specific aggregated Ethernet interface. Earlier, EX Series switches supported only Junos OS link protection for aggregated Ethernet interfaces. [See [Understanding Aggregated Ethernet Interfaces and LACP](#).]

- **Interfaces feature support added on EX3300 switches**—EX3300 switches now support:

- Unicast reverse-path forwarding (RPF)
- IP directed broadcast

[See [Understanding Unicast RPF for EX Series Switches](#) and [Understanding IP Directed Broadcast for EX Series Switches](#).]

- **Default logging for Ethernet ring protection switching (ERPS) (EX2200, EX2200-VC, EX3200, EX3300, EX3300-VC, EX4200, EX4200-VC, EX4500, EX4500-VC, EX4550, EX4550-VC, EX8200, EX8200-VC)**—Starting with Junos OS Release 12.3R9, the listed EX switches automatically log basic state transitions for the ERPS protocol. No configuration is required to initiate this logging. Basic state transitions include ERPS interface transitions from up to down, and down to up; and ERPS state transitions from idle to protection, and protection to idle.

The basic state transitions are logged in a single file named **erp-default**, located in the **/var/log** directory of the switch. The maximum size of this file is 15 MB.

Default logging for ERPS can capture initial ERPS interface and state transitions, which can help you troubleshoot issues that occur early in the ERPS protocol startup process. However, if more robust logging is needed, you can enable traceoptions for ERPS by entering the **traceoptions** statement in the **[edit protocols protection-group]** hierarchy level.

Be aware that for ERPS, only default logging or traceoptions can be active at a time on the switch. That is, default logging for ERPS is automatically enabled and if you enable traceoptions for ERPS, the switch automatically disables default logging. Conversely, if you disable traceoptions for ERPS, the switch automatically enables default logging.

IPv6

- **Compliance with RFC 4291**—EX Series switches drop the following types of illegal IPv6 packets:
 - Packets that have a link-local source or destination address. Because link-local addresses are intended to be used for addressing only on a single link, EX Series switches do not forward any packets with such addresses to other links.
 - Packets with the IPv6 unspecified source address 0:0:0:0:0:0:0:0.
 - Packets that are to be sent outside a node, but have the IPv6 loopback address 0:0:0:0:0:0:0:1 as the source address. When IPv6 packets are received on an interface, EX Series switches drop packets that have the loopback address as the destination address.
- **IPv6 neighbor redirect compliance with RFC 4861**—Routers use ICMP redirect messages to notify the users on the data link that a better route is available for a particular destination. All EX Series switches now support sending ICMP redirect messages for both IPv4 and IPv6 traffic. [See [Understanding the Protocol Redirect Mechanism on EX Series Switches](#).]

- **Added license support for EX2200 and EX4200 switches**—The enhanced feature license (EFL) for EX2200 switches now supports the EX-2200-24T-DC model. The advanced feature license (AFL) for EX4200 switches now supports EX4200-24PX and EX4200-48PX models. [See [Understanding Software Licenses for EX Series Switches](#).]
- **Support for IPv6 features on EX3300 switches**—EX3300 switches now support:
 - IPv6 path MTU discovery
 - IPv6 routing BGP, RIPng, MBGP, and OSPFv3
 - IPv6 routing PIM for IPv6 multicast
 - IPv6 routing MLDv1 and MLDv2
 - IPv6 routing IPv6 ping and IPv6 traceroute
 - IPv6 routing stateless autoconfiguration
 - IPv6 routing IPv6 Layer 3 forwarding in hardware

J-Web Interface

- **10-member EX4500 Virtual Chassis configuration through the J-Web interface**—Using the J-Web interface, you can configure an EX4500 Virtual Chassis that includes a maximum of 10 members. [See [Configuring a Virtual Chassis on an EX Series Switch \(J-Web Procedure\)](#).]
- **EX8200 Virtual Chassis configuration through the J-Web interface**—Using the J-Web interface, you can configure an EX8200 Virtual Chassis to include up to four EX8200 switches and one or two XRE200 External Routing Engines. [See [Configuring a Virtual Chassis on an EX Series Switch \(J-Web Procedure\)](#).]

Layer 2 and Layer 3 Protocols

- **VRF support on EX2200 switches**—Virtual routing and forwarding (VRF) is now supported on EX2200 switches. [See [Understanding Virtual Routing Instances on EX Series Switches](#).]
- **Feature support added on EX3300 switches**—EX3300 switches now support:
 - Virtual routing and forwarding (VRF)—virtual routing instances—with IPv6 for unicast traffic
 - Layer 3 filter-based forwarding for unicast traffic
 - Layer 3 VRF for unicast BGP, RIP, and OSPF traffic
 - Multiple VLAN Registration Protocol (MVRP, IEEE 802.1ak)

Management and RMON

- **MIB enhancements on EX8200 Virtual Chassis**—The Virtual Chassis MIB has been enhanced to enable monitoring of Virtual Chassis interface statistics for EX8200 Virtual Chassis. [See [Juniper Networks Enterprise-Specific MIBs](#).]

- **Support for 802.1ag Ethernet OAM CFM on EX3300 switches**—EX3300 switches now support 802.1ag Ethernet OAM connectivity fault management (CFM). [See [Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch.](#)]

MPLS

- **Re-mark the DSCP values for MPLS packets that exit an EX8200 standalone switch or an EX8200 Virtual Chassis**—In firewall filter configurations for EX8200 standalone switches and EX8200 Virtual Chassis, you can now apply the **dscp** action modifier on Layer 3 interfaces for IPv4 and IPv6 ingress traffic. This action modifier is useful specifically to re-mark the DSCP values for MPLS packets that leave an EX8200 standalone switch or an EX8200 Virtual Chassis, because these switches cannot re-mark the DSCP value on egress traffic. If you apply the **dscp** action modifier to ingress traffic, the DSCP value in the IP header is copied to the EXP value in the MPLS header, thus changing the DSCP value on the egress side. [See [Firewall Filter Match Conditions, Actions, and Action Modifiers for EX Series Switches.](#)]

Virtual Chassis

- **Member link enhancement for optical interfaces configured as Virtual Chassis ports between EX4500 and EX4550 member switches**—When you configure optical interfaces as Virtual Chassis ports (VCPs) that you then use to interconnect EX4500 or EX4550 switches in a Virtual Chassis, you can now configure up to 24 optical interface links into a link aggregation group (LAG). Earlier, you could configure a maximum of eight links into a LAG. You can increase the member link limit in the following configurations: when you interconnect EX4500 switches in an EX4500 Virtual Chassis; when you interconnect EX4550 switches in an EX4550 Virtual Chassis; and when you interconnect EX4500 or EX4550 switches to other EX4500 or EX4550 switches in a mixed Virtual Chassis.
- **Dedicated Virtual Chassis port link aggregation on EX4550 switches**—The dedicated Virtual Chassis ports (VCPs) on EX4550 switches automatically form a link aggregation group (LAG) bundle when two or more dedicated VCPs are used to interconnect the same Virtual Chassis member switches. This feature became available in Junos OS Release 12.3R2. The LAG provides more bandwidth than a single dedicated VCP can provide, and it provides VCP redundancy by load-balancing traffic across all available dedicated VCPs in the LAG. If one of the dedicated VCPs fails, the VCP traffic is automatically load-balanced across the remaining dedicated VCPs in the LAG. See also “[Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches](#)” on page 95.
- **Support for RJ-45 interfaces as Virtual Chassis ports on EX2200 and EX2200-C switches**—Starting with Junos OS Release 12.3R3, all RJ-45 interfaces, including built-in network ports with 10/100/1000BASE-T Gigabit Ethernet connectors and 1000BASE-T RJ-45 transceivers, on EX2200 and EX2200-C switches can now be configured into Virtual Chassis ports (VCPs). VCPs are used to interconnect EX2200 or EX2200-C switches into a Virtual Chassis.

Related Documentation

- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)

- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches

This section lists the changes in default behavior and syntax in Junos OS Release 12.3 for EX Series switches.

Authentication and Access Control

- Starting with Junos OS Release 12.3R10 for EX Series switches, the **accounting-port** CLI statement is supported at the **[edit access radius-server server-address]** hierarchy level on all EX Series switches. This command was previously supported only on EX4300, EX4600, and EX9200 switches. The **accounting-port** statement enables you to specify the port on which to contact the RADIUS accounting server. The default port number is 1813, as specified in RFC 2866.

Dynamic Host Configuration Protocol

- When the DHCP relay agent receives a DHCP DISCOVER packet from a client while the client already has a binding on the relay that is in BOUND state, an OFFER message in reply might be dropped when sent from a server other than one to which the client is bound. You can now prevent this by using the new configuration statement **delete-binding-on-renegotiation** at the **[edit forwarding-options dhcp-relay overrides]** hierarchy level. When this option is configured, the client binding is removed on receiving a new DHCP DISCOVER packet from a client which is already in BOUND state and all OFFER messages in reply are relayed back to the client. [PR/1031605]

Layer 2 Features

- On EX Series switches except for EX9200 switches, the **vlan-tagging** and **family ethernet-switching** statements cannot be configured on the same interface. Interfaces on EX2200, EX3200, EX3300, EX4200, and EX4500 switches are set to **family ethernet-switching** by the default factory configuration. EX6200 and EX8200 switch interfaces do not have a default family setting.

Hardware

- On EX Series switches, EX-SFP-1GE-LX40K SFP transceivers (length 40 km / 24.8 miles) are incorrectly recognized as LH transceivers—that is, in the **show chassis hardware** command output, the Description field lists them as **SFP-LH**. Starting with Junos OS Release 12.3R8, the Description field displays **SFP-EX** for these transceivers. [PR/977327]

High Availability

- Change in the automatically generated virtual-link-local-address for VRRP over IPv6—The seventh byte in the automatically generated virtual-link-local-address for VRRP over IPv6 will be 0x02. This change makes the VRRP over IPv6 feature in Junos OS Releases 12.2R5, 12.3R3, and later releases inoperable with Junos OS Releases 12.2R1, 12.2R2, 12.2R3, 12.2R4, 12.3R1, and 12.3R2 if automatically generated virtual-link-local-address IDs are used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

Infrastructure

- You can now configure the disk usage monitoring level for a disk partition by using the

set chassis disk-partition *partition* level *state* free-space *threshold-value* (mb | percent) configuration mode command. When the specified disk usage monitoring level is reached, a system alarm is activated. The partition can be **/config** or **/var**; the level of disk usage at which monitoring occurs can be **high** or **full**; and the threshold value can be either megabytes (**mb**) of disk space or a percentage (**percent**) of disk space. Here is a sample command: **set chassis disk-partition /var level high free-space 30 mb**.

- These EX Series switches now support a maximum of 111 link aggregation groups (LAGs): EX3300, EX4200, EX4500, EX4550, and EX6210 switches.
- On EX Series switches, the **request chassis routing-engine master switch** command erroneously showed the **check** option; the option does not apply and has been removed from the CLI.

Interfaces

- LLDP frames are validated only if the Network Address Family subtype of the Chassis ID TLV has a value of 1 (IPv4) or 2 (IPv6). For any other value, LLDP detects the transmitting device as a neighbor and displays it in the output of the **show lldp neighbors** command. Earlier, the frames with the Network Address Family subtype of the Chassis ID TLV having a value of 1 (IPv4) or 2 (IPv6) were discarded, and LLDP did not detect the device as a neighbor.

Network Management and Monitoring

- On EX Series switches that run Junos OS with Enhanced Layer 2 Software (ELS), you can now configure the adaptive sampling rate in sFlow monitoring technology configurations. You configure the rate with this command: **set protocols sflow adaptive-sampling-rate *rate***. The maximum value for the rate is 950 pps.
- **New system log message indicating the difference in the Packet Forwarding Engine counter value (EX Series)**—Effective in Junos OS Release 12.3R9, if the counter value of a Packet Forwarding Engine is reported lesser than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the **MIB2D_COUNTER_DECREASING** system log message for that specific counter.

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

Limitations in Junos OS Release 12.3 for EX Series Switches

This section lists the limitations in Junos OS Release 12.3 for EX Series switches. If the limitation is associated with an item in our bug database, the description is followed by the bug tracking number.

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

Ethernet Switching and Spanning Trees

- If the bridge priority of a VSTP root bridge is changed such that this bridge becomes a nonroot bridge, the transition might take more than 2 minutes, and you might see a loop during the transition. [PR/661691: This is a known software limitation.]
- On EX9200 switches, MVRP does not propagate the dynamically learned VLAN information that is associated with trunk interfaces. [PR/840390, PR/848600: This is a known software limitation.]
- On EX9200 switches, BFD on IRB interfaces flaps if BFD is configured for subsecond timers. [PR/844951: This is a known software limitation.]

Firewall Filters

- On EX3200 and EX4200 switches, when a very large number of firewall filters are included in the configuration, it might take a long time, possibly a few minutes, for the egress filter rules to be installed. [PR/468806: This is a known software limitation.]
- On EX3300 switches, if you add and delete filters with a large number of terms (on the order of 1000 or more) in the same commit operation, not all the filters are installed. As a workaround, add filters in one commit operation, and delete filters in a separate commit operation. [PR/581982: This is a known software limitation.]
- On EX8200 switches, if you configure an implicit or explicit discard action as the last term in an IPv6 firewall filter on a loopback (lo0) interface, all the control traffic from the loopback interface is dropped. To prevent this, you must configure an explicit **accept** action. [This is a known software limitation.]
- On EX9200 switches, you cannot configure VLAN firewall filters for traffic leaving a VLAN. [PR/850520: This is a known software limitation.]

Hardware

- On 40-port SFP+ line cards for EX8200 switches, the LEDs on the left of the network ports do not blink to indicate that there is link activity if you set the speed of the network ports to 10/100/1000 Mbps. However, if you set the speed to 10 Gbps, the LEDs blink. [PR/502178: This is a known limitation.]
- The [Uplink Modules in EX3200 Switches](#) topic notes the following behavior for the SFP and SFP+ uplink modules:
 - On an EX3200 switch, if you install a transceiver in an SFP uplink module, a corresponding network port from the last four built-in ports is disabled. For example, if you install an SFP transceiver in port 2 on the uplink module (ge-0/1/2) on 24-port

models, then ge-0/0/22 is disabled. The disabled port is not listed in the output of **show interfaces** commands.

- On an EX3200 switch, if you install a transceiver in an SFP+ uplink module when the uplink module is operating in 1-gigabit mode, a corresponding network port from the last four built-in ports is disabled. For example, if you install an SFP transceiver in port 2 on the uplink module (ge-0/1/2), then ge-0/0/22 is disabled. The disabled port is not listed in the output of **show interfaces** commands.

However, if you install an SFP uplink module or an SFP+ uplink module when the SFP+ uplink module is operating in 1-gigabit mode and no transceiver is installed in the uplink module port, then all the network ports from the last four built-in ports are disabled and remain disabled until you reboot the switch.

If transceivers are installed in the uplink module ports, then only the corresponding built-in network ports are disabled and are not displayed in the output of **show interfaces** commands.

[PR/686467: This is a known limitation.]

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

High Availability

- You cannot verify that nonstop bridging (NSB) is synchronizing Layer 2 protocol information to the backup Routing Engine even when NSB is properly configured. [PR/701495: This is a known software limitation.]
- On EX Series Virtual Chassis using nonstop software upgrade (NSSU) to upgrade from Junos OS Release 11.2 or earlier to Junos OS Release 11.3 or later, after the NSSU operation finishes, the same MAC address might be assigned to multiple Layer 2 or aggregated Ethernet interfaces on different member switches within the Virtual Chassis. To set all Layer 2 and aggregated Ethernet ports to have unique MAC addresses, reboot the Virtual Chassis after the upgrade operation. To avoid these MAC address assignment issues, upgrade to Junos OS Release 11.3 or later without performing an NSSU operation. Unique MAC address assignment for Layer 2 and aggregated Ethernet interfaces in a Virtual Chassis was introduced in Junos OS Release 11.3. If you are upgrading to Junos OS Release 11.2 or earlier, you should expect to see the same MAC address assigned to multiple ports on different member switches within the Virtual Chassis. [PR/775203: This is a known software limitation.]

Infrastructure

- Do not use nonstop software upgrade (NSSU) to upgrade the software on an EX8200 switch from Junos OS Release 10.4 to Junos OS Release 11.1 or later if you have configured the PIM, IGMP, or MLD protocols on the switch. If you attempt to use NSSU, your switch might be left in a nonfunctional state from which it is difficult to recover. If you have these multicast protocols configured, use the **request system software add** command to upgrade the software on an EX8200 switch from Junos OS Release 10.4 to Release 11.1 or later. [This is a known software limitation.]
- On EX Series switches, the **show snmp mib walk etherMIB** command does not display any output, even though the etherMIB is supported. This occurs because the values are not populated at the module level—they are populated at the table level only. You can issue the **show snmp mib walk dot3StatsTable**, **show snmp mib walk dot3PauseTable**, and **show snmp mib walk dot3ControlTable** commands to display the output at the table level. [This is a known software limitation.]
- Momentary loss of an inter-Routing Engine IPC message might trigger an alarm that displays the message **Loss of communication with Backup Routing Engine**. However, no functionality is affected. [PR/477943: This is a known software limitation.]
- Routing between virtual-routing instances for local direct routes is not supported. [PR/490932: This is a known software limitation.]
- On EX4500 switches, the maintenance menu is not disabled even if you include the **lcd maintenance-menu disable** statement in the configuration. [PR/551546: This is a known software limitation.]
- When you enable the **filter-id** attribute on the RADIUS server for a particular client, none of the required 802.1X authentication rules are installed in the IPv6 database. Therefore, IPv6 traffic on the authenticated interface is not filtered; only IPv4 traffic is filtered on that interface. [PR/560381: This is a known software limitation.]
- On EX8200 switches, if OAM link fault management (LFM) is configured on a member of a VLAN on which Q-in-Q tunneling is also enabled, OAM PDUs are not transmitted to the Routing Engine. [PR/583053: This is a known software limitation.]
- When you reconfigure the maximum transmission unit (MTU) value of a next hop more than eight times without restarting the switch, the interface uses the maximum value of the eight previously configured values as the next MTU value. [PR/590106: This is a known software limitation.]
- On EX8208 and EX8216 switches that have two Routing Engines, one Routing Engine cannot be running Junos OS Release 10.4 or later while the other one is running Release 10.3 or earlier. Ensure that both Routing Engines in a single switch run either Release 10.4 or later or Release 10.3 or earlier. [PR/604378: This is a known software limitation.]
- On EX9200 switches, if you configure DHCP relay on an integrated routing and bridging (IRB) interface, DHCP relay does not perform binding on the client's DHCP Discover messages. As a workaround, configure the relay agent by using the BOOTP helper in the **[edit forwarding-options helpers]** hierarchy level. [PR/847772: This is a known software limitation.]

- On EX4550 switches, you might see the message **UI_OPEN_TIMEOUT: Timeout connecting to peer 'dhcp'**, and the message might appear even though you have not configured DHCP services. The operation of the switch is not affected, and you can ignore the message. [PR/895320: This is a known software limitation.]
- On EX9200 switches, an SRAM parity error might be logged during normal operation. This behavior is expected. You can ignore the error as long as you do not see large numbers of error messages. [PR/958661: This is a known software limitation.]

Interfaces

- EX Series switches do not support IPv6 interface statistics. Therefore, all values in the output of the **show snmp mib walk ipv6IfStatsTable** command always display a count of 0. [PR/480651: This is a known software limitation.]
- On EX8216 switches, a link might go down momentarily when an interface is added to a LAG. [PR/510176: This is a known software limitation.]
- On EX Series switches, if you clear LAG interface statistics while the LAG is down, then bring up the LAG and pass traffic without checking for statistics, and finally bring the LAG interface down and check interface statistics again, the statistics might be inaccurate. As a workaround, use the **show interfaces interface-name** command to check LAG interface statistics before bringing down the interface. [PR/542018: This is a known software limitation.]
- In some instances on an EX9200 switch, tagged traffic is not dropped on access interfaces even though the traffic is processed in the correct VLAN (the VLAN to which the access port belongs). If the packet exits the switch on a trunk port, the packet might be tagged twice. [PR/838597: This is a known software limitation.]
- On an EX9200 switch with a single Routing Engine, when the Routing Engine is rebooted, the interfaces do not immediately shut down. In this case, use the **set chassis power-off-ports-on-no-master-re** command with the **enable** or **disable** option. [PR/843743: This is a known software limitation.]
- On EX9200 switches, after you perform an online insertion of a QSFP+ transceiver in a 40-Gigabit Ethernet interface, the interface might take more than 10 to 15 seconds to come up. [PR/847186: This is a known software limitation.]
- On EX9200 switches, dynamic ARP resolution is not supported over interchassis control links (ICLs). As a workaround, configure static ARP on both ends of the ICL. [PR/850741: This is a known software limitation.]
- On EX Series switches, member links within the same link aggregation group (LAG) bundle must be configured to operate at the same speed. The default interface speed for RJ-45 BASE-T copper interfaces on an EX4550 switch is 10 gigabits per second (Gbps). The default interface speed for RJ-45 BASE-T copper interfaces on all other EX Series switches is 1 Gbps. You must, therefore, configure the RJ-45 BASE-T copper interfaces on an EX4550 switch to 1 Gbps using the **set interfaces xe-x/y/z ether-options speed 1g** command when you create a static LAG between RJ-45 BASE-T copper interfaces on an EX4550 switch and RJ-45 BASE-T copper interfaces on any other EX Series switch. [PR/940027: This is a known software limitation.]

- For aggregated Ethernet interfaces on EX Series switches, the traffic statistics fields in **show interfaces** commands do not include broadcast packet information. Also, for aggregated Ethernet interfaces, the SNMP counters ifHCInBroadcastPkts and ifInBroadcastPkts are not supported. The counter values are always 0. [This is a known software limitation.]
- On EX9200 switches, the CLI command **set interfaces interface-name speed auto-10m-100m** is not supported. [This is a known software limitation.]

J-Web Interface

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

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

- In the J-Web interface, the Ethernet Switching Monitor page (Monitor > Switching > Ethernet Switching) might not display monitoring details if the switch has more than 13,000 MAC entries. [PR/425693: This is a known software limitation.]
- On EX Series switches, when you use the Microsoft Internet Explorer browser to open reports from the following pages in the J-Web interface, the reports open in the same browser session:
 - Files page (Maintain > Files)
 - History page (Maintain > Config Management > History)
 - Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
 - Static Routing page (Monitor > Routing > Route Information)
 - Support Information page (Maintain > Customer Support > Support Information)
 - View Events page (Monitor > Events and Alarms > View Events)

[PR/433883: This is a known software limitation.]

- In the J-Web interface, if you open configuration pages for class-of-service (CoS) classifiers and drop profiles (**Configure > Class of Service > Classifiers and Configure > Class of Service > Drop Profile**), and then exit the pages without editing the configuration, no validation messages are displayed and the configuration of the switch proceeds. [PR/495603: This is a known software limitation.]
- In the J-Web interface for EX4500 switches, the Ports Configuration page (Configure > Interfaces > Ports), the Port Security Configuration page (Configure > Security > Port Security), and the Filters Configuration page (Configure > Security > Filters) display features that are not supported on EX4500 switches. [PR/525671: This is a known software limitation.]

- When you use an HTTPS connection in the Microsoft Internet Explorer browser to save a report from the following pages in the J-Web interface, the error message **Internet Explorer was not able to open the Internet site** is displayed on the following pages:
 - Files page (Maintain > Files)
 - History page (Maintain > Config Management > History)
 - Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
 - Static Routing page (Monitor > Routing > Route Information)
 - Support Information page (Maintain > Customer Support > Support Information)
 - View Events page (Monitor > Events and Alarms > View Events)

[PR/542887: This is a known software limitation.]

- If you insert four or more EX8200-40XS line cards in an EX8208 or EX8216 switch, the Support Information page (Maintain > Customer Support > Support Information) in the J-Web interface might fail to load because the configuration might be larger than the maximum size of 5 MB. The error message that appears is **Configuration too large to handle**. [PR/552549: This is a known software limitation.]
- If you have accessed the J-Web interface using an HTTPS connection through the Microsoft Internet Explorer Web browser, you might not be able to download and save reports from some pages on the Monitor, Maintain, and Troubleshoot tabs. Some affected pages are at these locations:
 - Maintain > Files > Log Files > Download
 - Maintain > Config Management > History
 - Maintain > Customer Support > Support Information > Generate Report
 - Troubleshoot > Troubleshoot Port > Generate Report
 - Monitor > Events and Alarms > View Events > Generate Report
 - Monitor > Routing > Route Information > Generate Report

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

- The J-Web interface does not support role-based access control; it supports only users in the super-user authorization class. So a user who is not in the super-user class, such as a user with view-only permission, is able to launch the J-Web interface and is allowed to configure everything, but the configuration fails on the switch, and the switch displays access permission errors. [PR/604595: This is a known software limitation.]
- In mixed EX4200 and EX4500 Virtual Chassis, the J-Web interface does not list the features supported by the backup or linecard members. Instead, it lists only the features supported by the master. [PR/707671: This is a known software limitation.]
- After you remove or reboot a Virtual Chassis member (either the backup or a member in the linecard role), when you click other members in the J-Web interface, the chassis view for those members might not expand, and the dashboard might log the following

error: **stackimg is null or not an object**. As a workaround, manually refresh the dashboard. [PR/771415: This is a known software limitation.]

- If a Virtual Chassis contains more than six members, the Support Information page (Maintain > Customer Support > Support information) might not load. [PR/777372: This is a known software limitation.]
- On EX Series Virtual Chassis that have more than five members, logging in to the J-Web dashboard might take more than 30 seconds. [PR/785300: This is a known software limitation.]
- For EX Series switches, in the J-Web interface, the username field on the Login screen does not accept HTML tags or the < and > characters. The following error message appears: **A username cannot include certain characters, including < and >**. [This is a known software limitation.]
- When you use an HTTPS connection in the Microsoft Internet Explorer browser to save a report from some pages in the J-Web interface, the error message **Internet Explorer was not able to open the Internet site** is displayed. This problem occurs because the Cache-Control: no cache HTTP header is added on the server side, and Internet Explorer does not allow you to download the encrypted file with the Cache-Control: no cache HTTP header set in the response from the server.

As a workaround, refer to Microsoft Knowledge Base article 323308, which is available at <http://support.microsoft.com/kb/323308>. Alternatively, use HTTP in the Internet Explorer browser or use HTTPS in the Mozilla Firefox browser to save a file from one of these pages. [This is a known software limitation.]

- On EX2200-C switches, both the copper and the fiber uplink ports display as connected in the J-Web dashboard if either is connected. [PR/862411: This is a known software limitation.]

Layer 2 and Layer 3 Protocols

- On EX3200 and EX4200 switches, MPLS is not supported on Layer 3 tagged subinterfaces and routed VLAN interfaces (RVIs), even though the CLI allows you to commit a configuration that enables these features. [PR/612434: This is a known software limitation.]

Management and RMON

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

Multicast Protocols

- On EX9200 switches, multicast traffic might be momentarily duplicated on an mrouter port (the port that connects to a multicast router) when a new member is added to an aggregated Ethernet bundle (or link aggregation group [LAG]) and when that new member is in the Detached state. [PR/848390: This is a known software limitation.]

Software Installation and Upgrade

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

Virtual Chassis

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

The PIC mode on an EX4500 switch might have been set to virtual-chassis in one of the following ways:

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

To check the PIC mode for an EX4500 switch that has a Virtual Chassis module installed in it, use the **show chassis pic-mode** command.

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

[This is a known software limitation.]

- The automatic software update feature is not supported on EX4500 switches that are members of a Virtual Chassis. [PR/541084: This is a known software limitation.]
- When an EX4500 switch becomes a member of a Virtual Chassis, it is assigned a member ID. If that member ID is a nonzero value, then if that member switch is downgraded to a software image that does not support Virtual Chassis, you cannot change the member ID to 0. A standalone EX4500 switch must have a member ID of 0. The workaround is to convert the EX4500 Virtual Chassis member switch to a standalone EX4500 switch before downgrading the software to an earlier release, as follows:
 1. Disconnect all Virtual Chassis cables from the member to be downgraded.
 2. Convert the member switch to a standalone EX4500 switch by issuing the **request virtual-chassis reactivate** command.

3. Renumber the member ID of the standalone switch to 0 by issuing the **request virtual-chassis renumber** command.
4. Downgrade the software to the earlier release.

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

- When you add a new member switch to an EX4200 Virtual Chassis, EX4500 Virtual Chassis, or mixed EX4200 and EX4500 Virtual Chassis in a ring topology, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. The member switch returns to the operational state with no user intervention. Network traffic to the member switch is dropped during the downtime. To avoid this issue, follow this procedure:
 1. Cable one dedicated or user-configured Virtual Chassis port (VCP) on the new member switch to the existing Virtual Chassis.
 2. Power on the new member switch.
 3. Wait for the new switch to become operational in the Virtual Chassis. Monitor the **show virtual-chassis** command output to confirm the new switch is recognized by the Virtual Chassis and is in the Prsnt state.
 4. Cable the other dedicated or user-configured VCP on the new member switch to the Virtual Chassis.

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

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

Outstanding Issues in Junos OS Release 12.3 for EX Series Switches

The following are outstanding issues in Junos OS Release 12.3R11 for EX Series switches. The identifier following the description is the tracking number in our bug database.

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



NOTE: Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in “[Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers](#)” on page 216.

Access Control and Port Security

- On EX9200 switches, the LLDP database is not updated when you change the interface description or system name. [PR/848320]
- On EX9200 switches, an LLDP neighbor is not formed for Layer 3-tagged interfaces, although peer switches are able to form the neighbor. [PR/848721]

Ethernet Switching and Spanning Trees

- On EX9200 switches running the VLAN Spanning Tree Protocol (VSTP), incoming BPDUs might not be included in the output of the **show spanning-tree statistics interface** command. [PR/847405]
- On EX Series switches, if two interface ranges are configured on the same access interface with different VLANs for each range, the commit check succeeds instead of detecting an error in the configuration. [PR/957178]

High Availability

- On EX Series switches, there is no RPC equivalent for the CLI command **show chassis nonstop-upgrade**. [PR/872587]

Infrastructure

- When multicast traffic is transiting an EX8200 switch, kernel panic might occur on a new master Routing Engine, causing the string **rn_clone_unwire parent unreferenced** to be displayed during a nonstop software upgrade (NSSU) or after multiple graceful Routing Engine switchover (GRES) operations. [PR/734295]
- On EX9200 switches, if you configure multichassis link aggregation (MC-LAG) in active-active mode and then one of the provider edge (PE) routers is rebooted, Layer 3 traffic might be lost for more than 10 seconds. This loss occurs because, by default, whenever a peer PE router is rebooted, the Link Aggregation Control Protocol (LACP) system ID on the other peer router changes.

As a workaround, designate one of the PE routers to be **status-control active** using the following command on the PE router: **set interfaces aex aggregated-ether-options mc-ae events iccp-peer-down prefer-status-control-active**. As a result of issuing this command on the PE router, when a peer PE router goes down, the LACP system ID does not change.



NOTE: To configure the **prefer-status-control-active** statement, you must configure the **status-control active** statement. Do not configure **status-control as standby**.

[PR/853694]

- On EX Series switches, the **wildcard range set interfaces interface-name unit logical-unit-number vlan-id-list [vlan-id, vlan-id-vlan-id]** CLI command does not work. An error message is displayed on the console if you issue that command. As a

workaround, configure the VLAN list by using expanded configuration commands. [PR/1030369]

- On EX9200 switches, recurring LMEM data errors might cause a chip wedge. [PR/1033660]
- On EX Series switches, **fpc0 dfw_counter_get_by_name failed inst 0 policer index 0 status 7** error messages might appear in the log files when **show firewall counter** or **snmp mib get jnxFirewallCounterTable** is executed. [PR/1035113]
- On EX Series switches, if you issue the **show ethernet-switching table | display xml** CLI command, the **<l2ng-mac-entry>** tag does not appear in the output. [PR/1097532]

Interfaces

- If you configure a VLAN range on an access interface of an EX9200 switch, the Layer 2 address learning process (l2ald) might fail. As a workaround, do not configure VLAN ranges on EX9200 switch access interfaces. [PR/837608]

J-Web Interface

- In the J-Web interface on EX4200 switches, if you try to change the position of columns using the drag-and-drop method, only the column header moves to the new position instead of the entire column in the OSPF Global Settings table in the OSPF Configuration page, the Global Information table in the BGP Configuration page, or the Add Interface window in the LACP Configuration page. [PR/465030]
- If you configure an IPv6 address for a VLAN in the J-Web interface, you cannot then edit the VLAN configuration. [PR/466633]
- When a large number of static routes are configured and you have navigated to pages other than page 1 in the Route Information table in the Static Routing monitoring page in the J-Web interface (Monitor > Routing > Route Information), changing the Route Table to query other routes refreshes the page but does not return to page 1. For example, if you run a query from page 3 and the new query returns very few results, the Results table continues to display page 3 and shows no results. To view the results, navigate to page 1 manually. [PR/476338]
- If you access the J-Web interface by using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR/603669]
- On the process details page of the J-Web interface (Monitor > System View > Process Details), there are multiple entries listed for a few processes that do not impact any functionality. [PR/661704]
- In the J-Web interface, the Next Hop column in the Static Routing page (Monitor > Routing > Route Information) displays only the interface address; the corresponding IP address is missing. The title of the first column displays Static Route Address instead of Destination Address. As a workaround, use the CLI to execute the **show route detail** command to fetch the corresponding next-hop interface IP address. [PR/684552]

- In the J-Web interface, HTTPS access might work with an invalid certificate. As a workaround, after you change the certificate, issue the **restart web-management** command to restart the J-Web interface. [PR/700135]
- On EX2200-C switches, if you have changed the media type and committed the change, the Ports configuration page (Configure > Interfaces > Ports) might not list the uplink port. [PR/742847]
- On EX8200 Virtual Chassis, if you are using the Virtual Chassis Wizard in the J-Web interface in the Mozilla Firefox version 3.x browser, if you have selected more than six port pairs from the same member for conversion, the wizard might display the incorrect port conversion status. Also, if you double-click **Next** after deleting an active member in the Members page, the J-Web interface might stop working. [PR/796584]

Layer 2 and Layer 3 Protocols

- On EX Series switches, if a session is initiated with a peer that is not configured, and the peer AS is a member of a configuration, then an rpd core file is created. As a workaround, use explicitly configured peers for peers in the confederation AS. [PR/963565]

Multicast Protocols

- On EX9200 switches, the mcsnoopd process creates multiple core files at `rt_nexthops_free` during graceful Routing Engine switchover (GRES) with Layer 3 multicast traffic. [PR/848732]
- On EX9200 switches, Layer 3 multicast traffic loss might occur for about 20 seconds on the switch when the switch is acting as the last-hop router (LHR) and the software performs a graceful Routing Engine switchover (GRES). [PR/848861]
- If you configure a large number of PIM source-specific multicast (SSM) groups on an EX9200 switch, the switch might experience periodic IPv6 traffic loss. As a workaround, configure the **pim-join-prune-timeout** value on the last-hop router (LHR) to 250 seconds. [PR/853586]
- In the J-Web interface on EX Series switches, you cannot configure OSPFv3 by using the point-and-click function (Configure > Point&Click > Protocols > Configure > Ospf3). As a workaround, configure OSPFv3 options by using the CLI. You can then view and edit the OSPFv3 parameters by using the point-and-click function in the J-Web interface. [PR/857540]

Network Management and Monitoring

- EX Series switches do not notify users that a system log event has occurred. [PR/897200]

Routing Policy and Firewall Filters

- On EX Series switches, if the switch receives BFD control packets larger than the maximum supported size of 256 bytes, the BFD process (bfd) might generate a core file. [PR/1004482]

Software Upgrade and Installation

- On EX8200 Virtual Chassis, when an NSSU is initiated to upgrade to Junos OS Release 12.3R5, multiple pfem core files might be created on some member switches. [PR/917863]

Virtual Chassis

- On an EX9200 Virtual Chassis, with the Virtual Chassis member switches having multiple aggregated Ethernet (ae) interfaces configured for load balancing, if you reconfigure a Virtual Chassis port (VCP) as a network-traffic Ethernet port, you might see permanent traffic losses for Layer 3 traffic that transits the aggregated Ethernet interfaces. [PR/895058]
- On EX4550 Virtual Chassis, the output of the CLI command **show virtual-chassis vc-port** indicates the speed of each VCP port is 32G and bi-directional, but it should be 16G and uni-directional. [PR/913523]

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

Resolved Issues in Junos OS Release 12.3 for EX Series Switches

The following issues have been resolved in Junos OS Release 12.3 for EX Series switches. The identifier following the descriptions is the tracking number in our bug database.

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



NOTE: Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in “[Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers](#)” on page 216.

- [Issues Resolved in Release 12.3R1 on page 55](#)
- [Issues Resolved in Release 12.3R2 on page 65](#)
- [Issues Resolved in Release 12.3R3 on page 67](#)
- [Issues Resolved in Release 12.3R4 on page 72](#)
- [Issues Resolved in Release 12.3R5 on page 74](#)
- [Issues Resolved in Release 12.3R6 on page 77](#)
- [Issues Resolved in Release 12.3R7 on page 81](#)
- [Issues Resolved in Release 12.3R8 on page 84](#)
- [Issues Resolved in Release 12.3R9 on page 87](#)
- [Issues Resolved in Release 12.3R10 on page 92](#)
- [Issues Resolved in Release 12.3R11 on page 93](#)

Issues Resolved in Release 12.3R1

The following issues have been resolved since Junos OS Release 12.2. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- For LLDP, the values for the IEEE 802.3 - MAC/PHY Configuration/Status TLV might be incorrect. [PR/607533: This issue has been resolved.]
- If a Unified Access Control (UAC) infranet controller is unreachable, an 802.1X (dot1x) interface might not be able to access the server-fail VLAN. [PR/781586: This issue has been resolved.]
- If you enable 802.1X with MAC RADIUS authentication, that is, by including the **mac-radius** statement in the configuration, the authentication management process (authd) might reach a memory limit when there are approximately 250 users. As a workaround, reset the authd process when it reaches 85 percent of its RLIMIT_DATA value (that is, 85 percent of 130 MB). To check the amount of memory being used by the authd process, use the **show system processes extensive** operational mode command. [PR/783363: This issue has been resolved.]
- When access configuration is not required and the guest VLAN feature is configured, supplicants might not be able to authenticate using the guest VLAN and remain in the *connecting* state. [PR/783606: This issue has been resolved.]
- DHCP snooping might not allow DHCP Inform ACK packets to pass to the client. [PR/787161: This issue has been resolved.]
- If you configure a static MAC bypass for 802.1X (dot1x) and you add a new host to the exclusion list, the MAC addresses of existing hosts that have already been successfully

authenticated using static MAC bypass might move to an incorrect VLAN. [PR/787679: This issue has been resolved.]

- Traffic leaks might occur for unknown unicast and broadcast traffic from multiple VLANs when a MAC-RADIUS-assigned VLAN is set on a switch interface through a server-initiated attribute change. If the 802.1X interface has VLAN 100 assigned and the RADIUS server sends a different VLAN attribute (for example, 200 rather than 100), after the interface is assigned in VLAN 200, it also sends egress unknown unicast and broadcast traffic that belongs to VLAN 100. [PR/829436: This issue has been resolved.]
- On EX6200 switches, LLDP stops working if you execute the **set ethernet-switching-options voip interface access-ports vlan** command. [PR/829898: This issue has been resolved.]

Class of Service

- When you are configuring class-of-service (CoS) drop profiles, the commit operation might fail and might display the message **Missing mandatory statement: 'drop-probability'**. [PR/807885: This issue has been resolved.]

Converged Networks (LAN and SAN)

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

Ethernet Switching and Spanning Trees

- When you enable Q-in-Q tunneling and MLD snooping, no snooping database is present on the switch. [PR/693224: This issue has been resolved.]
- If a VLAN change occurs quickly, the client might not be able to obtain an IP address. [PR/746479: This issue has been resolved.]
- When you add a new virtual routing and forwarding (VRF) instance, existing firewall filters might not be applied to the new VRF instance. [PR/786662: This issue has been resolved.]
- You cannot configure a VLAN whose name contains a hyphen (-). As a workaround, use an underscore (_) in the name instead. [PR/753090: This issue has been resolved.]
- Ethernet ring protection switching (ERPS; G.8032) does not block PVST BPDUs. [PR/793891: This issue has been resolved.]
- If you delete an IPv6 configuration on a routed VLAN interface (RVI), ARP requests might not be trapped to the CPU and are not resolved. As a workaround, delete the RVI and then reconfigure it, or reboot the switch after you delete the IPv6 configuration. [PR/826862: This issue has been resolved.]
- After a software upgrade on the switch, Spanning Tree Protocol (STP) might not be distributed on some aggregated Ethernet links. [PR/822673: This issue has been resolved.]

Firewall Filters

- On all EX Series switches except EX8200 switches, if you have configured several policer settings in the same filter, they might all be overwritten when you change one of the settings. As a workaround, delete the setting and then add it back again with the desired changes. [PR/750497: This issue has been resolved.]
- On EX8200 Virtual Chassis, if you add and delete a firewall filter for traffic that enters on one Virtual Chassis member and is transmitted out another member, IPv6 traffic might be dropped. If the ingress and egress interfaces are on the same member, the firewall filter works correctly. [PR/803845: This issue has been resolved.]
- On EX8200 Virtual Chassis, when both dscp and ieee-802.1 rewrite rules are applied on a routed VLAN interface (RVI), deleting the filters and binding again on the same RVI or clearing interface statistics might create a pfem core file. [PR/828661: This issue has been resolved.]

Hardware

- When you remove the hard drive from an XRE200 External Routing Engine, an SNMP trap and a system alarm might not be generated. [PR/710213: This issue has been resolved.]
- Non-Juniper Networks DAC cables do not work on EX Series switches. [PR/808139: This issue has been resolved.]
- On EX4200 switches, high CPU usage might be due to console cable noise. [PR/818157: This issue has been resolved.]
- On EX4550 switches, the backlight on the LCD panel does not turn on. [PR/820473: This issue has been resolved.]
- When an uplink module in the switch is operating in 1-gigabit mode, a chassism core file might be created if you remove an SFP transceiver from one of the module's interfaces. As the chassism process restarts, all traffic passing through the interface is dropped. This problem happens with both copper and fiber SFPs. [PR/828935: This issue has been resolved.]

High Availability

- On an XRE200 External Routing Engine, when you perform a nonstop software upgrade (NSSU) operation that includes the **reboot** option, the physical link might flap, which causes traffic loss and protocol flapping. [PR/718472: This issue has been resolved.]
- After you perform a nonstop software upgrade (NSSU), you might notice a traffic outage of 150 seconds while the line cards are restarting. [PR/800460: This issue has been resolved.]

Infrastructure

- If you enable gratuitous ARP by including the **gratuitous-arp-reply**, **no-gratuitous-arp-reply**, or **no-gratuitous-arp-request** statement in the configuration, the switch might process gratuitous ARP packets incorrectly. [PR/518948: This issue has been resolved.]
- The output of the **show system users no-resolve** command displays the resolved hostname. [PR/672599: This issue has been resolved.]
- Rate limiting for management traffic (namely, FTP, SSH, and Telnet) arriving on network ports causes file transfer speeds to be slow. [PR/691250: This issue has been resolved.]
- In some cases, broadcast traffic that is received on the management port (me0) is broadcast to other subnets on the switch. [PR/705584: This issue has been resolved.]
- The **allow-configuration-regexps** statement at the **[edit system login class]** hierarchy level does not work exactly the same way as the deprecated **allow-configuration** statement at the same hierarchy level. [PR/720013: This issue has been resolved.]
- When you delete the VLAN mapping for an aggregated Ethernet (ae) interface, the Ethernet switching process (eswd) might crash and display the error message **No vlan matches vlan tag 116 for interface ae5.0**. [PR/731731: This issue has been resolved.]
- The **wildcard range unprotect** configuration statement might not be synchronized with the backup Routing Engine. [PR/735221: This issue has been resolved.]
- After you successfully install Junos OS, if you uninstall AI scripts, an mgd core file might be created. [PR/740554: This issue has been resolved.]
- When there is a large amount of NetBIOS traffic on the network, the switch might exhibit high latency while pinging between VLANs. [PR/748707: This issue has been resolved.]
- On EX4200 switches, a Packet Forwarding Engine process (pfem) core file might be created while the switch is running the Packet Forwarding Engine internal support script and saving the output to a file. [PR/749974: This issue has been resolved.]
- You might see the following message in log files: **Kernel/ (COMPOSITE NEXT HOP) failed, err 6 (No Memory)**. [PR/751985: This issue has been resolved.]
- On EX3300 switches, if you configure more than 20 BGPv6 neighbor sessions, the CLI might display the db> prompt. [PR/753261: This issue has been resolved.]
- On EX8200 switches, the master-only configuration for the management interface does not work. [PR/753765: This issue has been resolved.]
- The Junos OS kernel might crash because of a timing issue in the ttymodem() internal I/O processing routine. The crash can be triggered by simple remote access (such as Telnet or SSH) to the device. [PR/755448: This issue has been resolved.]
- On EX Series switches, after a flash memory initialization process for the **/var** or **/var/tmp** directory has been caused by severe corruption, SSH and HTTP access might not work correctly. As a workaround for SSH access, create a **/var/empty** folder. [PR/756272: This issue has been resolved.]

- On EX8200 switch line cards, a Packet Forwarding Engine process (pfem) core file might be created as the result of a memory segmentation fault. [PR/757108: This issue has been resolved.]
- EX4500 switches and EX8200-40XS line cards do not forward IP UDP packets when their destination port is 0x013f (PTP) or when the fragmented packet has the value 0x013f at the same offset (0x2c). [PR/775329: This issue has been resolved.]
- After you upgrade to Junos OS Release 11.4R3, EX Series switches might stop responding to SNMP ifIndex list queries. As a workaround, restart the switch. If restarting the switch is not an option, restart the shared-memory process (shm-rtssdbd). [PR/782231: This issue has been resolved.]
- When EX Series switches receive packets across a GRE tunnel, they might not generate and send ARP packets to the device at the other end of the tunnel. [PR/782323: This issue has been resolved.]
- On EX4550 switches, if you configure the management (me0) interface and a static route, the switch is unable to connect to a gateway. [PR/786184: This issue has been resolved.]
- After you remove an IPv6 interface configuration and then perform a rollback operation, the IPv4 label might change to explicit null. [PR/786537: This issue has been resolved.]
- When many packets are queued to have their next hop resolved, some packets might become corrupted. [PR/790201: This issue has been resolved.]
- If you configure IPv6 and VRRP, the IPv6 VRRP MAC address might be used incorrectly as the source MAC address when the switch routes traffic across VLANs. [PR/791586: This issue has been resolved.]
- The `/var/log/messages` file might fill up with the following message: **caff_sf_rd_reg ret:00000 slot:1 chip:1 addr:02b45c data:0**. [PR/792396: This issue has been resolved.]
- When you restart a line card, the BFD session might go down. [PR/793194: This issue has been resolved.]
- After the system has been up for days, EX8200 line cards might reach 100 percent CPU usage and then stay at 100 percent. [PR/752454: This issue has been resolved.]
- On an EX8200 Virtual Chassis, the dedicated Virtual Chassis port (VCP) link between the XRE200 External Routing Engine and the Routing Engine on a member switch might be down after an upgrade. As a workaround, manually disable and then enable the physical link. [PR/801507: This issue has been resolved.]
- After you upgrade Junos OS, a pppmd core file might be created, and protocols that use pppmd might not work correctly. [PR/802315: This issue has been resolved.]
- On EX3300 switches, when you are configuring BGP authentication, after you have configured the authentication key, BGP peering is never established. [PR/803929: This issue has been resolved.]
- An EX6200 switch might send 802.1Q tagged frames out of access ports when DHCP snooping is configured. This might prevent certain vendors' end devices from receiving proper IP addresses from the DHCP server. [PR/804010: This issue has been resolved.]

- On EX Series switches that have Power over Ethernet (PoE) capability, chassisd (the chassis process) might crash when running SNMP requests (for example, SNMP get, get-next, and walk requests) on pethMainPse objects. This is caused by the system trying to free memory that is already freed. As a workaround, avoid running SNMP requests on pethMainPse objects. [PR/817311: This issue has been resolved.]
- If you reboot the switch with the routed VLAN interface (RVI) disabled, then even if you reenables the RVI, the RVI traffic is not routed in the Packet Forwarding Engine; the traffic is trapped to the CPU and is policed by the rate limit in the Packet Forwarding Engine. [PR/838581: This issue has been resolved.]

Interfaces

- EX4200 and EX4500 switches support 64 aggregated Ethernet interfaces even though the hardware can support 111 interfaces. [PR/746239: This issue has been resolved.]
- When VRRP is running between two EX8200 switches on a VLAN, after a master switchover, both switches might act as master. [PR/752868: This issue has been resolved.]
- After you change the physical speed on a Virtual Chassis member interface, an aggregated Ethernet (ae) interface might flap after you issue the next **commit** command to commit configuration changes. [PR/779404: This issue has been resolved.]
- On EX4500 switches, link-protection switchover or revert might not work as expected. [PR/781493: This issue has been resolved.]
- On aggregated Ethernet (ae) interfaces, the Link Layer Discovery Protocol (LLDP) might not work. [PR/781814: This issue has been resolved.]
- When you issue the **show vrrp brief** command, a VRRP process (vrrpd) core file might be created. [PR/782227: This issue has been resolved.]
- On EX8200 switches, when you issue the **request system reboot other-routing-engine** command, a timeout error might be displayed before the Routing Engine initiates its reboot operation. [PR/795884: This issue has been resolved.]
- On EX4550 switches, link autonegotiation does not work on 1-Gb SFP interfaces. [PR/795626: This issue has been resolved.]
- On EX Series switches, if you have configured a link aggregation group (LAG) with link protection, an interface on the backup member might drop ingress traffic. [PR/796348: This issue has been resolved.]
- If you apply a policer to an interface, the policer might not work, and messages similar to the following are logged: **dfw_bind_policer_template_to_filter:205 Binding policer fails**. [PR/802489: This issue has been resolved.]
- An interface on an EX4550-32F switch might go up and down randomly even when no cable is plugged in. [PR/803578: This issue has been resolved.]
- On EX3300 switches, when you configure VRRP with MD5 authentication with the **preempt** option on a routed VLAN interface (RVI), a vmcore file might be created. As

a workaround, delete the **preempt** option and disable MD5 authentication for VRRP. [PR/808839: This issue has been resolved.]

- On EX4550 Virtual Chassis, the **show chassis environment power-supply-unit** operational mode command does not show the power supply status of all member interfaces. Use the **show chassis hardware** command instead. [PR/817397: This issue has been resolved.]

J-Web Interface

- In the J-Web interface, you cannot upload a software package using the HTTPS protocol. As a workaround, use either the HTTP protocol or the CLI. [PR/562560: This issue has been resolved.]
- In the J-Web interface, the link status might not be displayed correctly in the Port Configuration page or the LACP (Link Aggregation Control Protocol) Configuration page if the Commit Options preference is set to *single commit* (the Validate configuration changes option). [PR/566462: This issue has been resolved.]
- If you have created dynamic VLANs by enabling MVRP from the CLI, then in the J-Web interface, the following features do not work with dynamic VLANs or static VLANs:
 - In the Port Configuration page (Configure > Interface > Ports)—Port profile (select the interface, click **Edit**, and select **Port Role**) or the VLAN option (select the interface, click **Edit**, and select **VLAN Options**).
 - VLAN option in the LACP (Link Aggregation Control Protocol) Configuration page (Configure > Interface > Link Aggregation)—Select the aggregated interface, click **Edit**, and click **VLAN**.
 - In the 802.1X Configuration page (Configure > Security > 802.1x)—VLAN assignment in the exclusion list (click **Exclusion List** and select **VLAN Assignment**) or the move to guest VLAN option (select the port, click **Edit**, select **802.1X Configuration**, and click the **Authentication** tab).
 - Port security configuration (Configure > Security > Port Security).
 - In the Port Mirroring Configuration page (Configure > Security > Port Mirroring)—Analyzer VLAN or ingress or egress VLAN (click **Add** or **Edit** and then add or edit the VLAN).

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

- On EX4500 Virtual Chassis, if you use the CLI to switch from virtual-chassis mode to intraconnect mode, the J-Web interface dashboard might not list all the Virtual Chassis hardware components, and the image of the master and backup switch chassis might not be visible after an autorefresh occurs. The J-Web interface dashboard also might not list the vcp-0 and vcp-1 Virtual Chassis ports in the rear view of an EX4200 switch (in the linecard role) that is part of an EX4500 Virtual Chassis. [PR/702924: This issue has been resolved.]
- The J-Web interface is vulnerable to HTML cross-site scripting attacks, also called XST or cross-site tracing. [PR/752398: This issue has been resolved.]

- When you configure the **no-tcp-reset** statement, the J-Web interface might be slow or unresponsive. [PR/754175: This issue has been resolved.]
- In the J-Web interface, you cannot configure the TCP fragment flag for a firewall filter in the Filters Configuration page (Configure > Security > Filters). [PR/756241: This issue has been resolved.]
- In the J-Web interface, you cannot delete a term from a filter and simultaneously add a new term to that filter in the Filters configuration page (Configure > Security > Filters). [PR/769534: This issue has been resolved.]
- Some component names shown by the tooltip on the Temperature in the Health Status panel of the dashboard might be truncated. As a result, you might see many components that have the same name displayed. For example, the components GEPHY Front Left, GEPHY Front Middle, and GEPHY Front Right might all be displayed as GEPHYFront. [PR/778313: This issue has been resolved.]
- In the J-Web interface, the Help page for the Install package in the Software Maintenance page (Maintain > Software) might not appear. [PR/786654: This issue has been resolved.]
- If you issue the **set protocols rstp interface *logical-interface-name* edge** configuration command from the CLI, the J-Web interface might show that the configuration in the Configuration detail for Desktop and Phone window is not applicable for the port profile. However, no functionality for the Desktop and Phone port profile is affected. [PR/791323: This issue has been resolved.]
- In the J-Web interface, if you enable a spanning-tree protocol (STP, RSTP, or MSTP) and then exclude some ports from the spanning tree, you might not be able to include these ports as part of a redundant trunk group (RTG). [PR/791759: This issue has been resolved.]
- In the J-Web interface on EX4500 and EX4550 switches, you can configure temporal and exact-temporal buffers, which are not supported by Junos OS. [PR/796719: This issue has been resolved.]
- In a mixed Virtual Chassis in which an EX4550 switch is the master and at least one Virtual Chassis member supports Power over Ethernet (PoE), if you click **Configure > POE** and then click another tab, a javascript error might be displayed. [PR/797256; This issue has been resolved.]
- In the J-Web interface on EX4550 switches, if you are using in-band management and select EZSetup, the error message **undefined configuration delivery failed** is displayed even though the configuration has been successfully committed. [PR/800523: This issue has been resolved.]
- On EX2200 switches, in the dashboard in the J-Web interface, the flash memory utilization graph might show an incorrect value of 0%. As a workaround, to view utilization, click **Monitor > System View > System Information** and then click the **Storage Media** tab. [PR/823795: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- On EX8200 switches with OSPF configured, after a nonstop software upgrade (NSSU) to Junos OS Release 12.1R1, OSPF adjacency might not be established for some RVIs

across link aggregation group (LAG) interfaces because the flooding entry is not programmed correctly. As a workaround, disable or enable the problematic interface by issuing the following commands:

- **user@switch# set interface *interface-name* disable**
- **user@switch# delete interface *interface-name* disable**

[PR/811178: This issue has been resolved.]

- A BFD session might flap if there are stale BFD entries. [PR/744302: This issue has been resolved.]
- On XRE200 External Routing Engines on which PIM is configured, a nonstop software upgrade (NSSU) operation might fail when performed when an MSDP peer is not yet up. As a workaround, either disable nonstop active routing (NSR) for PIM using the **set protocols pim nonstop-routing disable** configuration command or ensure that MSDP has reached the Established state before starting an NSSU operation. [PR/799137: This issue has been resolved.]
- Multicast packets might be lost when the user switches from one IPTV channel to another. [PR/835538: This issue has been resolved.]

Management and RMON

- On EX8200 Virtual Chassis, when you perform an snmpwalk operation on the jnxPsuMIB, the output shows details only for the power supplies on a single line card member. [PR/689656: This issue has been resolved.]
- When you are using IS-IS for forwarding only IPv6 traffic and IPv4 routing is not configured, if you perform an SNMP get or walk operation on an IS-IS routing database table, the routing protocol process (rpd) might crash and restart, possibly causing a momentary traffic drop. [PR/753936: This issue has been resolved.]
- When an SNMP string is longer than 30 characters, it is not displayed in Junos OS command output. [PR/781521: This issue has been resolved.]
- The incorrect ifType might be displayed for counters on physical interfaces. [PR/784620: This issue has been resolved.]
- For sFlow monitoring technology traffic on the switches, incorrect information might be displayed for output ports. [PR/784623: This issue has been resolved.]
- After a Routing Engine switchover, LACP and MIB II process (mib2d) core files might be created. [PR/790966: This issue has been resolved.]
- An SNMP MIB walk might show unwanted data for newly added objects such as jnxVirtualChassisPortInPkts or jnxVirtualChassisPortInOctets. [PR/791848: This issue has been resolved.]
- On EX Series switches, sFlow monitoring technology packets might be dropped when the packet size exceeds 1500 bytes. [PR/813879: This issue has been resolved.]
- In EX3300 Virtual Chassis, if you perform an SNMP poll of jnxOperatingState for fan operation, the information for the last two members in the Virtual Chassis is incorrect. [PR/813881: This issue has been resolved.]

- On EX8200 switches, sFlow monitoring technology packets were being generated with an incorrect source MAC address of 20:0b:ca:fe:5f:10. This issue has been fixed, and the EX8200 switches now use the outbound port's MAC address as the source MAC address for sFlow monitoring technology traffic. [PR/815366: This issue has been resolved.]
- An SNMP poll might not return clear information for some field-replaceable units (FRUs), such as fans and power supplies. The FRU description might not indicate which physical switch contains the FRU. [PR/837322: This issue has been resolved.]

Multicast Protocols

- When an EX Series switch is routing multicast traffic, that traffic might not exit from the multicast router port in the source VLAN. [PR/773787: This issue has been resolved.]
- While multicast is resolving routes, the following SPF-related error might be displayed: **SPF:spf_change_sre(),383:jt_change() returned error-code (Not found:4)!** [PR/774675: This issue has been resolved.]
- On EX8200 switches, multicast MDNS packets with the destination address 224.0.0.251 are blocked if IGMP snooping is enabled. [PR/782981: This issue has been resolved.]
- In MPLS implementations on EX Series switches, EXP bits that are exiting the provider edge switch are copied to the three least-significant bits of DSCP—that is, to IP precedence—rather than to the most-significant bits. [PR/799775: This issue has been resolved.]

Power over Ethernet (PoE)

- Power over Ethernet (PoE) and Power over Ethernet Plus (PoE+) cannot be configured by using the EX8200 member switches in an EX8200 Virtual Chassis. [PR/773826: This issue has been resolved.]

Software Installation and Upgrade

- EX4550 switches might not load the configuration file after you perform an automatic image upgrade. [PR/808964: This issue has been resolved.]
- On EX8200 Virtual Chassis, nonstop software upgrade (NSSU) with the no-reboot option is not supported. [PR/821811: This issue has been resolved.]

Virtual Chassis

- On EX8200 Virtual Chassis, when you swap the members of a link aggregation group (LAG), a vmcore or ksyncd core file might be created on the backup Routing Engine. [PR/711679: This issue has been resolved.]
- On EX8200 Virtual Chassis, after you ungracefully remove the master Routing Engine from the member switch, traffic might be interrupted for up to 2 minutes. [PR/742363: This issue has been resolved.]
- On EX3300 switches, when a Virtual Chassis is formed, the Virtual Chassis backup member's console CLI is not automatically redirected to the Virtual Chassis master's console CLI. As a workaround, manually log out from the Virtual Chassis backup member. [PR/744241: This issue has been resolved.]

- On EX8200 Virtual Chassis, the **request system snapshot** command does not take a snapshot on the backup Routing Engine of both members. [PR/750724: This issue has been resolved.]
- On EX8200 Virtual Chassis, the switch might incorrectly send untagged packets. As a result, some hosts in the VLAN might experience connectivity issues. [PR/752021: This issue has been resolved.]
- On EX8200 Virtual Chassis, after one Virtual Chassis member is rebooted, the line card of the corresponding rebooted member switch is not brought down immediately, and hence the peer sees that the interfaces remain in the Up state. Additionally, the interface state is not cleared immediately in the switch card chassis kernel. The result is that the protocol session goes down, and traffic loss occurs even if you have configured nonstop active routing (NSR). [PR/754603: This issue has been resolved.]
- On XRE200 External Routing Engines, when you issue the **show chassis hardware** command and specify **display xml**, duplicate occurrences of the **<name>** and **<serial-number>** tags under the **<chassis>** tag might result in malformed XML output. [PR/772507: This issue has been resolved.]
- In a mixed EX4200 and EX4500 Virtual Chassis, the master chassis view might display the temperature indicator of the backup. [PR/783052: This issue has been resolved.]
- On XRE200 External Routing Engines, a chassis core file might be created. [PR/791959: This issue has been resolved.]
- On EX8200 Virtual Chassis, when you swap the members of a link aggregation group (LAG), a **vmcore** or **ksyncd** core file might be created on the backup Routing Engine. [PR/793778: This issue has been resolved.]
- On XRE200 External Routing Engines on which DHCP snooping and dynamic ARP inspection are enabled, when packets are transmitting out a different line card type from the ingress interface, an **sfid** core file might be created. [PR/794293: This issue has been resolved.]
- On EX8200 Virtual Chassis, the devbuf process might leak memory, eventually bringing the switch down to a halt. As a workaround, perform a hard shutdown by issuing the **ifconfig em[0-8] down** command on the em interfaces that are in the down state. [PR/823045: This issue has been resolved.]

Issues Resolved in Release 12.3R2

The following issues have been resolved since Junos OS Release 12.3R1. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- On EX Series switches except EX9200, with 802.1X user-based dynamic firewall filters enabled, a stale firewall filter that is properly authenticated after a server timeout might not be purged from an interface even after that interface is disconnected. When this issue occurs, 802.1X authentication fails. [PR/833712: This issue has been resolved.]
- On EX Series switches, the LLDP-MED media endpoint class is shown as invalid. This problem is just a display issue—there is no functional impact. [PR/840915: This issue has been resolved.]

Class of Service

- On EX Series switches, EXP CoS classification does not occur if EXP CoS classifiers are deleted and then added. [PR/848273: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- On an EX4200 switch configured for VLAN translation, Windows NetBIOS traffic might not be translated. [PR/791131: This issue has been resolved.]
- On EX Series switches, the Cisco Discovery Protocol (CDP) and the VLAN Trunking Protocol (VTP) do not work through Layer 2 protocol tunneling (L2PT). [PR/842852: This issue has been resolved.]
- On EX Series switches, the Q-BRIDGE-MIB OID 1.3.6.1.2.1.17.7 reports the VLAN internal index instead of the VLAN ID. [PR/850299: This issue has been resolved.]
- If an EX Series switch has a redundant trunk group (RTG) link, a MAC Refresh message might be sent on a new active link of the RTG when RTG failover occurs. The switch sends the RTG MAC Refresh message with a VLAN tag even though RTGs are configured on access ports. [PR/853911: This issue has been resolved.]

Firewall Filters

- In the case of a stateful proxy, SIP hairpinning does not function, because of which two SIP users behind the NAT device might be unable to connect through a phone call. [PR/832364: This issue has been resolved.]
- On EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6210 switches, a firewall filter with family set to ethernet-switching and configured for IPv4 blocks specific transit IPv6 traffic if the ether_type match condition in the filter is not explicitly set to ipv4. As a workaround, set ether_type to ipv4 in the filter. [PR/843336: This issue has been resolved.]

Infrastructure

- The **unlink** option in the **request system software add package-name unlink** command does not work on EX Series switches. [PR/739795: This issue has been resolved.]
- On EX Series switches, if the sfid process receives dequeuing packets from various queues, the queue indexes do not increment properly, which might cause the sfid process to generate a core file. [PR/835535: This issue has been resolved.]
- On EX8200 switches, multiple rpd process core files might be created on the backup Routing Engine after a nonstop software upgrade (NSSU) has been performed while multicast traffic is on the switch. [PR/841848: This issue has been resolved.]
- On EX8200 switches, the **commit synchronize** command might fail with the error message **error: could not open configuration database (juniper.data+)**. [PR/844315: This issue has been resolved.]

Interfaces

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

Management and RMON

- On EX Series switches, a configured OAM threshold value might be reset when the chassis is rebooted. [PR/829649: This issue has been resolved.]
- An SNMP query or walk on ipNetToMediaPhysAddress does not match the **show arp** command output. [PR/850051: This issue has been resolved.]

Virtual Chassis

- On EX2200 Virtual Chassis, when there are multiple equal-cost paths, the **show virtual-chassis vc-path source-interface *interface-name* destination-interface *interface-name*** command displays the first discovered shortest path, even though traffic might be flowing in an alternate path. [PR/829752: This issue has been resolved.]
- In a mixed EX4200 and EX4500 Virtual Chassis, link aggregation might generate a PFEM core file in some member switches. [PR/846498: This issue has been resolved.]
- On EX4200 Virtual Chassis, **CHASSISD_SNMP_TRAP6: SNMP trap generated: Fan/Blower Removed** messages might be generated periodically, even when member switches cited in the messages are not present in the Virtual Chassis. [PR/858565: This issue has been resolved.]

Issues Resolved in Release 12.3R3

The following issues have been resolved since Junos OS Release 12.3R2. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

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

Class of Service

- On EX Series switches, EXP CoS classification does not occur if EXP CoS classifiers are deleted and then added back. [PR/848273: This issue has been resolved.]
- On EX4500 switches and EX4500 Virtual Chassis, MPLS CoS classifications and rewrites might not work. [PR/869054: This issue has been resolved.]

Ethernet Switching and Spanning Trees

- On EX Series switches, when you issue the **show spanning-tree interface vlan-id vlan-id detail** command, the **vlan-id** parameter is ignored, and the output displays information for all interfaces instead of only for interfaces that are associated with the VLAN ID. [PR/853632: This issue has been resolved.]
- On EX Series switches, when a topology change is detected on an MSTP-enabled interface, there might be a delay of several seconds before a BPDU is sent out with a topology change flag to all the other interfaces. When such a change is detected on RSTP-enabled interfaces, a BPDU is sent out immediately with the topology change flag. [PR/860748: This issue has been resolved.]

Hardware

- EX2200 switches are intermittently not recognizing the Redundant Power System (RPS) after the configuration has been changed and a power supply has been reseated in the RPS. [PR/841785: This issue has been resolved.]
- On EX3200, EX4200, EX8200, EX4500, and EX4550 switches, the receiver signal average optical power is shown as 0.0000 in output for the **show interfaces diagnostics optics** command. [PR/854726: This issue has been resolved.]

High Availability

- On EX8200 Virtual Chassis, a nonstop software upgrade (NSSU) might fail. [PR/871288: This issue has been resolved.]

Infrastructure

- After you successfully install Junos OS, if you uninstall AI scripts, an mgd core file might be created. [PR/740554: This issue has been resolved.]
- Rate limiting for management traffic (namely, SSH and Telnet) arriving on network ports causes file transfer speeds to be slow. [PR/831545: This issue has been resolved.]
- On EX8200 Virtual Chassis, a disabled routed VLAN interface (RVI) might send gratuitous ARP requests. [PR/848852: This issue has been resolved.]
- On EX4200 Virtual Chassis, **CHASSISD_SNMP_TRAP6: SNMP trap generated: Fan/Blower Removed** messages might be generated periodically, even when member switches cited in the messages are not present in the Virtual Chassis. [PR/858565: This issue has been resolved.]
- On EX4500 Virtual Chassis, an **SNMP trap generated for Power Supply Removed** message might be sent for a nonexistent power supply in an active member of the Virtual Chassis. [PR/864635: This issue has been resolved.]
- On EX4200 Virtual Chassis, a **/var partition is full** alarm and a **CHASSISD_RE_CONSOLE_ME_STORM** log might occur, caused by a console error storm, even though the **/var partition** is not full. You can ignore this alarm; it has no effect on the system. [PR/866863: This issue has been resolved.]

Interfaces

- For EX4500 switches, queue counters are not updated for member interfaces of a LAG when the **monitor interface aex** command is running. As a workaround, use the **monitor interfaces traffic** command. [PR/846059: This issue has been resolved.]
- When you boot up an EX2200 or EX3300 switch with Junos OS Release 12.2R1 or later, the message **?dog: ERROR - reset of uninitialized watchdog** appears. The message appears even if you reboot the switch by using the proper reboot procedure. The error does not cause a system reset; thus, you can ignore this message. [PR/847469: This issue has been resolved.]
- On EX3200 and EX4200 switches, high traffic on management Ethernet (me0) interfaces might affect switch control and management plane functions. [PR/876110: This issue has been resolved.]
- On a device that is in configuration private mode, when you attempt to deactivate a previously defined VLAN members list and then commit the change, the mgd process creates a core file. [PR/855990: This issue has been resolved.]

Layer 2 and Layer 3 Protocols

- If you have configured PIM nonstop active routing (NSR), a core file might be created on an upstream router because of high churn in unicast routes or a continuous clearing of PIM join-distribution in the downstream router. To prevent this possibility, disable NSR for PIM. [PR/707900: This issue has been resolved.]
- On a device that is running Protocol Independent Multicast (PIM) and with nonstop active routing (NSR) enabled on the device, if a PIM corresponding interface flaps continuously, a PIM thread might attempt to free a pointer that has already been freed. This attempt causes the routing protocol process (rpd) to crash and create a core file. [PR/801104: This issue has been resolved.]
- If an invalid PIM-SSM multicast group is configured on the routing device, then when you issue the **commit** or **commit check** command, a routing protocol process (rpd) core file is created. There is no traffic impact because the main rpd process spawns another rpd process to parse the corresponding configuration changes, and the new rpd process crashes and creates a core file. When this problem occurs, you might see the following messages:

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

[PR/856925: This issue has been resolved.]
- On EX2200 switches, the periodic packet management process (ppmd) might create a core file. [PR/859625: This issue has been resolved.]

Management and RMON

- When a graceful Routing Engine switchover (GRES) is executed on an EX Series Virtual Chassis, CHASSISD_SNMP_TRAP6: SNMP trap generated: Power Supply Removed traps are generated periodically for all possible members of the Virtual Chassis—that is, the power supply status is checked for the maximum number of members that the Virtual Chassis can contain, even though some of those members might not exist in the configured Virtual Chassis. [PR/842933: This issue has been resolved.]
- The sFlow monitoring technology feature is not supported on EX2200, EX2200-C, and EX3300 switches. [PR/872292: This issue has been resolved.]

Multicast

- On EX4500 switches, multicast packet fragments might be dropped. [PR/835855: This issue has been resolved.]

Power over Ethernet (PoE)

- On EX2200, EX3200, EX3300 and EX4200 switches, when PoE fails to initialize, the chassis process might cause a memory leak by repeatedly calling a file open without closing it. When this issue occurs, the chassis process, which is responsible for managing hardware inventory, might generate a core file periodically every 8-9 hours. This might cause interfaces to flap, and impact service performance. [PR/845809: This issue has been resolved]

Software Installation and Upgrade

- On an EX2200-24T-DC-4G switch model, autoinstallation is not activated during initial installation because this model is missing a configuration file.

As a workaround, on the switch, starting with the shell prompt, execute these commands:

```
root@LC:0% cp /etc/config/ex2200-24t-4g-factory.conf
/etc/config/ex2200-24t-dc-4g-factory.conf
root@LC:0% cli
root>edit
root#load factory-default
{linecard:0}[edit]
root#:set system root-authentication plain-text-password
New password:
Retype new password:
[PR/873689: This issue has been resolved.]
```

Virtual Chassis

- On EX Series Virtual Chassis, if you configure a physical interface on the master switch as a member of an interface range, associate that interface with a VLAN, and then delete the interface from the interface range, the interface is not removed from the VLAN. [PR/811773: This issue has been resolved.]
- The **request system scripts add** command does not install the AI-Scripts bundle package on all nodes of an EX8200 Virtual Chassis. [PR/832975: This issue has been resolved.]
- On EX4200 Virtual Chassis, if the MAC persistence timer is configured for 0 minutes, the system MAC base address changes when a master switchover occurs and you issue the **request chassis routing-engine master switch** command. As a workaround, configure a value in the range of 1 through 60 for the **mac-persistence-timer** statement. [PR/858330: This issue has been resolved.]
- On EX8200 Virtual Chassis, NetBIOS traffic might be dropped when it crosses the non-dedicated Virtual Chassis port (that is, fiber-optic ports configured as VCPs) connections. The NetBIOS traffic is dropped because of a conflict on the Packet Forwarding Engine of the Virtual Chassis member with the VCPs. [PR/877503: This issue has been resolved.]

Issues Resolved in Release 12.3R4

The following issues have been resolved since Junos OS Release 12.3R3. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- On an EX Series switch, when you configure LLDP-MED on a trunk interface and set that interface as a member of both a voice VLAN and another VLAN, and you then change the mode of that interface to port (access) mode, the switch might send two different voice VLAN TLVs in an LLDP advertisement, and a VoIP phone connected to that interface might randomly select a VLAN to join. Use the **monitor traffic interface interface-name** command to check this issue. [PR/884177: This issue has been resolved.]

Class of Service

- On EX4200 switches, if you configure and apply more than 32 CoS rewrite rules, the Packet Forwarding Engine manager (pfem) process creates core files continuously. [PR/893911: This issue has been resolved.]

High Availability

- On EX8200 Virtual Chassis, during an NSSU, BGP neighbors might flap during the master switchover. [PR/892219: This issue has been resolved.]
- On EX8200 Virtual Chassis, during NSSU, all interfaces, including LAGs, might go down during FRU upgrades, resulting in traffic loss. [PR/893440: This issue has been resolved.]

Infrastructure

- On EX4550 switches, high-temperature alarms are triggered not on the thresholds displayed in the output of the **show chassis temperature-thresholds** command, but on other internal thresholds. [PR/874506: This issue has been resolved.]
- On EX3200 switches, an SNMP trap for pethPsePortDetectionStatus is not sent when a VoIP phone is disconnected from a PoE port. [PR/877768: This issue has been resolved.]
- On EX2200 and EX3300 switches, storm control does not limit traffic to the set value when that traffic enters through uplink ports; instead, the traffic is limited to 10 times the set value. [PR/879798: This issue has been resolved.]
- On EX4550 switches, the log message **PFC is supported only on 10G interfaces** is generated over and over again in logs. [PR/880571: This issue has been resolved.]
- On EX2200 switches, the CPU is completely consumed by the swi7: clock and chassism processes when the Redundant Power System (RPS) is powered off but is connected to the switch. At the same time, link LEDs blink continuously. When the RPS is powered up, CPU utilization and switch function becomes normal. [PR/890194: This issue has been resolved.]
- On EX4500 switches, the TLV type 314 is sent as a notification of the DCBX state of a port. In a link flap scenario, the kernel sends a DCBX PFC state TLV to the Packet

Forwarding Engine even if there is no change in the DCBX state. Also, the kernel synchronizes this state to the backup Routing Engine. On the backup Routing Engine, this message is not processed, and the system shows an Unknown TLV type 314 error. The message in itself is harmless, but it fills up the logs unnecessarily. [PR/893802: This issue has been resolved.]

- On EX4200 switches, if you issue the **request system zeroize media** command, the system boots from the backup partition and displays the following message: **WARNING: THIS DEVICE HAS BOOTED FROM THE BACKUP JUNOS IMAGE**. If the auto-snapshot feature is not enabled, reinstall Junos OS to recover the primary copy in case it has been corrupted. [PR/894782: This issue has been resolved.]
- On an EX3300 switch, when another vendor's access point is connected to one of the EX3300 interfaces, LLDP negotiation might fail and the access point is unable to boot. The system is storing the organization-specific TLV's OUI and subtype values in the parsed TLV-to-value buffer, and due to this, the offset for reading PoE power negotiation from the buffer has been changed. As a workaround:
 1. Unplug the access point.
 2. Wait until the interface power goes to 0, and verify that the physical interface is down.
 3. Issue the **set protocol lldp interface *AP-interface-name* power-negotiation disable** CLI command and commit the command. This disables power negotiation.
 4. Connect the access point.

The access point powers on in IEEE class mode power (not negotiated power). [PR/898234: This issue has been resolved.]

- On EX Series switches, after you issue the **request system zeroize media** command, SSH access fails when the switches boot from the backup root partition. This issue does not affect the primary root partition. [PR/898268: This issue has been resolved.]

Interfaces

- On EX Series switches, if you have configured a link aggregation group (LAG) with link protection, ingress traffic does not pass through the backup port. [PR/886205: This issue has been resolved.]
- On EX4200 switches, an aggregated Ethernet interface is not supported as a match condition in a firewall filter. [PR/886476: This issue has been resolved.]
- On EX Series switches, configuration of a static LACP system ID is not supported. [PR/889318: This issue has been resolved.]
- EX4500 switches might reboot suddenly because they have accessed an invalid register value for a port; this problem might occur when you insert or remove SFPs, or exchange 10-gigabit and 1-gigabit SFPs in a specific port. [PR/891733: This issue has been resolved.]
- On EX Series switches, the **request interface revert *interface-name*** command might not work. If you issue the command on the switch, the following message appears: **error:**

the **redundancy-interface-process subsystem** is not running. [PR/892976: This issue has been resolved.]

Management and RMON

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

Virtual Chassis

- If you unplug the management cable from the master switch of an EX2200 Virtual Chassis, a remote session through the management port is lost even if the backup switch has a management cable. [PR/882135: This issue has been resolved.]

Issues Resolved in Release 12.3R5

The following issues have been resolved since Junos OS Release 12.3R4. The identifier following the description is the tracking number in our bug database.

Class of Service

- Class of service on EX2200 and EX3300 Virtual Chassis ports (VCPs) might not work properly. [PR/902224: This issue has been resolved.]

Firewall Filters

- On EX4200 switches, if you change a firewall filter term and commit or roll back the firewall filter configuration, policer counter restoration might occur. [PR/900078: This issue has been resolved.]
- On EX Series switches, when two interfaces share the same firewall filter, combining two nodes into one node, and then unbinding the filter from one bind point splits the combined nodes. If the node operation type is UNBIND or DESTROY, the operation wrongly destroys the filter associated with the other node and creates a pfem process core file. [PR/927063: This issue has been resolved.]

Hardware

- On EX2200 and EX3300 switches, for some types of SFP transceivers, the output of the **show interfaces diagnostics optics** CLI command contains an incorrect value of **0.0000 mW / - Inf dBm** for the **Receiver signal average optical power** field. [PR/909334: This issue has been resolved.]
- On EX4550-32T switches, some ports might not link up correctly. [PR/901513: This issue has been resolved.]
- On EX6200 switches, if the LCD backlight is off and you then press the **menu** or **enter** buttons on the LCD panel, the LCD is reinitialized. During this reinitialization, the switch might drop some packets. [PR/929356: This issue has been resolved.]

High Availability (HA) and Resiliency

- On EX Series Virtual Chassis, an upgrade with NSSU might cause a mismatch in the physical interface index numbers between the master and backup Packet Forwarding Engines, causing result packets to be dropped as they pass through the Virtual Chassis. [PR/882512: This issue has been resolved.]

Infrastructure

- On EX Series switches, the messages **CMLC: connection in progress for long** and **pfem: devrt_gencfg_rtsock_msg_handler Incorrect major_type 8** might be displayed, but the messages do not impact switch functionality. [PR/890633: This issue has been resolved.]
- On EX2200 switches, a primary file system corruption might not be detected and the system might not fail over to the backup partition. Some functional problems might occur. [PR/892089: This issue has been resolved.]
- On EX8200 switches, an NSSU might cause some hosts to become unreachable because the ARP index for the impacted host route is incorrectly programmed. The host route references the old ARP index and fails to update the new ARP index. [PR/894436: This issue has been resolved.]
- On EX8200 switches equipped with EX8200-40XS line cards, when a port on a 40XS line card connects to another device and the port is then disabled, the carrier transition count might increase continuously, which might cause high CPU utilization. The carrier transition count is displayed in the output of the **show interfaces interface-name extensive** command. [PR/898082: This issue has been resolved.]
- On EX4550 switches running Junos OS Release 12.2R5 or Release 12.3R3, commit operations might cause a spike in CPU utilization, resulting in a timeout of LACP, BFD, and other protocols. [PR/898097: This issue has been resolved.]
- On EX2200 switches, the system log (syslog) messages might show IP addresses in reverse. For example, an ICMP packet from 10.0.1.114 to 10.0.0.7 might be shown in the log as **PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 114.1.0.10 7.0.0.10 0 0 (1 packets)** instead of **PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 10.0.1.114 10.0.0.7 0 0 (1 packets)**. [PR/898175: This issue has been resolved.]
- On EX Series switches, if the eventd process is not restarted gracefully, the process might crash or exit and the **SYSTEM_ABNORMAL_SHUTDOWN: System abnormally shut down** message might be generated. [PR/901924: This issue has been resolved.]
- On an EX6200 switch, if you disconnect the master Routing Engine (RE0) and reconnect it, the backup Routing Engine (RE1) becomes the master, and then when the original RE0 is rebooted, it becomes the backup; however, that new backup does not appear in **show chassis routing-engine** command output on RE0 (the new master). [PR/919242: This issue has been resolved.]
- On EX Series switches that are running Junos OS Release 12.1 and later releases, if you install AI-Scripts package releases earlier than 3.6R4 and 3.7R3 and then execute a reboot/commit sequence, the switch might generate a FIPS core file and might crash. [PR/920478: This issue has been resolved.]

- On EX Series switches with DHCP snooping enabled, the DHCP reply packets without any DHCP options (BOOTP reply packets) might be dropped. [PR/925506: This issue has been resolved.]
- Polling the OID mib-2.17.7.1.4.3.1.5...: dot1qPortVlan on an EX9200 switch might cause a memory leak on the l2ald process, and the process might create core files. [PR/935981: This issue has been resolved.]

Interfaces

- On EX6200 switches, an interface might not be able to come up after the interface flaps due to a discrepancy on the physical channel. [PR/876512: This issue has been resolved.]
- On EX9200 switches, Layer 3 unicast traffic losses might be seen for a few seconds during graceful Routing Engine switchover (GRES) for host prefixes learned over MC-LAG interfaces. [PR/880268: This issue has been resolved.]
- On an EX3300 switch, when another vendor's AP is connected to one of the EX3300 interfaces, LLDP negotiation might fail and the AP is unable to boot. The system is storing the organization-specific TLV's OUI and subtype values in the parsed TLV-to-value buffer, and due to this, the offset for reading PoE power negotiation from the buffer has been changed. As a workaround:
 1. Unplug the AP.
 2. Wait until the interface power goes to 0, and verify that the physical interface is down.
 3. Issue the **set protocol lldp interface power-negotiation disable** CLI command and commit the command. This will disable power negotiation.
 4. Connect the AP.

The AP will power on in IEEE class mode power (not negotiated power). [PR/898234: This issue has been resolved.]

Layer 2 Protocols

- On EX8200 switches or EX8200 Virtual Chassis with nonstop bridging (NSB) enabled, continuously adding and deleting VLAN members along with continuously creating and deleting VLANs might cause the Ethernet switching process (eswd) to leak memory and create a core file. [PR/878016: This issue has been resolved.]

Multicast

- On EX Series switches, the multicast route cache timer might not be cleared in some situations. As a workaround, issue the **show multicast route** command several times. [PR/937695: This issue has been resolved.]

Network Management and Monitoring

- On an EX9200 switch, if you configure port mirroring, the feature might not work and the switch might not be able to mirror Layer 2 and Layer 3 traffic. [PR/920213: This issue has been resolved.]

Software Installation and Upgrade

- On EX8200 Virtual Chassis, the licensing policy specifies that you install the Advanced Feature Licenses (AFLs) on the master and backup XRE200 External Routing Engines. In Junos OS 12.3 releases, a warning message might appear at commit indicating that the AFLs have not been installed on the Routing Engines on the EX8200 member switches even though the AFLs have been installed on the external Routing Engines. [PR/919605: This issue has been resolved.]

Virtual Chassis

- On EX Series Virtual Chassis, if you convert a physically down Virtual Chassis port (VCP) to a network port, broadcast and multicast traffic might be dropped on the VCP interface. [PR/905185: This issue has been resolved.]

Issues Resolved in Release 12.3R6

The following issues have been resolved since Junos OS Release 12.3R5. The identifier following the description is the tracking number in our bug database.

Class of Service

- On EX4200-48PX switch models, configuring the traffic shaping rate on an interface using the **set class-of-service interfaces *interface-name* shaping-rate** command might return the error message **shaping rate not allowed on interface *interface-name***. [PR/944172: This issue has been resolved.]

Hardware

- On EX Series switches, an SFP might stop working unexpectedly with i2c errors and the switch might not recognize the SFP in its existing port. [PR/939041: This issue has been resolved.]

High Availability

- On EX Series Virtual Chassis with a link aggregation group (LAG) interface configured, if one member link of the LAG is on the backup Routing Engine, traffic loss on the LAG interface might be observed during an NSSU. Traffic resumes after the graceful Routing Engine switchover (GRES) occurs in the last state of the NSSU. [PR/916352: This issue has been resolved.]

Infrastructure

- EX3200 and EX4200 switches might stop forwarding traffic when the traffic exits from interfaces. [PR/856655: This issue has been resolved.]
- On EX2200, EX2200-C, and EX3300 switches, if you configure more than one domain-search attribute under the **[edit system services dhcp pool]** hierarchy level, the dhcpd process might create a core file. [PR/900108: This issue has been resolved.]
- On EX4550 Virtual Chassis, SFPs might not be detected, causing continuous EEPROM read failed errors. [PR/911306: This issue has been resolved.]
- In EX4200 Virtual Chassis, a member of the Virtual Chassis might reboot and create a pfem core file. [PR/912889: This issue has been resolved.]
- On EX Series switches except EX9200, the network interfaces information about **Receiver signal average optical power** that is displayed in command output might be incorrect when you reconfigure a fiber network interface to a Virtual Chassis port (VCP). You can see this information display by issuing the show virtual-chassis vc-port diagnostics optics command. [PR/916444: This issue has been resolved.]
- On EX Series switches, when an RSTP-enabled interface that becomes active is a member of a VLAN that has a Layer 3 interface, if this interface does not receive any BPDUs, gratuitous ARP is not sent out. [PR/920197: This issue has been resolved.]
- On EX Series switches, when a packet is received that matches a firewall filter term with action **syslog**, configured to send the log to a remote syslog server, the switch might not send logs to the syslog server. [PR/926891: This issue has been resolved.]
- On EX Series switches with a router firewall filter configured, the filter might not work if it is applied to an IPv6 VRRP-enabled interface; also, features corresponding to the filter, such as policers, do not work. [PR/926901: This issue has been resolved.]

- On an EX Series switch with TACACS+ authentication and accounting enabled, when the TACACS+ server is in an unresponsive state and sends an erroneous response with an End of File (EOF) that indicates that no data can be read from a data source, this circumstance causes the client to fail to decrement the sequence number that it manages locally. During that time, any TACACS+ authentication might fail. [PR/929273: This issue has been resolved.]
- On EX6200 switches running Junos OS Release 11.3R1 or later, if the LCD backlight is off and then you press the buttons on the LCD panel, the LCD is reinitialized. During this reinitialization, the switch might drop some packets. [PR/929356: This issue has been resolved.]
- On an EX9200 switch configured for DHCP relay, if an IRB interface walks through a Layer 2 trunk interface and the corresponding DHCP relay is configured in a routing instance, and if you deactivate or activate (or delete or add) a hierarchy that contains a DHCP relay-related configuration, DHCP relay might not work as expected. As a workaround, restart DHCP services after you make any changes to DHCP configurations. [PR/935155: This issue has been resolved.]
- On EX3200 and EX4200 switches, if multicast traffic is bursty or cyclical with no traffic for continuous 30-second periods, then the multicast keepalive timer might age out, thus deleting that particular route and causing multicast traffic loss.

As a workaround, use one of the following options:

- Set a large timeout value for multicast forwarding cache entries using the **set routing-options multicast forwarding-cache timeout** command.
- Using a script, issue the **show multicast route** command continuously every 25 seconds.

[PR/937695: This issue has been resolved.]

- On EX9200 switches that are configured for DHCP relay, if you deactivate or activate an IRB interface, DHCP relay for that interface might stop working and might drop DHCP packets. [PR/937996: This issue has been resolved.]
- On EX Series switches with dual Routing Engines, with the switch configured with VRRP, if VRRP is configured under an interface subnet, the kernel might create a core file on the backup Routing Engine because states are out of sync on the master and backup Routing Engines. If this issue occurs on an EX Series Virtual Chassis, it will cause a service impact. [PR/939418: This issue has been resolved.]
- On EX4500 or EX4550 switches, if you apply a firewall filter to a loopback interface, transit packets that match Precise Time Protocol (PTP) errata might be dropped. [PR/949945: This issue has been resolved.]
- On an EX Series Virtual Chassis that is configured for DHCP services and configured with a DHCP server, when a client sends DHCP INFORM packets and then the same client sends the DHCP RELEASE packet, an IP address conflict might result because the same IP address has been assigned to two clients. As a workaround:
 - 1. Clear the binding table:
`user@switch> clear system services dhcp binding`

- 2. Restart the DHCP service:
`user@switch> restart dhcp`

[PR/953586: This issue has been resolved.]

- When the SNMP mib2d process polls system statistics from the kernel, the kernel might cause a memory leak (mbuf leak), which in turn might cause packets such as ARP packets to be dropped at the kernel. [PR/953664: This issue has been resolved.]

Interfaces

- On EX9200 switches that are equipped with EX9200-32XS or EX9200-2C-8XS line cards, 10-gigabit ports on these cards might stay offline after a link flaps or after an SFP+ is inserted. [PR/905589: This issue has been resolved.]
- On EX9200 switches, an inter-IRB route might not work if Q-in-Q tunneling is enabled, because the TPID (0x9100) is not set on egress dual-tagged packets, and other devices that receive these untagged packets might drop them. [PR/942124: This issue has been resolved.]
- On an EX9200 switch that is configured for DHCP relay, with the switch acting as the DHCP relay agent, the switch might not be able to relay broadcast DHCP inform packets, which are used by the client to get more information from the DHCP server. [PR/946038: This issue has been resolved.]
- On an EX Series switch, if you remove an SFP+ and then add it back or reboot the switch, and the corresponding disabled 10-gigabit interface is a member of a LAG, the link on that port might be activated. [PR/947683: This issue has been resolved.]

Layer 2 Features

- On EX Series switches, the following log message might appear after every commit operation for a configuration change: **Aug 20 12:06:35.224 2013 UKLDNHASTST5B01 eswd[1309]: Bridge Address: add ffffffb0:fffffc6:ffff9a:69:ffff9d:ffff81 Aug 20 12:36:35.423 2013 UKLDNHASTST5B01 eswd[1309]: Bridge Address: add ffffffb0:fffffc6:ffff9a:69:ffff9d:ffff81**. The MAC address is that of the chassis. This is an informational message and does not impact any service. [PR/916522: This issue has been resolved.]
- On EX Series switches that are configured with Ethernet Ring Protection Switching (ERPS), if the switch is configured as the RPS owner and is in a topology with other vendors' switches that are running ERPSv2 (ERPS version 2), when an indirect link failure occurs on the Ethernet ring, the ring protection link (RPL) end interface might not be able to get into the forwarding state. [PR/944831: This issue has been resolved.]
- On EX Series switches with RSTP enabled at the global level, when a VoIP-enabled interface is also enabled with VSTP, if you deactivate VSTP on this interface, the interface might stop forwarding traffic. [PR/952855: This issue has been resolved.]
- On EX Series switches (except EX9200) with VSTP configured, if a switch has two access ports looped back that connect to another switch over a trunk port, this might cause an incorrect STP state (BLK or DESG) in the same VLAN on the trunk port. When this issue occurs, service is impacted. [PR/930807: This issue has been resolved.]

Network Management and Monitoring

- An EX Series switch might send sFlow monitoring technology packets with source port 0. [PR/936565: This issue has been resolved.]

Port Security

- On EX Series switches with VoIP configured, if the switch receives an IP source guard (IPSG) or dynamic ARP inspection (DAI) route-delete message on an interface, voice VLAN traffic on these interfaces might be dropped. [PR/937992: This issue has been resolved.]

Software Installation and Upgrade

- On EX8200 Virtual Chassis, an NSSU from Junos OS Release 11.4R9 to Release 12.3R4 brings down LAGs and other interfaces during the member-switch upgrades, and thus large traffic losses occur. [PR/914048: This issue has been resolved.]

Virtual Chassis

- In a protocol-mastership transition, the ksyncd process might fail to clean up the kernel VPLS routing tables due to dependencies such as VLANs not being cleaned up first, leaving the tables in an inconsistent state. [PR/927214: This issue has been resolved.]

Issues Resolved in Release 12.3R7

The following issues have been resolved since Junos OS Release 12.3R6. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On an EX Series switch that has both 802.1X authentication (dot1x) and a dynamic firewall filter enabled, when the server-timeout value is set to a short time (for example, 3 seconds), if many clients try to authenticate at the same time, a **delay success authentication success** message might be received on the switch due to a RADIUS server timeout, the firewall filter might corrupt the interfaces on which the authentication attempts were made, and the subsequent client authentications might fail due to the stale firewall filter. As a workaround, configure a server-timeout value that is greater than 30 seconds. [PR/967922: This issue has been resolved.]

Class of Service (CoS)

- On an EX Series switch, when you configure both inet and inet6 on an interface and both dscp and dscp-ipv6 classifiers are configured on the switch, you might see this system log message: **Jan 22 15:56:54.932 2014 EX4200 cosd[1306]: Classifier CLASSIFIER is not supported on ge-0/0/1.0 interface for inet6 family. Jan 22 15:56:54.932 2014 EX4200 cosd[1306]: Classifier CLASSIFIER6 is not supported on ge-0/0/1.0 interface for inet family.** This message has no operational effect on the switch, as this function is supported. You can ignore the message. [PR/956708: This issue has been resolved.]

High Availability

- On EX Series switches with dual Routing Engines that are configured with IS-IS, if **traceoptions** is configured under **[edit protocols isis]**, a Routing Engine switchover might cause IS-IS to flap or an rpd core file to be generated. [PR/954433: This issue has been resolved.]

Infrastructure

- On an EX Series Virtual Chassis that has a virtual management Ethernet (vme) interface, when the Virtual Chassis is initially formed, you might be unable to access the Virtual Chassis through the vme interface if the management cable is connected to a Virtual Chassis member other than the master. As a workaround, reboot the Virtual Chassis. [PR/934867: This issue has been resolved.]
- On EX8200 switches with Multicast Listener Discovery (MLD) snooping enabled, the number of MLD snooping entries might grow in the kernel, increasing the number of multicast groups to such an extent that eventually the forwarding table is filled, causing a service impact. [PR/940623: This issue has been resolved.]
- On EX Series switches with 802.1X enabled, when the RADIUS server is unreachable, 802.1X-enabled interfaces might stop forwarding traffic after you have deactivated the 802.1X protocol by deactivating **[edit protocols dot1x]**. As a workaround, deactivate 802.1X. [PR/947882: This issue has been resolved.]
- On EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6200 switches, DHCPv6 unicast packets might be dropped after you enable a firewall filter on the loopback interface (lo0.0) to protect the Routing Engine. As a workaround, add a term to accept DHCPv6 packets in the loopback filter. [PR/960687: This issue has been resolved.]

- On EX9200 switches acting as DHCP relays, broadcast BOOTP reply messages that are received might be dropped. [PR/961520: This issue has been resolved.]
- On EX Series switches, starting in Junos OS Release 12.3, the file `/var/log/wtmp` is not rotated once a month or every 10 MB. As a workaround, manually rotate `/var/log/wtmp` by issuing the command `set system syslog file wtmp archive files 10 size 1M`. [PR/964118: This issue has been resolved.]
- On EX Series switches except for EX9200, when an IPv6 firewall filter that has Layer 4 match conditions (for example, `tcp-established`) configured, is applied to routed VLAN interfaces (RVIs) in the egress direction, these match conditions might not work as expected. [PR/972405: This issue has been resolved.]
- On EX9200 switches, in a DHCP relay scenario, in cases where a binding entry already exists for a client, if the client sends a DHCP discover packet, the device might not relay DHCP offers from any server other than the server used to establish the existing binding. [PR/974963: This issue has been resolved.]

Interfaces and Chassis

- On EX9200 switches with DHCP relay configured and with permanent ARP entries for relay clients installed, if a client is reachable through a different preferred path due to STP topology changes, MC-LAG changes, and so on, the forwarding state is not refreshed, and traffic might be dropped until the relay binding is cleared. As a workaround, issue the following configuration command to suppress installation of destination routes (DHCPv4 only): `set forwarding-options dhcp-relay route-suppression destination`. [PR/961479: This issue has been resolved.]
- On EX Series switches with scaled ARP entries (for example, 48K entries), in a normal state, an ARP entry's current time is less than the expiry time. Some events might cause the current time to be greater than the expiry time, which prevents this ARP entry from being flushed, which causes a connectivity issue. One possible trigger event is an ICL flap in an MC-LAG scenario. [PR/963588: This issue has been resolved.]
- On EX9200 switches, the configuration statement `mcae-mac-flush` is not available in the CLI; it is missing from the `[edit vlans]` hierarchy level. [PR/984393: This issue has been resolved.]

Layer 3 Features

- On EX Series switches, EBGP neighborship might go down and an `rpd` core file might be created. [PR/960829: This issue has been resolved.]
- On EX8200 switches, when the MTU value on a Layer 3 interface is configured as 1518 and you execute the `clear pim join` command or reboot the switch, multicast traffic might be dropped when packet sizes are greater than 1500, because the multicast route might eventually point to a smaller MTU value and packets cannot pass, even though the packet size is smaller than the MTU-configured value. As a workaround, configure all the Layer 3 interface MTUs to 9192. [PR/966704: This issue has been resolved.]

Network Management and Monitoring

- On EX Series switches, an OAM CFM interface might not recover automatically if the action in **[edit oam ethernet connectivity-fault-management action-profile link-down action action]** is **interface-down**. As a workaround, do not use link-down in the action profile. [PR/948082: This issue has been resolved.]

Platform and Infrastructure

- On an EX9200 switch working as a DHCP server, when you delete an IRB interface or change the VLAN ID of a VLAN corresponding with an IRB interface, the DHCP process (jdhcpd) might create a core file after commit, because a stale interface entry in the jdhcpd database has been accessed. [PR/979565: This issue has been resolved.]

Port Security

- On EX Series switches that are configured for voice over IP (VoIP), if dynamic ARP inspection (DAI) is enabled with the voice VLAN, ARP packets might get dropped for this VLAN. [PR/946502: This issue has been resolved.]

Routing Protocols and Firewall Filters

- On EX Series switches that are configured for filter-based forwarding (FBF), if you configure a maximum transmission unit (MTU) on an egress interface, packets that are larger than the configured MTU size might be dropped. [PR/922581: This issue has been resolved.]
- On EX9200 switches with IGMP snooping enabled on an IRB interface, some transit TCP packets might be treated as IGMP packets, causing packets to be dropped. As a workaround, disable IGMP snooping. [PR/979671: This issue has been resolved.]

Spanning-Tree Protocols

- On EX Series switches except EX2200 and EX9200, when Rapid Spanning Tree Protocol (RSTP) and VLAN Spanning Tree Protocol (VSTP) are enabled at the same time, an RSTP topology change might delete MAC entries learned on VLANs managed by VSTP. [PR/900600: This issue has been resolved.]

Virtual Chassis

- On EX4550 Virtual Chassis and EX4550 mixed mode Virtual Chassis, the chassis manager process (chassism) might crash when the **request support information** is executed. [PR/977011: This issue has been resolved.]
- On EX Series switches except EX9200, the Ethernet switching process (eswd) might crash when receiving Link Layer Discovery Protocol (LLDP) packets on a member interface of a LAG. This is because the LAG fails to handle LLDP packets. [PR/983330: This issue has been resolved.]

Issues Resolved in Release 12.3R8

The following issues have been resolved since Junos OS Release 12.3R7. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On EX Series switches configured for accounting based on 802.1X RADIUS, if the RADIUS server is enabled with the **User-Name** attribute and a new username is used to send account information, the switches might ignore this attribute and not send accounting information with the authentication username. [PR/950562: This issue has been resolved.]
- On EX Series switches with 802.1X authentication enabled, if you associate an 802.1X-enabled interface in single-secure mode with a VLAN, when a client is authenticated on that VLAN and is later authenticated on a dynamic VLAN (a guest VLAN or a VLAN assigned by a RADIUS server), the client might still be associated with the interface-associated VLAN and receive broadcast and multicast traffic of the VLAN associated with the interface. [PR/955141: This issue has been resolved.]

Class of Service

- On EX2200 or EX3300 switches that are running Junos OS Release 12.3R5 or later releases, CoS settings might remain active on interfaces after you have removed the CoS-related configuration. [PR/992075: This issue has been resolved.]

Infrastructure

- On EX9200 switches, in a BOOTP relay agent scenario, DHCPACK messages that are sent in response to DHCPINFORM messages might not be forwarded to the DHCP client if these ACK messages are sent from a DHCP server other than the DHCP server that is in the DHCP relay agent's binding table. [PR/994735: This issue has been resolved.]
- On EX8200 switches with link aggregation groups (LAGs) configured, high CPU utilization might be observed on line cards (FPCs) after you change the configuration of the LAGs. [PR/976781: This issue has been resolved.]
- On EX Series switches with Protocol Independent Multicast (PIM) configured, when the upstream interface on a rendezvous point (RP) changes between a Layer 3 interface and the PIM de-encapsulation interface for the multicast route, the earlier route entry might be deleted twice, which causes a loss of multicast traffic on the RP. [PR/982883: This issue has been resolved.]
- On EX4200 switches, the system date and time might change after you reboot the switch. [PR/985819: This issue has been resolved.]
- On EX Series switches, the software forwarding infrastructure process (sfid) might create a core file while processing a packet for which the TTL has expired, because the packet pointer is freed twice. [PR/988640: This issue has been resolved.]
- After you reboot an EX Series switch, the interface initial sequence on an aggregated Ethernet member interface might not proceed properly and the interface might not come up. [PR/989300: This issue has been resolved.]
- On EX Series switches except EX9200 switches, ARP reply packets might get dropped when the switch receives a large amount of reverse-path forwarding (RPF) multicast failure packets (for example, 300 pps). As a workaround, create a static ARP entry for the next-hop device. [PR/1007438: This issue has been resolved.]

- On an EX2200 or EX3300 switch that is backed up by a Redundant Power System, if the configuration is committed on the switch side, the RPS might stop providing backup power, and the switch might be powered off. [PR/1011821: This issue has been resolved.]
- On EX Series switches that are configured for filter-based forwarding (FBF) and are running Junos OS 12.3R7 or a later release, configuring the accept action for a firewall filter to forward matched traffic to a specific routing instance might not work as expected, and all traffic is dropped. [PR/1014645: This issue has been resolved.]

Interfaces and Chassis

- On an EX4500 or EX4550 switch with an MPLS circuit cross-connect (CCC) interface configured, there might be high CPU utilization by the software forwarding infrastructure process (sfid) while large amounts of IPv6 neighbor solicitation packets (for example, 1000 pps) are received on the MPLS CCC interface. [PR/961807: This issue has been resolved.]
- On EX Series switches with Link Aggregation Control Protocol (LACP) enabled on a LAG interface, after the master Routing Engine is rebooted, if the first LACP packet is dropped during switchover, the LACP state might get stuck in the same state for a long time (about 10s), which causes the LAG interface to flap and traffic to drop on the interface. [PR/976213: This issue has been resolved.]
- On EX8200 Virtual Chassis running Junos OS Release 12.3R6 or later releases, multicast and broadcast traffic might be dropped on Virtual Chassis ports (VCPs) in the following scenarios:
 - When a few link aggregation group (LAG) member interfaces go up and down continuously.
 - When LAG member interfaces go up and down simultaneously.
 - After you delete a VCP that corresponds to a LAG member interface.[PR/993369: This issue has been resolved.]
- On EX8200 switches with a generic routing encapsulation (GRE) tunnel configured, packets might be dropped permanently on GRE interfaces when you create a logical GRE interface. The existing GRE interfaces are not affected by the addition of more GRE interfaces. [PR/995990: This issue has been resolved.]
- On EX9200 switches, if you configure an interface with the **mac-rewrite** statement, the Layer 2 address learning process (l2ald) might create a core file. [PR/997978: This issue has been resolved.]
- If you configure autoinstallation on an EX3300 switch, you might see the error message **Unaligned memory access error**. [PR/999982: This issue has been resolved.]
- On EX4200 switches, some interface diagnostic optical values might be inconsistent between Junos OS Releases 12.3R6.6 and 12.3R7.7. [PR/1007055: This issue has been resolved.]

Layer 2 Features

- On EX Series switches, on the backup Routing Engine, in the Ethernet-switching process (eswd), there might be a scenario that causes the backup Routing Engine to miss the

sync update from kernel and get into an inconsistent state with respect to flood next-hop. An eswd core file is created on the backup Routing Engine. No outage occurs as the core file is on the backup Routing Engine. [PR/936567: This issue has been resolved.]

- On EX Series switches, after a switch reboot, a Q-in-Q tunneling interface might not function as expected. The problem occurs when the interface is a member of a PVLAN with mapping set to **swap** and is also a member of a non-private VLAN. The PVID of the access interface does not get set when the PVLAN is configured before the non-private VLAN. The problem does not occur when the non-private VLAN is configured before the PVLAN. [PR/937927: This issue has been resolved.]
- On EX Series switches with L2PT and Q-in-Q tunneling enabled, some of the MAC addresses might not be learned. The problem occurs when there is a high volume of L2PT packets. As a workaround, restart the eswd and sfid processes. [PR/996638: This issue has been resolved.]

MPLS

- On EX Series switches running Junos OS Release 12.1R1 or later releases, the MPLS TTL might change to 1 on a transit MPLS switch, causing packets to be dropped on the egress MPLS tunnel due to TTL expiration. As a workaround, enable the **no-decrement-ttl** statement in the **[edit protocols mpls]** hierarchy level. [PR/1005436: This issue has been resolved.]

Port Security

- On EX Series switches except EX9200, the port security **allowed-mac** feature might not work as expected. When this issue occurs, the traffic from an unauthorized host is unimpeded. [PR/1001124: This issue has been resolved.]

Routing Policy and Firewall Filters

- On EX3300 switches, when you use a wildcard mask in firewall filters, the error message **Unaligned memory access by pid 87736 [dfwc] at 1000f5 PC[4adec]** might be displayed at commit. [PR/996083: This issue has been resolved.]

Spanning-Tree Protocols

- On EX4550 Virtual Chassis switches with xSTP (RSTP, VSTP or MSTP) enabled, multiple xSTP-enabled interfaces might go into the STP Disabled state on the Packet Forwarding Engine because of the overlap of STP identifiers. Traffic is dropped on these problematic interfaces. [PR/980551: This issue has been resolved.]

Issues Resolved in Release 12.3R9

The following issues have been resolved since Junos OS Release 12.3R8. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On EX Series switches with 802.1X enabled, if the voice VLAN is authenticated using MAC-based authentication and the data VLAN is authenticated using 802.1X-based

authentication, traffic loss might occur on the voice VLAN during re-authentication. [PR/1011985: This issue has been resolved.]

- On an EX Series Virtual Chassis with 802.1X enabled, if the Software Forwarding Infrastructure process (sfid) generates a core file, it causes the FPC to disconnect from the Routing Engine. The 802.1X process (dot1xd) receives a delete message for the physical interface (ifd) from the kernel, but does not clear the sessions associated with the interface. When those sessions expire and the corresponding timer attempts to access any interface data, then the dot1xd generates a core file. [PR/1016027: This issue has been resolved.]
- On EX Series switches, captive portal authentication is used to redirect Web browser requests to a login page. After the client is successfully authenticated, there might be a delay of 1-3 minutes before captive portal redirects the browser to the login page, and at times, the redirection might fail. [PR1026305: This issue has been resolved.]

Class of Service

- On EX4500 switches, if CoS is configured on an interface and MPLS is enabled on the same interface, the configured CoS mapping might disappear. [PR/1034599: This issue has been resolved.]

Dynamic Host Configuration Protocol

- On EX Series switches, DHCP option 125 cannot be configured for use as the byte-stream option. [PR/895055: This issue has been resolved.]

Firewall Filters

- On EX Series switches, a firewall filter configured on a loopback interface (lo0) containing a match condition for an IPv6 destination address with a prefix length longer than 8 leading bits might not work as expected. [PR/962009: This issue has been resolved.]

High Availability

- On EX4550 Virtual Chassis with LAG interfaces configured, the LAG interfaces might go down for approximately 30 seconds during an NSSU if two switches in the linecard role have consecutive member IDs and are members of a LAG interface with only two child members. LAG downtime will increase with the number of SFPs. [PR/1005024: This issue has been resolved.]
- On EX Series switches, after you change the VRRP advertise interval, the VRRP **Master is Dead** timer value might still be based on the previous advertise interval. As a workaround, reboot the switch. [PR/1017319: This issue has been resolved.]

Infrastructure

- On EX4500 or EX4550 switches, the software forwarding infrastructure process (sfid) might continuously create core files, causing interruptions in traffic, because packets are erroneously freed twice. A possible trigger is the handling of Layer 2 protocol tunneling packets. [PR/941482: This issue has been resolved.]
- On EX4500 switches with an uplink module installed, if the uplink module is removed and then installed in less than 10 seconds, the chassis manager (chassism) might create a core file. [PR/941499: This issue has been resolved.]
- On EX Series switches, if you use **apply-groups** 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 does not affect the configuration; if the configuration validation passes, **apply-groups** is expanded correctly only on interfaces on which **apply-groups** is configured. [PR/967233: This issue has been resolved.]
- On EX8200 switches, a kernel memory leak might occur and core files might be created when a next-hop device is changed (for example, when MAC or ARP entries from Layer 3 interfaces that span multiple Packet Forwarding Engines are flushed). You can view the log files for the memory leak by issuing the **show system virtual-memory | match temp** command multiple times. [PR/977285: This issue has been resolved.]
- On EX2200 switches, if CoS is configured on the VCP and network ports, the log message **devrt_ifd getting linkstate failed for dev 1 port 0** might appear continuously. These messages have no service impact. [PR/988063: This issue has been resolved.]
- On EX Series switches, **GENCFG: op 8 (COS BLOB) failed; err** error messages might appear in the log files. These messages have no service impact. [PR/997946: This issue has been resolved.]
- On EX Series Virtual Chassis, if the master switch is rebooted or halted while traffic is flowing through one of its interfaces, the FDB entry remains in an incomplete or discard state for about 30 seconds. During that time, traffic that uses the FDB entry is lost. [PR/1007672: This issue has been resolved.]
- On EX2200-C switches, the **log-out-on-disconnect** command might not work, causing the previous console session users to be seen on the switch. [PR/1012964: This issue has been resolved.]
- If the **disable-logging** option is the only configured option under the **[edit system ddos-protection global]** hierarchy level, and if this option is deleted, the kernel might generate a core file. [PR/1014219: This issue has been resolved.]
- On EX Series switches, hosts might lose connectivity to switches when the ARP entry ages out because of a programming error in the ARP entry for the Packet Forwarding Engine hardware. [PR/1025082: This issue has been resolved.]
- On EX Series switches, the ptopoConnLastVerify MIB returns a wrong value. [PR/1049860: This issue has been resolved.]
- On EX Series switches, the ptopoConnRemotePort MIB returns a wrong value. [PR/1052129: This issue has been resolved.]

Interfaces and Chassis

- On EX9200 switches that are configured in a multicast scenario with PIM enabled, an (S,G) discard route might stop programming if the switch receives resolve requests from an incorrect reverse-path-forwarding (RPF) interface. After this issue occurs, the (S,G) state might not be updated when the switch receives multicast traffic from the correct RPF interfaces, and multicast traffic might be dropped. [PR/1011098: This issue has been resolved.]
- On EX9200 switches, in an MC-LAG scenario, a MAC address might incorrectly point to an inter-chassis control link (ICL) after a MAC move from a single-home LAG to the MC-LAG. [PR/1034347: This issue has been resolved.]

J-Web Interface

- On EX Series switches, the J-Web service might become slow or unresponsive. [PR/1017811: This issue has been resolved.]

Layer 2 Features

- On EX Series switches, an Ethernet switching process (eswd) memory leak might occur if the following conditions are met:
 - If a VLAN has the VLAN index 0, and the VLAN is deleted, but the memory is not freed accordingly.
 - In a Multiple VLAN Registration Protocol (MVRP) scenario, when a VLAN map entry is deleted, but the memory is not freed accordingly.

[PR/956754: This issue has been resolved.]

- On EX Series switches running Junos OS Release 12.1R1 or later with Layer 2 protocol tunneling (L2PT) configured, if the switch receives a burst of more than 10 L2PT packets, the excessive L2PT packets might be dropped. [PR/1008983: This issue has been resolved.]
- On EX Series switches with private VLAN (PVLAN) and DHCP snooping configured, if the interface configured with PVLAN flaps, the Ethernet switching process (eswd) might stop responding to management requests, and high eswd and Software Forwarding Infrastructure process (sfid) utilization might be observed. [PR/1022312: This issue has been resolved.]

Network Management and Monitoring

- On EX Series switches, the connectivity fault management process (cfmd) might generate a core file. This happens when the Ethernet switching process (eswd) sends information to the cfmd to update its VLAN database, but because of a timing issue, the VLAN ID that the cfmd has is no longer current. [PR/961662: This issue has been resolved.]
- On EX Series Virtual Chassis, if one member of the Virtual Chassis is rebooted or if there is a switch failover, the connectivity fault management process (cfmd) might continue to send next-hop add requests to the kernel, which results in traffic being dropped

when the next-hop space index is exhausted. [PR/1016587: This issue has been resolved].

Port Security

- On EX Series switches, there might be traffic loss under any of the following conditions:
 - IP source guard is enabled and then disabled.
 - An interface belonging to a VLAN that has IP source guard enabled is changed to another VLAN that does not have IP source guard enabled.
 - 802.1X authentication is enabled or disabled on an interface belonging to a VLAN that has IP source guard enabled.

[PR/1011279: This issue has been resolved.]

Routing Protocols

- On EX9200 Virtual Chassis, the chassisd on the protocol master RE and the protocol backup Routing Engine connect to the main SNMP process (snmpd) on the protocol master using the following methods:
 - Chassisd on the protocol master Routing Engine connects, using a local socket because snmpd is running locally.
 - Chassisd on the protocol backup Routing Engine connects, using a TNP socket because snmpd is not local.

As a result, all processes that run on the protocol master (other than chassisd) attempt to connect to snmpd by using the TNP socket instead of a local socket. The snmpd does not accept these connections. [PR/986009: This issue has been resolved.]

- On EX Series switches with dual Routing Engines, in a multicast scenario, the routing protocol process (rpd) might generate a core file when the backup Routing Engine processes a multicast **resolve** request to add a multicast route entry that is already present. [PR/1018896: This issue has been resolved.]

Virtual Chassis

- On EX Series Virtual Chassis, if VLAN pruning is enabled on a VLAN, traffic on that VLAN might be dropped on the Virtual Chassis port (VCP) if the link is changed from trunk to access mode and then back to trunk mode. [PR/1012049: This issue has been resolved.]
- In an EX8200 Virtual Chassis with three members and link aggregation group interfaces configured, traffic might be dropped on the LAG interfaces after one member of the Virtual Chassis is rebooted. [PR/1016698: This issue has been resolved.]
- In an EX8200 Virtual Chassis with tunnel interfaces configured, for example, for GRE or a LAG, traffic might be dropped on tunnel interfaces after an upgrade using NSSU. [PR/1028549: This issue has been resolved.]

Issues Resolved in Release 12.3R10

The following issues have been resolved since Junos OS Release 12.3R9. The identifier following the description is the tracking number in our bug database.

Authentication and Access Control

- On EX4500 and EX8200 switches with LLDP enabled and Edge Virtual Bridging (EVB) configured, when a switch is connected to a virtual machine (VM) server using Virtual Ethernet Port Aggregator (VEPA) technology, the EVB TLV in LLDP packets might be sent to the incorrect multicast MAC address of 01:80:c2:00:00:0e instead of the correct address 01:80:c2:00:00:00. [PR/1022279: This issue has been resolved.]
- On EX Series switches, the output for the ptopoConnRemotePort MIB might display an incorrect value for portIDMacAddr. [PR/1061073: This issue has been resolved.]

Infrastructure

- On EX Series switches, when the **auto-negotiation** statement is configured at the **[edit interfaces interface-name ether-options]** hierarchy level, and the **configured-flow-control** statement is configured at the **[edit interfaces interface-name aggregated-ether-options]** hierarchy level, and both objects have the same attribute ID (**aid**), the CoS process (**cosd**) might generate a core file. [PR/837458: This issue has been resolved.]
- On EX Series switches, if multiple L3 and non-L2 sub-interfaces are enabled on a physical interface, and the family is deleted on a sub-interface or a sub-interface itself is deleted, traffic might be sent to the Routing Engine instead of the Packet Forwarding Engine, which can impact performance. [PR/1032503: This issue has been resolved.]
- On an EX8216 switch, if the Switch Interface Board (SIB) or the Switch Fabric (SF) module fails, there are no spare fabric planes available for switchover, which might cause a traffic outage. Depending on the nature of the SIB failure, the plane might need to be taken offline to resolve the issue. [PR/1037646: This issue has been resolved.]
- On EX4200 switches, high levels of traffic bound for the Routing Engine might cause the watchdog timer to expire, which in turn, causes the switch to reboot. This issue is seen with Protocol Independent Multicast (PIM) configurations when the multicast route is not present in the Packet Forwarding Engine for some amount of time, during which the multicast traffic for that route is routed to the CPU. [PR/1047142: This issue has been resolved.]
- On EX4200 and EX3200 switches, a high number of pause frames received on the switch interfaces might cause a soft reset of the switch. The following messages will be seen in **/var/log** when the switch undergoes a soft reset:

```
/kernel: simulated intr  
chassis[1293]: cm_java_pfe_critical_error_check: Soft-resetting device 1
```

If the pause frames are continuous and frequent, this might result in continuous soft reset of the Packet Forwarding Engine. The impact on traffic of a soft reset of the Packet Forwarding Engine is minor; however, if the switch is a member of a Virtual Chassis, continuous soft resets due to pause frames might cause the FPC to detach

from the Virtual Chassis, leading to other traffic related issues. [PR/1056787: This issue has been resolved.]

- On EX9200 switches, a software upgrade might cause firewall filters to redirect packets to an incorrect routing instance. [PR/1057180: This issue has been resolved.]

Layer 2 Features

- On EX Series switches, SNMP MAC notification traps are not generated if an interface goes down after a cable has been removed or disconnected, though traps are generated after the interface comes back up for MAC address removal, and also when a MAC address has been learned. [PR/1070638: This issue has been resolved.]
- On EX Series switches except EX4600 and EX9200 switches, when MSTP is configured, the Ethernet switching process (eswd) might generate multiple types of core files in the large-scale VLANs that are associated with Multiple Spanning-Tree Instances (MSTIs). [PR/1083395: This issue has been resolved.]

MPLS

- On EX4500 and EX8200 switches, if the switch is configured as a P router for MPLS, MPLS labels might be seen on the P router where the packets transit the Routing Engine on both input and output MPLS interfaces. This might lead to high CPU usage and can impact performance. [PR/1038618: This issue has been resolved.]

Virtual Chassis

- On an EX4550 Virtual Chassis with VLAN pruning enabled, if the LACP child interfaces span different Virtual Chassis members, then changing the Routing Engine mastership and rebooting the backup Routing Engine might cause the LACP child interfaces to remain in detached or passive state. [PR/1021554: This issue has been resolved.]

Issues Resolved in Release 12.3R11

The following issues have been resolved since Junos OS Release 12.3R10. The identifier following the description is the tracking number in our bug database.

Infrastructure

- On EX4500 and EX4550 switches, if you disable an interface on an EX-SFP-10GE-LR uplink module by issuing the CLI command **set interface interface-name disable**, and then the interface through which a peer device is connected to the interface on the uplink module goes down, the CPU utilization of the chassis manager process (chassism) might spike, causing chassism to create a core file. [PR/1032818: This issue has been resolved.]
- On EX4200, EX4500, EX6200, and EX8200 switches that are configured with distributed periodic packet management (PPM) mode, if you configure the Bidirectional Forwarding Detection (BFD) **minimum-receive-interval** value to a custom interval, BFD packets might be sent to a remote neighbor at a rate that exceeds the remote **minimum-receive-interval** value. As a workaround, configure PPM in centralized mode. [PR/1055830: This issue has been resolved.]

- On EX Series switches except EX9200 switches, if you configure both **family ethernet-switching** and **vlan-tagging** on the same interface, traffic might be dropped. [PR/1059480: This issue has been resolved.]
- On an EX8200 Virtual Chassis, if you configure **vlan-tagging** on an interface without configuring a family for the interface, the Packet Forwarding Engine might program an improper MAC address (the local chassis MAC) instead of the router MAC, which is used for routing. As a workaround, configure family **inet** on the interface. [PR/1060148: This issue has been resolved.]
- On EX Series switches except EX9200, configuring more than 1000 IPv4 addresses might prevent gratuitous ARP packets from being sent to peers. [PR/1062460: This issue has been resolved.]
- On EX4200 switches, if CoS scheduler maps are configured on all interfaces with the loss-priority value set to high, traffic between different Packet Forwarding Engines might be dropped. [PR/1071361: This issue has been resolved.]
- On EX3200 and EX4200 switches, if you apply a firewall filter to or remove it from a large number of interfaces, the Packet Forwarding Engine manager process (pfem) might generate a core file. [PR/1073055: This issue has been resolved.]
- On EX4500 and EX4550 Virtual Chassis, NFS/UDP fragmented packets might be dropped if these packets ingress over an aggregated bundle and traverse VCP links. [PR/1074105: This issue has been resolved.]
- On EX3300 switches, the output for the **show system license** command displays **invalid** for **connectivity-fault-management**. You can ignore this output; CFM is included in the EFL license. [PR/1087581: This issue has been resolved.]
- On EX Series switches, if you change the PIM mode from sparse to dense or dense to sparse, a pfem core file might be generated. [PR/1087730: This issue has been resolved.]
- On EX Series switches, the Packet Forwarding Engine Manager process (pfem) might crash and generate a core file when the TCAM is full. [PR/1107305: This issue has been resolved.]
- On EX4500 or EX4550 Virtual Chassis, if an NFS/UDP fragmented packet enters the Virtual Chassis through a LAG and traverses a Virtual Chassis port (VCP) link, CPU utilization might become high, and the software forwarding infrastructure process (sfid) might generate a core file. [PR/1109312: This issue has been resolved.]

Layer 2 Features

- On an EX3300 switch, in a broadcast storm situation in which DHCP snooping is enabled and there are repeated DHCP requests and acknowledgements arriving on the switch as a result of IP addresses not being accepted by clients, the eswd process might create a core file. [PR/1109312: This issue has been resolved.]
- On EX4200 switches with DHCP snooping configured, when a host moves from one interface to another interface and then renews its DHCP lease, the DHCP snooping database might not get updated, and thus the host might not connect on the new interface. [PR/1112811: This issue has been resolved.]

Platform and Infrastructure

- On EX4300 and EX9200 switches, the **show ethernet-switching table vlan-name vlan-name | display xml** CLI command does not have the **vlan-name** attribute in the `<l2ng-l2ald-rtb-macdb>` xml tag. [PR/955910: This issue has been resolved.]

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches

- [Changes to Junos OS for EX Series Switches Documentation on page 95](#)
- [Errata on page 96](#)

Changes to Junos OS for EX Series Switches Documentation

The following changes have been made to the documentation for Junos OS Release 12.3 for EX Series switches since it was published:

- The EZ Touchless Provisioning feature has been renamed *Zero Touch Provisioning (ZTP)*. The feature was introduced on EX Series switches in Junos OS Release 12.2. For more information, see [Understanding Zero Touch Provisioning](#).
- The EX2200 Virtual Chassis and the EX2200-C Virtual Chassis no longer require a software license. The document describing the software licenses for EX Series switches has been updated with this information. See [Understanding Software Licenses for EX Series Switches](#).
- The **request system software validate** command is not supported on EX Series switches. The documentation for the **request system software validate** command has been updated with this information. See [request system software validate](#). [This issue is being tracked by PR/821244.]
- The **request system software add** command **validate** option is not supported on EX Series switches. The documentation for the **request system software add** command has been updated with this information. See [request system software add](#). [This issue is being tracked by PR/821244.]

Errata

This section lists outstanding issues with the published documentation for Junos OS Release 12.3 for EX Series switches.

- The EX4500 switch models that support Converged Enhanced Ethernet (CEE) now also support IEEE Data Center Bridging Capability Exchange protocol (IEEE DCBX). These switches previously supported only DCBX version 1.01. The documentation does not reflect this support update. See [“New Features in Junos OS Release 12.3 for EX Series Switches” on page 28](#) for more information about the feature.
- You can configure VN_Port to VN_Port FIP snooping if the hosts are directly connected to the same EX4500 switch. See [Example: Configuring VN2VN_Port FIP Snooping \(FCoE Hosts Directly Connected to the Same FCoE Transit Switch\)](#) for details about this configuration. The documentation does not yet reflect this support update for EX4500 switches.
- The documentation for firewall filters on the switches states: *By default, a configuration that does not contain either ether-type or ip-version in a term applies to IPv4 traffic.* This is incorrect; the configuration must include the match condition **ether_type = ipv4** for an Ethernet-switching filter to be applied to only IPv4 traffic.
- The documentation for the EX9200 switches does not mention that the EX9200 switches do not process media access control (MAC) PAUSE frames.
- The documentation for the EX9200 switches does not mention that the EX9200 switches calculate the IRB interface **family inet** MTU by taking the minimum MTU of its Layer 2 members.
- The documentation for the 12.3 release does not mention the dedicated Virtual Chassis port (VCP) link aggregation feature on EX4550 switches. Starting in Junos OS Release 12.3R2, the dedicated VCPs on EX4550 switches automatically form a link aggregation group (LAG) bundle when two or more dedicated VCPs are used to interconnect the same Virtual Chassis member switches. An EX4550 switch can include up to four dedicated VCPs, and all four dedicated VCPs can act as member links in a LAG when they are used to interconnect to the same Virtual Chassis member switch. Dedicated VCPs and optical ports configured as VCPs cannot be member links in the same LAG and are placed into different LAGs when both are configured to connect to the same EX4550 member switch.
- The *OSPF Configuration Guide* incorrectly includes the **transmit-interval** statement at the **[edit protocols ospf area area interface interface-name]** hierarchy level. The **transmit-interval** statement at this hierarchy level is deprecated in the Junos OS CLI.
- The *NETCONF XML Management Protocol Guide* incorrectly states that when performing a confirmed commit operation using the **<commit>** element, the **<confirm-timeout>** value specifies the number of minutes for the rollback deadline. The value of the **<confirm-timeout>** element actually specifies the number of seconds for the rollback deadline.

[NETCONF XML Management Protocol Guide]

- Related Documentation**
- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
 - [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)
 - [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
 - [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
 - [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
 - [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 97](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches

This section discusses the following topics:

- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 97](#)
- [Upgrading EX Series Switches Using NSSU on page 98](#)
- [Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations on page 99](#)
- [Upgrading from Junos OS Release 10.4R3 or Later on page 99](#)
- [Upgrading from Junos OS Release 10.4R2 or Earlier on page 101](#)

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Upgrading EX Series Switches Using NSSU

You can use nonstop software upgrade (NSSU) to upgrade Junos OS releases on standalone EX6200 and EX8200 switches with dual Routing Engines and on EX3300, EX4200, EX4500, EX4550, and EX8200 Virtual Chassis. NSSU enables you to upgrade the software with a single command and minimal disruption to network traffic.

To optimize the quality and reliability of NSSU across multiple EX Series platforms and Junos OS releases, starting with Junos OS Release 12.3R6, NSSU support is limited to specific release combinations using the following guidelines:

- NSSU support is limited to $N-1$ to N and N to $N+1$ major release versions, where N represents a major release version such as 12.1, 12.2, or 12.3.
For example, NSSU from Release 11.4 to Release 12.2 is supported. NSSU from Release 11.x to Release 13.x is *not* supported.
- NSSU support is limited to $N.3$ and $N.6$ minor release versions.
For example, NSSU from Release 11.4R11 to Releases 12.1R3, 12.1R6, 12.2R3, 12.2R6, 12.3R3, and 12.3R6 is supported.
- Additional NSSU support for EOL releases is provided within the same major release version and is limited to two consecutive prior minor releases.
For example, NSSU from Release 11.4R9 or 11.4R10 to Release 11.4R11 is supported.

For details on the supported Junos OS release combinations for upgrading EX Series switches using NSSU, see the tables in [Junos OS Release Support for Upgrading EX Series Switches Using NSSU](#).

For details on NSSU, including procedures, see [Understanding Nonstop Software Upgrade on EX Series Switches](#).



NOTE: On a Virtual Chassis, you can use NSSU to upgrade from a domestic version of Junos OS to a controlled (MACsec) version of Junos OS. You cannot, however, use NSSU to upgrade from the controlled version of Junos OS to a domestic version of Junos OS.



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



NOTE: Do not use NSSU to upgrade the software on an EX8200 switch from Junos OS Release 10.4 if you have configured the IGMP, MLD, or PIM protocols on the switch. If you attempt to use NSSU, your switch might be left in a nonfunctional state from which it is difficult to recover. If you have these multicast protocols configured, upgrade the software on the EX8200 switch from Junos OS Release 10.4 by following the instructions in [Installing Software on an EX Series Switch with Redundant Routing Engines \(CLI Procedure\)](#). This issue does not apply to upgrades from Junos OS Release 11.1 or later.



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

Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations

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

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

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

Upgrading from Junos OS Release 10.4R3 or Later

This section contains the procedure for upgrading from Junos OS Release 10.4R3 or later to Junos OS Release 12.2. You can use this procedure to upgrade Junos OS on a standalone EX Series switch with a single Routing Engine and to upgrade all members of a Virtual Chassis or a single member of a Virtual Chassis.

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

On switches with dual Routing Engines or on Virtual Chassis, you might also be able to use nonstop software upgrade (NSSU) to upgrade Junos OS. See [“Upgrading EX Series Switches Using NSSU” on page 98](#) for more information.

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

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

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

4. Install the new software package on the switch:

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

Replace **package** with one of the following paths:

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

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

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

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

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

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

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



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

5. Reboot the switch to start the new software:

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

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

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

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

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

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

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

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

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

Upgrading from Junos OS Release 10.4R2 or Earlier

To upgrade to Junos OS Release 12.3 from Junos OS Release 10.4R2 or earlier, first upgrade to Junos OS Release 11.4 by following the instructions in the Junos OS Release 11.4 release notes. See *Upgrading from Junos OS Release 10.4R2 or Earlier* or *Upgrading from Junos OS Release 10.4R3 or Later* in the [Junos OS 11.4 Release Notes](#).

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 28](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 39](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 42](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 50](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 54](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 95](#)

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

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)
- [Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 181](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 357](#)

New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

The following features have been added to Junos OS Release 12.3. Following the description is the title of the manual or manuals to consult for further information.

- [Hardware on page 103](#)
- [Class of Service on page 104](#)
- [Forwarding and Sampling on page 110](#)
- [High Availability \(HA\) and Resiliency on page 110](#)
- [Interfaces and Chassis on page 111](#)
- [Junos OS XML API and Scripting on page 130](#)
- [Layer 2 Features on page 131](#)
- [Layer 2 Tunneling Protocol on page 133](#)
- [MPLS on page 134](#)
- [Multicast on page 136](#)
- [Power Management on page 139](#)
- [Routing Policy and Firewall Filters on page 139](#)
- [Routing Protocols on page 141](#)
- [Security on page 142](#)
- [Subscriber Access Management on page 142](#)
- [System Logging on page 153](#)
- [User Interface and Configuration on page 154](#)

- [VPLS on page 157](#)
- [VPNs on page 160](#)

Hardware

- **SFP-GE80KCW1470-ET, SFP-GE80KCW1490-ET, SFP-GE80KCW1510-ET, SFP-GE80KCW1530-ET, SFP-GE80KCW1550-ET, SFP-GE80KCW1570-ET, SFP-GE80KCW1590-ET, and SFP-GE80KCW1610-ET (MX Series)**—These transceivers provide a duplex LC connector and support operation and monitoring with links up to a distance of 80 km. Each transceiver is tuned to a different transmit wavelength for use in CWDM applications. These transceivers are supported on the following interface module. For more information about interface modules, see the *Interface Module Reference* for your router.
 - Gigabit Ethernet MIC with SFP (model number: MIC-3D-20GE-SFP) in all versions of MX-MPC1, MX-MPC2, and MX-MPC3 —Supported in Junos OS Release 12.3R5, 13.2R3, 13.3R1, and later.

[See [Gigabit Ethernet SFP CWDM Optical Interface Specification](#).]

- **CFP-GEN2-CGE-ER4 (MX Series, T1600, and T4000)**—The CFP-GEN2-CGE-ER4 transceiver (part number: 740-049763) provides a duplex LC connector and supports the 100GBASE-ER4 optical interface specification and monitoring. The “GEN2” optics have been redesigned with newer versions of internal components for reduced power consumption. The following interface modules support the CFP-GEN2-CGE-ER4 transceiver. For more information about interface modules, see the *Interface Module Reference* for your router.

MX Series routers:

- 100-Gigabit Ethernet MIC with CFP (model number: MIC3-3D-1X100GE-CFP)—Supported in Junos OS Release 12.1R1 and later
- 2x100GE + 8x10GE MPC4E (model number: MPC4E-3D-2CGE-8XGE)—Supported in Junos OS Release 12.3R2 and later

T1600 and T4000 routers:

- 100-Gigabit Ethernet PIC with CFP (model numbers: PD-1CE-CFP-FPC4 and PD-1CGE-CFP)—Supported in Junos OS Releases 12.3R5, 13.2R3, 13.3R1, and later

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#).]

- **CFP-GEN2-100GBASE-LR4 (T1600 and T4000)**—The CFP-GEN2-100GBASE-LR4 transceiver (part number: 740-047682) provides a duplex LC connector and supports the 100GBASE-LR4 optical interface specification and monitoring. The “GEN2” optics have been redesigned with newer versions of internal components for reduced power consumption. The following interface modules support the CFP-GEN2-100GBASE-LR4 transceiver. For more information about interface modules, see the *Interface Module Reference* for your router.

- 100-Gigabit Ethernet PIC with CFP (model numbers: PD-1CE-CFP-FPC4 and PD-1CGE-CFP)—Supported in Junos OS Releases 12.3R5, 13.2R3, 13.3R1, and later

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#).]

- **Support for 32 GB RE-S-1800 Routing Engine (MX Series)**—Starting with Junos OS Release 12.3R4, the 32-GB RE-S-1800 Routing Engine (MX240, MX480, and MX960 model number RE-S-1800X4-32G-S; MX2010 and MX2020 model number REMX2K-1800-32G-S) is supported on MX Series routers. The 32 GB RE-S-1800 Routing Engine has a 1800 GHz processor with 32-GB of memory, and supports both 32-bit and 64-bit Junos OS builds. On the MX2010 and MX2020 routers, the Routing Engine supports only 64-bit Junos OS build. The new Routing Engine supports a 4-GB CompactFlash card. All CLI commands supported on the older Routing Engines are supported on the new Routing Engine.

Class of Service

- **Support for Class-of-Service Features to Ensure Quality of Service for Real-Time Traffic That Is Sensitive to Latency on a Network (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**—The new Application Services Modular Line Card (AS MLC) supports the following CoS features on MX240, MX480, and MX960 routers:
 - Code-point aliases—A code-point alias is a meaningful name that can be associated with CoS values such as Differentiated Services code points (DSCPs), DSCP IPv6, IP precedence, IEEE 802.1, and MPLS experimental (EXP) bits that can then be used while configuring CoS components.
 - Classification—Packet classification associates the packet with a particular CoS servicing level. In Junos OS, classifiers associate incoming packets with a forwarding class and loss priority and, based on the associated forwarding class, assign packets to output queues.
 - Behavior Aggregate—A method of classification that operates on a packet as it enters the router.
 - Multifield Classification— A method of classification that can examine multiple fields in the packet.
 - Fixed Classification—A method of classification that refers to the association of a forwarding class with a packet regardless of its packet contents.

[See [Class of Service on Application Services Modular Line Card Overview](#).]

- Scheduling—Schedulers are used to define the properties of output queues. On the AS modular carrier card (AS MCC), the following scheduling features are supported (physical interfaces only):
 - Buffer sizes
 - Delay buffer size
 - Drop profile map
 - Excess priority

- Excess rate percentage
- Output-traffic-control profile
- Priority
- Scheduler-map
- Shaping rate
- Transmit rate
- WRED rules

[*Junos OS Class-of-Service Configuration Guide*]

- **Setting the 802.1p field for host-generated traffic**—On MPCs and Enhanced Queuing DPCs, you can now configure the IEEE 802.1p bits in the 802.1p field—also known as the Priority Code Point (PCP) field—in the Ethernet frame header for host outbound packets (control plane traffic). In earlier releases, this field is not configurable; instead it is set by CoS automatically for host outbound traffic.

To configure a global default value for this field for all host outbound traffic, include the **default value** statement at the [**edit class-of-service host-outbound-traffic ieee-802.1**] hierarchy level. This configuration has no effect on data plane traffic; you configure rewrite rules for these packets as always.

You cannot configure a default value for the 802.1p bits for host outbound traffic on a per-interface level. However, you can specify that the CoS 802.1p rewrite rules already configured on egress logical interfaces are applied to all host outbound packets on that interface. To do so, include the **rewrite-rules** statement at the [**edit class-of-service host-outbound-traffic ieee-802.1**] hierarchy level. This capability enables you to set only the outer tags or both the outer and the inner tags on dual-tagged VLAN packets. (On Enhanced Queuing DPCs, both inner and outer tags must be set.)

This feature includes the following support:

- Address families—IPv4 and IPv6
- Interfaces—IP over VLAN demux, PPP over VLAN demux, and VLAN over Gigabit Ethernet
- Packet types—ARP, ANCP, DHCP, ICMP, IGMP, and PPP
- VLANs—Single and dual-tagged

[*Class of Service*]

- **Software feature support on the MX2020 routers**—Starting with Release 12.3, all MPCs and MICs supported on the MX Series routers in Junos OS Release 12.3 continue to be supported on the MX2020 routers. Also, the MX2020 routers support all software features that are supported by other MX Series routers in Junos OS Release 12.1.

The following key Junos OS features are supported:

- Basic Layer 2 features including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
- Class-of-service (CoS)
- Firewall filters and policers
- Integrated Routing and Bridging (IRB)
- Interoperability with existing DPCs and MPCs
- Layer 2 protocols
- Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs
- Layer 3 routing protocols and MPLS
- Multicast forwarding
- Port mirroring
- Synchronous Ethernet and Precision Time Protocol (IEEE 1588)
- Tunnel support
- Spanning Tree Protocols (STP)

[*Class of Service, Ethernet Interfaces Configuration Guide, System Basics and Services Command Reference*]

- **Ingress CoS on MIC and MPC interfaces (MX Series routers)**—You can configure ingress CoS parameters, including hierarchical schedulers, on MX Series routers with MIC and MPC interfaces. In general, the supported configuration statements apply to per-unit schedulers or to hierarchical schedulers.

To configure ingress CoS for per-unit schedulers, include the following statements at the **[edit class-of-service interfaces interface-name]** hierarchy level:

```
input-scheduler-map
input-shaping-rate
input-traffic-control-profile
input-traffic-control-profile-remaining
```

To configure ingress CoS for hierarchical schedulers, include the **interface-set interface-set-name** statement at the **[edit class-of-service interfaces]** hierarchy level.



NOTE: The interface-set statement supports only the following options:

```
input-traffic-control-profile
input-traffic-control-profile-remaining
```

To configure ingress CoS at the logical interface level, include the following statements at the **[edit class-of-service interfaces interface interface-name unit logical-unit-number]** hierarchy level:

```
input-scheduler-map
input-shaping-rate
input-traffic-control-profile
```

[See [Configuring Ingress Hierarchical CoS on MIC and MPC interfaces.](#)]

- **Extends explicit burst size configuration support on IQ2 and IQ2E interfaces**—The burst size for shapers can be configured explicitly in a traffic control profile for IQ2 and IQ2E interfaces. This feature is supported on M71, M10i, M40e, M120, M320, and all T Series routers.

To enable this feature, include the **burst-size** statement at the following hierarchy levels:

```
[edit class-of-service traffic-control-profiles shaping-rate]
[edit class-of-service traffic-control-profiles guaranteed-rate]
```



NOTE: The **guaranteed-rate** burst size value cannot be greater than the **shaping-rate** burst size.

[See [Configuring Traffic Control Profiles for Shared Scheduling and Shaping.](#)]

- **Classification and DSCP Marking of Distributed Protocol Handler Traffic**—The scope of traffic affected by the **host-outbound-traffic** statement is expanded. When it was introduced in Junos OS Release 8.4, the **host-outbound-traffic** statement at the **[edit class-of-service]** hierarchy level enabled you to specify the forwarding class assignment and DiffServ code point (DSCP) value for egress traffic sent from the Routing Engine. Affected traffic included control plane packets (such as OSPF hello and ICMP echo reply [ping] packets) and TCP-related packets (such as BGP and LDP control packets).

In Junos OS 12.2R2, the same configuration applies to *distributed protocol handler traffic* in addition to Routing Engine traffic. Distributed protocol handler traffic refers to traffic from the router's periodic packet management process (ppm) sessions, and it includes both IP (Layer 3) traffic such as BFD keepalive (KA) messages and non-IP (Layer 2) traffic such as LACP control traffic on aggregated Ethernet. DSCP changes do not apply to MPLS EXP bits or IEEE 802.1p bits. The specified queue must be correctly configured. The affected traffic includes distributed protocol handler traffic as well as Routing Engine traffic for egress interfaces hosted on MX Series routers with Trio-based or I-chip based Packet Forwarding Engines, and on M120, M320, and T Series routers.

If you need the Routing Engine traffic and distributed protocol handler traffic to be classified in different forwarding classes or marked with different DSCP values, then you need to configure some additional steps. Apply a standard firewall filter to the loopback interface and configure the filter actions to set the forwarding class and DSCP value that override the **host-outbound-traffic** settings.

For interfaces on MX80 routers, LACP control traffic is sent through the Routing Engine rather than through the Packet Forwarding Engine.



NOTE: Any DSCP rewrite rules configured on a 10-Gigabit Ethernet LAN/WAN PIC with SFP+ overwrite the DSCP value rewritten as specified under the **host-outbound-traffic** statement.

The following partial configuration example classifies egress traffic from the Routing Engine as well as distributed protocol handler traffic:

```
[edit]
class-of-service {
  host-outbound-traffic {
    forwarding-class my_fc_control-traffic_dph;
    dscp-code-point 001010;
  }
  forwarding-classes {
    queue 5 my_fc_control-traffic_dph;
    queue 6 my_fc_control-traffic_re;
  }
}
interfaces {
  lo0 {
    unit 0 {
      family inet {
        filter {
          output my_filter_reclassify_re;
        }
      }
    }
  }
}
firewall {
  filter my_filter_reclassify_re {
    term 1 {
      then {
        forwarding-class my_fc_control-traffic_re;
        dscp code-points 000011;
        accept;
      }
    }
  }
}
```

The statements in the example configuration cause the router to classify egress traffic from the Routing Engine and distributed protocol handler traffic as follows:

- Distributed protocol handler traffic is classified to the `my_fc_control-traffic_dph` forwarding class, which is mapped to queue 5. Of those packets, Layer 3 packets are marked at egress with DSCP bits 001010 (10 decimal), which is compatible with ToS bits 00101000 (40 decimal).
- Routing Engine traffic is classified to the `my_fc_control-traffic_re` forwarding class, which is mapped to queue 6. Of those packets, Layer 3 packets are marked at egress with DSCP bits 001100 (12 decimal), which is compatible with ToS bits 00110000 (48 decimal).

If you do not apply the firewall filter to the loopback interface, Routing Engine-sourced traffic is classified and marked using the forwarding class and DSCP value specified in the **host-outbound-traffic** configuration statement.

If you omit both the firewall filter and the **host-outbound-traffic** configuration shown in the previous configuration, then all network control traffic—including Routing

Engine-sourced and distributed protocol handler traffic—uses output queue 3 (the default output queue for control traffic), and DSCP bits for Layer 3 packets are set to the default value 0 (Best Effort service).

- **Enhancements to scheduler configuration on FRF.16 physical interfaces**—Starting with Release 12.3R2, Junos OS extends the class-of-service scheduler support on FRF.16 physical interfaces to the **excess-rate**, **excess-priority**, and **drop-profile-map** configurations. The **excess-rate**, **excess-priority**, and **drop-profile-map** statements are configured at the **[edit class-of-service schedulers scheduler-name]** hierarchy level.
 - Support for the **drop-profile-map** configuration enables you to configure random early detection (RED) on FRF.16 bundle physical interfaces.
 - Support for the **excess-rate** configuration enables you to specify the percentage of the excess bandwidth traffic to share.
 - Support for the **excess-priority** configuration enables you to specify the priority for excess bandwidth traffic on a scheduler.

This feature is supported only on multiservices PICs installed on MX Series routers.

- **Accurate reporting of output counters for MLFR UNI NNI bundles**—Starting with Release 12.3R2, Junos OS reports the actual output counters in the multilink frame relay (MLFR) UNI NNI bundle statistics section of the **show interfaces lsq-interface statistics** command output. From this release on, Junos OS also provides per-DLCI counters for logical interfaces. In earlier releases, there was a discrepancy between the actual output counters and the reported value because of errors in calculating the output counters at the logical interface level. That is, at the logical interface level, the output counter was calculated as the sum of frames egressing at the member links instead of providing the output counter as the sum of per-DLCI output frames.
- **Extended MPC support for per-unit schedulers**—Enables you to configure per-unit schedulers on the non-queuing 16x10GE MPC and the MPC3E, meaning you can include the **per-unit-scheduler** statement at the **[edit interfaces interface name]** hierarchy level. When per-unit schedulers are enabled, you can define dedicated schedulers for logical interfaces by including the **scheduler-map** statement at the **[edit class-of-service interfaces interface name unit logical unit number]** hierarchy level. Alternatively, you can include the **scheduler-map** statement at the **[edit class-of-service traffic-control-profiles traffic control profile name]** hierarchy level and then include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces interface name unit logical unit number]** hierarchy level.

Enabling per-unit schedulers on the 16x10GE MPC and the MPC3E adds additional output to the **show interfaces interface name [detail | extensive]** command. This additional output lists the maximum resources available and the number of configured resources for schedulers.

[Applying Scheduler Maps and Shaping Rate to DLCIs and VLANs]

Forwarding and Sampling

- **Increased forwarding capabilities for MPCs and Multiservices DPCs through FIB localization (MX Series routers)**—Forwarding information base (FIB) localization characterizes the Packet Forwarding Engines in a router into two types: FIB-Remote and FIB-Local. FIB-Local Packet Forwarding Engines install all of the routes from the default route tables into Packet Forwarding Engine forwarding hardware. FIB-Remote Packet Forwarding Engines create a default (0.0) route that references a next hop or a unilist of next hops to indicate the FIB-Local that can perform full IP table looks-ups for received packets. FIB-Remote Packet Forwarding Engines forward received packets to the set of FIB-Local Packet Forwarding Engines.

The capacity of MPCs is much higher than that of Multiservices DPCs, so an MPC is designated as the local Packet Forwarding Engine, and a Multiservices DPC is designated as the remote Packet Forwarding Engine. The remote Packet Forwarding Engine forwards all network-bound traffic to the local Packet Forwarding Engine. If multiple MPCs are designated as local Packet Forwarding Engines, then the Multiservices DPC will load-balance the traffic using the unilist of next hops as the default route.

High Availability (HA) and Resiliency

- **Protocol Independent Multicast Nonstop Active Routing Support for IGMP-Only Interfaces**—Starting with Release 12.3, Junos OS extends the Protocol Independent Multicast (PIM) nonstop active routing support to IGMP-only interfaces.

In Junos OS releases earlier than 12.3, the PIM joins created on IGMP-only interfaces were not replicated on the backup Routing Engine and so the corresponding multicast routes were marked as pruned (meaning discarded) on the backup Routing Engine. Because of this limitation, after a switchover, the new master Routing Engine had to wait for the IGMP module to come up and start receiving reports to create PIM joins and to install multicast routes. This causes traffic loss until the multicast joins and routes are reinstated.

However, in Junos OS Release 12.3 and later, the multicast joins on the IGMP-only interfaces are mapped to PIM states, and these states are replicated on the backup Routing Engine. If the corresponding PIM states are available on the backup, the multicast routes are marked as forwarding on the backup Routing Engine. This enables uninterrupted traffic flow after a switchover. This enhancement covers IGMPv2, IGMPv3, MLDv1, and MLDv2 reports and leaves.

[High Availability]

- RSVP

Nonstop active routing support for RSVP includes:

- Point-to-Multipoint LSPs
 - RSVP Point-to-Multipoint ingress, transit, and egress LSPs using existing non-chained next hop.
 - RSVP Point-to-Multipoint transit LSPs using composite next hops for Point-to-Multipoint label routes.

- Point-to-Point LSPs
 - RSVP Point-to-Point ingress, transit, and egress LSPs using non-chained next hops.
 - RSVP Point-to-Point transit LSPs using chained composite next hops.
- **Configuration support to include GTP TEID field in hash key for load-balancing GTP-U traffic (MX Series routers with MPCs and MX80)**—On an MX Series router with MPCs, when there are multiple equal-cost paths to the same destination for the active route, Junos OS uses a hash algorithm to choose one of the next-hop addresses from the forwarding table when making a forwarding decision. Whenever the set of next hops for a destination changes in any way, the next-hop address is rechosen using the hash algorithm. For GPRS tunneling protocol (GTP)-encapsulated traffic, the tunnel endpoint identifier (TEID) field changes for traffic traversing through peer routers. To implement load balancing for GTP-encapsulated traffic on the user plane (GTP-U), the TEID should be included in the hash key.

In Junos OS Release 12.3R2, you can configure GTP hashing on MX Series routers with MPCs and on MX80, to include the TEID field in hash calculations for IPv4 and IPv6 packets. To configure GTP hashing and include the GTP TEID field in hash calculations, configure the **gtp-tunnel-end-point-identifier** statement at the **[edit forwarding-options enhanced-hash-key family]** hierarchy level. GTP hashing is supported for both IPv4 and IPv6 packets received for GTP-U traffic at the MPC. For bridging and MPLS packets, GTP hashing is supported for IPv4 and IPv6 packets that are carried as payload for GTP-encapsulated traffic.



NOTE: For IPv4 packets, GTP hashing is supported only for the nonfragmented packets.

[\[Overview of Per-Packet Load Balancing\]](#)

Interfaces and Chassis

- **Support for Fabric Management Features (MX240, MX480, MX960 Routers with Application Services Modular Carrier Card)**—The Application Services Module Line Card (AS MLC) is supported on MX240, MX480, and MX960 routers. The AS MLC consists of the following components:
 - Application Services Modular Carrier Card (AS MCC)
 - Application Services Modular Processing Card (AS MXC)
 - Application Services Modular Storage Card (AS MSC)

The AS MCC plugs into the chassis and provides the fabric interface. On the fabric management side, the AS MLC provides redirection functionality using a demultiplexer. The following CLI operational mode commands display fabric-related information for the AS MCC:

- show chassis fabric fpcs
- show chassis fabric map

- show chassis fabric plane
- show chassis fabric plane-location
- show chassis fabric reachability
- show chassis fabric summary

[*Junos OS System Basics Configuration Guide, Junos OS System Basics and Services Command Reference*]

[See [Fabric Plane Management on AS MLC Modular Carrier Card Overview](#).]

- **Support for Chassis Management (MX240, MX480, MX960 Routers with Application Services Modular Line Card)**—The Application Services Modular Line Card (AS MLC) is a Modular Port Concentrator (MPC) that is designed to run services and applications on MX240, MX480, and MX960 routers.

The following CLI operational mode commands support the chassis management operations of the modular carrier card on the AS MLC:

- show chassis environment fpc
- show chassis firmware
- show chassis fpc
- show chassis hardware
- show chassis pic
- show chassis temperature-thresholds
- request chassis fpc
- request chassis mic
- request chassis mic fpc-slot mic-slot

[*Junos OS System Basics Configuration Guide, Junos OS System Basics and Services Command Reference*]

- **16-Port Channelized E1/T1 Circuit Emulation MIC (MX Series routers)**—Starting with Junos OS Release 12.3, the 16-Port Channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE) is supported on MX80, MX240, MX480, and MX960 routers. [See [16-Port Channelized E1/T1 Circuit Emulation MIC Overview](#).]

- **Extends signaling support for SAToP/CESoPSN for E1/T1 interfaces (MX Series routers)**—Starting with Junos OS Release 12.3, the E1/T1 interfaces support signaling for Structure-Agnostic TDM over Packet (SAToP) and Circuit Emulation Services over Packet-Switched Network (CESoPSN) through Layer 2 VPN using BGP.

[See [Configuring SAToP on Channelized E1/T1 Circuit Emulation MIC](#) and [Configuring CESoPSN on Channelized E1/T1 Circuit Emulation MIC](#).]

- **Extends support for diagnostic, OAM, and timing features to 16-port Channelized E1/T1 Circuit Emulation MIC (MX Series routers)**—Starting with Junos OS Release 12.3, the 16-port Channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE) supports the following features:

- Diagnostic features:
 - Loopback: Support for E1/T1-level payload, local line, remote line, and NxDSO payload loopbacks.
 - Bit error rate test (BERT): Support for the following BERT algorithms:
 - pseudo-2e11-o1520
 - pseudo-2e15-o152
 - pseudo-2e20-o150
- Operation, Administration, and Maintenance (OAM) features:
 - Performance monitoring: Supports the following Layer 1 performance-monitoring statistics at the E1/T1 interface level for all kinds of encapsulations:
 - E1 interfaces
 - BPV—Bipolar violation
 - EXZ—Excessive zeros
 - SEF—Severely errored framing
 - BEE—Bit error event
 - LCV—Line code violation
 - PCV—Pulse code violation
 - LES—Line error seconds
 - ES—Errored seconds
 - SES—Severely errored seconds
 - SEFS—Severely errored framing seconds
 - BES—Bit error seconds
 - UAS—Unavailable seconds
 - FEBE—Far-end block error
 - CRC—Cyclic redundancy check errors
 - LOFS—Loss of frame seconds
 - LOSS—Loss of signal seconds
 - T1 interfaces
 - BPV—Bipolar violation
 - EXZ—Excessive zeros
 - SEF—Severely errored framing
 - BEE—Bit error event
 - LCV—Line code violation

- PCV—Pulse code violation
 - LES—Line error seconds
 - ES—Errored seconds
 - SES—Severely errored seconds
 - SEFS—Severely errored framing seconds
 - BES—Bit error seconds
 - UAS—Unavailable seconds
 - LOFS—Loss of frame seconds
 - LOSS—Loss of signal seconds
 - CRC—Cyclic redundancy check errors
 - CRC Major—Cyclic redundancy check major alarm threshold exceeded
 - CRC Minor—Cyclic redundancy check minor alarm threshold exceeded
- Timing features: Support for the following transmit clocking options on the E1/T1 interface:
 - Looped timing
 - System timing



NOTE: In Junos OS Release 12.3, IMA Link alarms are not supported on the 16-port Channelized E1/T1 MIC.

[See [Configuring E1 Loopback Capability](#), [Configuring E1 BERT Properties](#), and [Interface Diagnostics](#).]

- **Extends support for SAToP features to 16-port Channelized E1/T1 Circuit Emulation MIC (MX Series routers)**—Starting with Junos OS Release 12.3, the 16-port Channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE) supports E1/T1 SAToP features.

[See [Configuring SAToP on Channelized E1/T1 Circuit Emulation MIC](#).]

- **CESoPSN encapsulation support extended to 16-Port Channelized E1/T1 Circuit Emulation MIC (MX Series routers)**—Starting with Junos OS Release 12.3, support for CESoPSN encapsulation is extended to the 16-port Channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE).

[See [Configuring CESoPSN on Channelized E1/T1 Circuit Emulation MIC](#).]

- **SNMP and MIB support (MX2020 routers)**—Starting with Junos OS Release 12.3, the enterprise-specific Chassis Definitions for Router Model MIB, `jnx-chas-defines.mib`, is updated to include information about the new MX2020 routers. The Chassis Definitions for Router Model MIB contains the object identifiers (OIDs) used by the Chassis MIB to identify platform and chassis components of each router.

[See [jnxBoxAnatomy](#), [Chassis Definitions for Router Model MIB](#), and [MIB Objects for the MX2020 3D Universal Edge Router](#).]

- **Junos OS support for FRU management of MX2020 routers**—Starting with Release 12.3, Junos OS supports the new MX2020 routers. The MX2020 routers are the next generation of MX Series 3D Universal Edge Routers. The Junos OS chassis management software for the MX2020 routers provides enhanced environmental monitoring and field-replaceable unit (FRU) control. FRUs supported on the MX2020 routers include:
 - RE and CB—Routing Engine and Control Board including a Processor Mezzanine Board (PMB)
 - PDM—Power distribution module
 - PSM—Power supply module
 - Fan trays
 - SFB—Switch Fabric Board
 - Front panel display
 - Adapter cards
 - Line cards

The MX2020 router supports up to two Control Boards (CBs) with the second CB being used as a redundant CB. The CB provides control and monitoring functions for the router. Adapter card and switch fabric board FRU management functionality is controlled by a dedicated processor housed on the Processor Mezzanine Board. The MX2020 router supports 20 adapter cards and 8 Switch Fabric Boards (SFBs).

The MX2020 chassis has two cooling zones. Fans operating in one zone have no impact on cooling in another zone, enabling the chassis to run fans at different speeds in different zones. The chassis can coordinate FRU temperatures in each zone and the fan speeds of the fan trays in these zones.

The power system on the MX2020 routers consists of three components: the power supply modules (PSMs), the power distribution module (PDM), and the power midplane. The MX2020 router chassis supplies $N + N$ feed redundancy, $N + 1$ power supply redundancy for line cards, and $N + N$ power supply redundancy for the critical FRUs. The critical FRUs include two CBs, eight SFBs, and three fan trays (two fan trays in one zone and one fan tray in the other zone.) In cases where all PSMs are not present, or some PSMs fail or are removed during operation, service interruption is minimized by keeping the affected FPCs online without supplying redundant power to these FPCs. You can use the following configuration statement to monitor power management on the switch chassis:

- **fru-poweron-sequence**—Include the **fru-poweron-sequence** statement at the **[edit chassis]** hierarchy level to configure the power-on sequence for the FPCs in the chassis.

Table 1: Maximum FRUs Supported on the MX2020 Router

FRU	Maximum Number
Routing Engines and CB	2

Table 1: Maximum FRUs Supported on the MX2020 Router (*continued*)

FRU	Maximum Number
PDM	4
PSM	18
Fan trays	4
SFB	8
Front panel display	1
Adapter cards	20
Line cards	20

The following CLI operational mode commands support the various FRU and power management operations on MX2020 routers:

Show commands:

- `show chassis adc`
- `show chassis alarms`
- `show chassis environment`
- `show chassis environment adc adc-slot-number`
- `show chassis environment cb cb-slot-number`
- `show chassis environment fpc fpc-slot-number`
- `show chassis environment fpm fpm-slot-number`
- `show chassis environment monitored`
- `show chassis environment psm psm-slot-number`
- `show chassis environment routing-engine routing-engine-slot-number`
- `show chassis environment sfb sfb-slot-number`
- `show chassis craft-interface`
- `show chassis ethernet-switch < errors | statistics >`
- `show chassis fabric destinations`
- `show chassis fabric fpcs`
- `show chassis fabric plane`
- `show chassis fabric plane-location`
- `show chassis fabric summary`
- `show chassis fan`

- `show chassis firmware`
- `show chassis fpc < detail | pic-status | fpc-slot-number >`
- `show chassis hardware < clei-models | detail | extensive | models >`
- `show chassis in-service-upgrade`
- `show chassis mac-addresses`
- `show chassis network-services`
- `show chassis pic fpc-slot fpc-slot-number pic-slot pic-slot-number`
- `show chassis power`
- `show chassis power sequence`
- `show chassis routing-engine < routing-engine-slot-number | bios >`
- `show chassis sfb < slot sfb-slot-number >`
- `show chassis spmb`
- `show chassis temperature-thresholds`
- `show chassis zones < detail >`

Request commands:

- `request chassis cb (offline | online) slot slot-number`
- `request chassis fabric plane (offline | online) fabric-plane-number`
- `request chassis fpc (offline | online | restart) slot fpc-slot-number`
- `request chassis fpm resync`
- `request chassis mic (offline | online) fpc-slot fpc-slot-number mic-slot mic-slot-number`
- `request chassis routing-engine master (acquire | release | switch) < no-confirm >`
- `request chassis sfb (offline | online) slot sfb-slot-number`
- `request chassis spmb restart slot spmb-slot-number`

Restart command:

- `restart chassis-control < gracefully | immediately | soft >`

For details of all system management operational mode commands and the command options supported on the MX2020 router, see the System Basics and Services Command Reference.

[See [System Basics: Chassis-Level Features Configuration Guide](#).]

- **SAToP support extended to Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX Series routers)**—Starting with Junos OS Release 12.3R1, support for Structure-Agnostic Time-Division Multiplexing over Packet (SAToP) is extended to MIC-3D-4COC3-1COC12-CE. You can configure 336 T1 channels on each COC12 interface on this MIC.

[See [Configuring SAToP on Channelized OC3/STM1 \(Multi-Rate\) Circuit Emulation MIC with SFP](#) and [Configuring SAToP Encapsulation on T1/E1 Interfaces on Channelized OC3/STM1 \(Multi-Rate\) Circuit Emulation MIC with SFP](#).]

- **CESoPSN support extended to Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX Series routers)**—Starting with Junos OS Release 12.3, support for Circuit Emulation Service over Packet-Switched Network (CESoPSN) is extended to MIC-3D-4COC3-1COC12-CE. You can configure 336 CT1 channels on each COC12 interface on this MIC.

[See [Configuring CESoPSN on Channelized OC3/STM1 \(Multi-Rate\) Circuit Emulation MIC with SFP](#) and [Configuring CESoPSN Encapsulation on DS Interfaces on Channelized OC3/STM1 \(Multi-Rate\) Circuit Emulation MIC with SFP](#).]

- **Support for ATM PWE3 on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX80 routers with a modular chassis, and MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.3, ATM Pseudowire Emulation Edge to Edge (PWE3) is supported on channelized T1/E1 interfaces of the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE). The following PWE3 features are supported:
 - ATM pseudowire encapsulation. The pseudowire encapsulation can be either cell-relay or AAL5 transport mode. Both modes enable the transport of ATM cells across a packet-switched network (PSN).
 - Cell-relay VPI/VCI swapping. The Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP can overwrite the virtual path identifier (VPI) and virtual channel identifier (VCI) header values on egress and on both ingress and egress.



NOTE: Cell-relay VPI swapping on both ingress and egress is not compatible with the ATM policing feature.

To configure the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP to modify both the VPI and VCI header values on both ingress and egress, you must specify the **psn-vci** statement at the following hierarchy level:

[edit interface at-*interface-name*/pic/port unit *logical-unit-number*]

To configure the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP to modify only the VPI header values on both ingress and egress, you must specify the **psn-vpi** statement at the following hierarchy level:

[edit interface at-*interface-name*/pic/port unit *logical-unit-number*]

To configure the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP to pass the VPI and VCI header values transparently, you must specify the **no-vpivci-swapping** statement at the following hierarchy level:

[edit interface at-*interface-name*/pic/port unit *logical-unit-number*]

If none of the aforementioned configuration statements are included, for virtual path pseudowires, VPI values are modified on egress, whereas for virtual channel pseudowires, both VPI and VCI header values are modified on egress.

[See [Configuring ATM Cell-Relay Pseudowire](#).]

- **Pseudowire ATM MIB support for Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX80 routers with a modular chassis, and MX240, MX480, and MX960 routers)**—Starting with Release 12.3, Junos OS extends Pseudowire ATM MIB support to the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE).

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

- **Multiple VRRP owners per physical port**—Support for multiple owner addresses per physical interface, allowing users to reuse interface address identifiers (IFAs) as virtual IP addresses (VIPs).
- **Chassis daemon enhancements for the MFC application on the Routing Engine**—The **chassisd** (chassis daemon) process runs on the Routing Engine to communicate directly with its peer processes running on the Packet Forwarding Engine. Starting with Junos OS Release 12.1, the **chassisd** process has been enhanced to enable the Media Flow Controller (MFC) application to run on a Dense Port Concentrator (DPC) with an x86 blade for high application throughput and a large amount of solid state storage on MX Series routers. The **chassisd** process detects the installation of the modular x86 blade for MFC services and monitors the physical status of hardware components and the field-replaceable units (FRUs) that enable MFC to be run on the x86 blade.

[*System Basics*]

- **Support for aggregated SONET/SDH Interfaces (MX Series Routers)**—Starting with Junos OS Release 12.3, you can configure aggregated SONET bundles with the member links of SONET/SDH OC3/STM1 (Multi-Rate) MICs with SFP, that is, MIC-3D-8OC3OC12-4OC48 and MIC-3D-4OC3OC12-1OC48.

Junos OS enables link aggregation of SONET/SDH interfaces; this is similar to Ethernet link aggregation, but is not defined in a public standard. Junos OS balances traffic across the member links within an aggregated SONET/SDH bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-packet load balancing.

The following features are supported on MIC-3D-8OC3OC12-4OC48 and MIC-3D-4OC3OC12-1OC48:

- Encapsulation—Point-to-Point Protocol (PPP) and Cisco High-Level Data Link Control (Cisco HDLC)
- Filters and policers—Single-rate policers, three-color marking policers, two-rate three-color marking policers, hierarchical policers, and percentage-based policers. By default, policer bandwidth and burst size applied on aggregated bundles are not matched to the user-configured bandwidth and burst size.
- Mixed mode links
- **Support for synchronizing an MX240, MX480, or MX960 router chassis with an Enhanced MX SCB to an external BITS timing source**—This feature uses the Building Integrated Timing Supply (BITS) external clock interface (ECI) on the Enhanced MX SCB. The BITS ECI can also be configured to display the selected chassis clock source (SETS) or a recovered line clock source (Synchronous Ethernet or Precision Time

Protocol). You can configure the BITS ECI by using the **synchronization** statement at the **[edit chassis]** hierarchy level. You can view the BITS ECI information with the **show chassis synchronization extensive** command.

- **Aggregated interfaces support increased to 64 links (MX Series)**—This feature adds support for specifying up to 64 links for aggregated devices. You set the number of links in the new **maximum-links** statement at the **[chassis aggregated-devices]** hierarchy level.
- **Junos OS support for new MX2010 routers**—Starting with Release 12.3, Junos OS supports the new MX2010 routers. The MX2010 routers are an extension of the MX2020 routers and support all features supported by the MX2020 routers. Also, the MX2010 routers support all software features that are supported by other MX Series routers in Junos OS Release 12.1.

The power system on the MX2010 routers consists of three components: the power supply modules (PSMs), the power distribution module (PDM), and the power midplane. The power feed (AC or DC) is connected to the PDM. The PDM delivers the power from the feeds to the power midplane. The power from the power midplane is provided to the PSMs. Output from the PSMs is sent back to the power midplane and then eventually to the field-replaceable units (FRUs). The MX2010 router chassis supplies $N + N$ feed redundancy and $N + 1$ PSM redundancy for line cards. In case some PSMs fail or are removed during operation, service interruption is minimized by keeping as many affected FPCs online by supplying redundant power to these FPCs. Unlike the MX2020 router chassis, the MX2010 router chassis does not provide redundancy for the critical FRUs because there is only one power zone.

Include the following existing configuration statement at the **[edit chassis]** hierarchy level to configure the power-on sequence for the FPCs in the chassis:

[edit chassis]

fru-poweron-sequence *fru-poweron-sequence*

Junos OS also supports the following CLI operational mode commands for chassis management of MX2010 routers:

<i>Show commands</i>	<i>Request commands</i>	<i>Restart commands</i>
show chassis adc	request chassis cb (offline online) slot <i>slot-number</i>	restart chassis-control < gracefully immediately soft >
show chassis alarms	request chassis fabric plane (offline online) <i>fabric-plane-number</i>	—
show chassis environment adc <adc-slot-number>	request chassis fpc (offline online restart) slot <i>fpc-slot-number</i>	—
show chassis environment cb <cb-slot-number>	request chassis fpm resync	—
show chassis environment fpc <fpc-slot-number>	request chassis mic (offline online) <i>fpc-slot fpc-slot-number mic-slot mic-slot-number</i>	—

<i>Show commands</i>	<i>Request commands</i>	<i>Restart commands</i>
show chassis environment fpm	request chassis routing-engine master (acquire release switch) <no-confirm>	—
show chassis environment monitored	request chassis sfb (offline online) slot <i>sfb-slot-number</i>	—
show chassis environment psm < <i>psm-slot-number</i> >	request chassis spmb restart slot <i>spmb-slot-number</i>	—
show chassis environment routing-engine < <i>routing-engine-slot-number</i> >	—	—
show chassis environment sfb < <i>sfb-slot-number</i> >	—	—
show chassis environment <adc cb fpc fpm monitored psm routing-engine sfb>	—	—
show chassis craft-interface	—	—
show chassis ethernet-switch <(errors statistics)>	—	—
show chassis fabric destinations <fpc <i>fpc-slot-number</i> >	—	—
show chassis fabric (destinations fpcs plane plane-location summary)	—	—
show chassis fan	—	—
show chassis firmware	—	—
show chassis fpc <slot> detail <detail <slot>> <pic-status <slot>> < <i>fpc-slot-number</i> >	—	—
show chassis hardware < (clei-models detail extensive models)>	—	—
show chassis in-service upgrade	—	—
show chassis mac-addresses	—	—
show chassis network-services	—	—
show chassis pic fpc-slot <i>fpc-slot-number</i> pic-slot <i>pic-slot-number</i>	—	—

Show commands	Request commands	Restart commands
show chassis power <sequence>	—	—
show chassis routing-engine <slot-number bios>	—	—
show chassis sfb <slot slot-number>	—	—
show chassis spmb	—	—
show chassis temperature-thresholds	—	—
show chassis zones <detail>	—	—

For details of all system management operational mode commands and the command options supported on the MX2010 router, see the *System Basics and Services Command Reference*.

[*System Basics and Services Command Reference*]

- **SNMP and MIB support for MX2010 routers**—Starting with Junos OS Release 12.3, the enterprise-specific Chassis Definitions for Router Model MIB, `jnx-chas-defines.mib`, is updated to include information about the new MX2010 routers. The Chassis Definitions for Router Model MIB contains the object identifiers (OIDs) used by the Chassis MIB to identify platform and chassis components of each router.

[See *jnxBoxAnatomy*, *Chassis Definitions for Router Model MIB*, and *MIB Objects for the MX2010 3D Universal Edge Router*.]

- **Improvements to Interface Transmit Statistics Reporting (MX Series devices)**—On MX Series devices, the logical interface-level statistics show only the offered load, which is often different from the actual transmitted load. To address this limitation, Junos OS introduces a new configuration option in Releases 11.4 R3 and 12.3 R1 and later. The new configuration option, **interface-transmit-statistics** at the `[edit interface interface-name]` hierarchy level, enables you to configure Junos OS to accurately capture and report the transmitted load on interfaces.

When the **interface-transmit-statistics** statement is included at the `[edit interface interface-name]` hierarchy level, the following operational mode commands report the actual transmitted load:

- `show interface interface-name <detail | extensive>`
- `monitor interface interface-name`
- `show snmp mib get objectID.ifIndex`



NOTE: This configuration is not supported on Enhanced IQ (IQE) and Enhanced IQ2 (IQ2E) PICs.

The **show interface *interface-name*** command also shows whether the interface-transmit-statistics configuration is enabled or disabled on the interface.

[See [Improvements to Interface Transmit Statistics Reporting](#).]

- **Extends support for encapsulating TDM signals as pseudowires for E1/T1 Circuit Emulation MIC (MX Series routers)**—Starting with Junos OS Release 12.3, the Channelized E1/T1 Circuit Emulation MIC (MIC-3D-16CHE1-T1-CE) supports encapsulating structured (NxDSO) time division multiplexed (TDM) signals as pseudowires over packet-switch networks (PSNs).

[See [Configuring SAToP Emulation on T1/E1 Interfaces on Circuit Emulation PICs](#).]

- **RE-JCS-1X2400-48G-S Routing Engine**—The JCS-1200 Control System now supports the RE-JCS-1X2400-48G-S Routing Engine. The RE-JCS-1X2400-48G-S Routing Engine requires the enhanced management module (model number MM-E-JCS-S). The RE-JCS-1X2400-48G-S Routing Engine provides a 2.4-GHz dual core Xeon processor, 48 GB of memory, and two 128 GB hot-pluggable solid state drives. The RE-JCS-1X2400-48G-S Routing Engine supports the same functionality as the other routing engines supported on the JCS-1200.

[See [JCS1200 Control System Hardware Guide](#).]

- **SFPP-10GE-ZR transceiver**—The following PICs on the T640, T1600, and T4000 routers now support the SFPP-10GE-ZR transceiver. The SFPP-10GE-ZR transceiver supports the 10GBASE-Z optical interface standard. For more information, see “Cables and connectors” in the PIC guide.

T640 Router:

- 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (Model number: PD-5-10XGE-SFPP)

T1600 Router:

- 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (Model number: PD-5-10XGE-SFPP)

T4000 Router:

- 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (Model number: PF-12XGE-SFPP)
- 10-Gigabit Ethernet LAN/WAN PIC with Oversubscription and SFP+ (Model numbers: PD-5-10XGE-SFPP for 10-Port Type 4 PIC and PF-24XGE-SFPP for 24-Port Type 5 PIC)

[See [10-Gigabit Ethernet 10GBASE Optical Interface Specifications](#), [T640 Core Router PIC Guide](#), [T1600 Core Router PIC Guide](#), and [T4000 Core Router PIC Guide](#).]

- **CFP-100GBASE-ER4 and CFP-100GBASE-SR10 Transceivers**—The following PICs on the T1600 and T4000 routers now support the CFP-100GBASE-ER4 and CFP-100GBASE-SR10 transceivers. The CFP-100GBASE-ER4 transceiver supports the 100GBASE-ER4 optical interface standard. The CFP-100GBASE-SR10 transceiver supports the 100GBASE-SR10 optical interface standard. For more information, see “Cables and connectors” in the PIC guide.
- **T1600 Router:** 100-Gigabit Ethernet PIC with CFP (Model number: PD-1CE-CFP-FPC4)

- **T4000 Router:** 100-Gigabit Ethernet PIC with CFP (Model numbers: PF-1CGE-CFP for Type 5 and PD-1CE-CFP-FPC4 for Type 4)

[See *100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications*, *T1600 Core Router PIC Guide*, and *T4000 Core Router PIC Guide*.]

- **Accounting of system statistics for IPv4 and IPv6 traffic**—On MX Series routers, you can enable accounting of system statistics for IPv4 and IPv6 traffic by including the **extended-statistics** statement at the **[edit chassis]** hierarchy level. By default, accounting of system statistics is disabled.

[See [extended-statistics](#).]

- **Fabric enhancements for Juniper Networks MX2020 and MX2010 routers**—Juniper Networks MX2020 and MX2010 routers now support all existing fabric hardening enhancements.
- **Support for unified in-service software upgrade (TX Matrix Plus router)**—Starting with Junos OS Release 12.3R2, unified in-service software upgrade (unified ISSU) is supported on a routing matrix based on a TX Matrix Plus router with the TXP-T1600 configuration.

Unified ISSU is a process to upgrade the system software with minimal disruption of transit traffic and no disruption on the control plane. In this process, the new system software version must be later than the previous system software version. When unified ISSU completes, the new system software state is identical to that of the system software when the system upgrade is performed by powering off the system and then powering it back on.

- Enhancement to **ping ethernet** command—Enables you to specify a multicast MAC address. For example:

```
user@host> ping ethernet maintenance-domain md3 maintenance-association ma3
01:80:c2:00:00:33
```

- **Symmetric Load Balancing on MX Series routers with MPCs**—Enables support for symmetrical load balancing over 802.3ad link aggregation groups (LAGs) on MX Series routers with MPCs. [See [Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers](#).]
- **Computation of the Layer 2 overhead attribute in interface statistics (MX Series routers)**—On MX Series routers, you can configure the physical interface and logical interface statistics to include the Layer 2 overhead size (header and trailer bytes) for both ingress and egress interfaces. Both the transit and total statistical information are computed and displayed for each logical interface. This functionality is supported on 1-Gigabit and 10-Gigabit Ethernet interfaces on Dense Port Concentrators (DPCs) and Modular Port Concentrators (MPCs).

You can enable the Layer 2 overhead bytes for computation in the logical interface statistics by configuring the **account-layer2-overhead (value | <ingress bytes | egress bytes>)** statement at the **[edit interface interface-name unit logical-unit-number]** hierarchy level. If you configure this capability, all the Layer 2 header details (L2 header and cyclic redundancy check [CRC]) based on the Layer 2 encapsulation configured for an interface are calculated and displayed in the physical and logical interface statistics for ingress and egress interfaces in the output of the **show interfaces**

interface-name commands. For physical and logical interfaces, the **Input bytes** and **Output bytes** fields under the Traffic statistics section in the output of the **show interfaces interface-name <detail | extensive>** command include the Layer 2 overhead of the packets. For physical and logical interfaces, the **Input Rate** and **Output Rate** fields under the Traffic statistics section in the output of the **show interfaces interface-name <media | statistics>** command include the Layer 2 overhead of the packets. For logical interfaces, the values for the newly added **Egress accounting overhead** and **Ingress accounting overhead** fields display the Layer 2 overhead size for transmitted and received packets respectively.

The ifInOctets and the ifOutOctets MIB objects display statistics that include Layer 2 overhead bytes if you configured the setting to account for Layer 2 overhead at the logical interface level.

- **New label-switching router (LSR) FPC (model number T4000-FPC5-LSR)**—The new LSR FPC in a T4000 core router provides LSR capability with the following scaling numbers:

Feature	LSR FPC Scale
RIB capacity	28 million
FIB (IPv4 and IPv6 unicast)	64,000
MPLS label push	48,000
MPLS label FIB/MPLS swap table	256,000
IP multicast route capacity	256,000
Multicast forwarding table (S, G)	128,000
RSVP LSPs	32,000 (ingress/egress)
	64,000 (transit)
Layer 2 VPN ingress/egress with family CCC/TCC	8000

The LSR FPC operates in the following modes:

- **Packet transport mode**—When the LSR FPC operates as an LSR only, the LSR FPC scaling numbers are supported.
- **Converged P/PE mode**—In a mixed provider (P) and provider edge (PE) router deployment, the LSR FPC might receive routes and next hops that exceed the LSR scaling numbers that are supported. In that case, extended scaling numbers, such as for the T4000-FPC5-3D, are supported.
- **PE router mode**—In PE router mode, running a Layer 3 VPN or peering services from the LSR FPC is not supported.

[See the [T4000 Core Router Hardware Guide](#).]

- **Support for new fixed-configuration MPC on MX240, MX480, MX960, and MX2020 routers**—MX2020, MX960, MX480, and MX240 routers support a new MPC, MPC4E. MPC4E provides scalability in bandwidth and services capabilities of the routers. MPC4E, like other MPCs, provides the connection between the customer's Ethernet interfaces and the routing fabric of the MX Series chassis. MPC4E is a fixed-configuration MPC and does not contain separate slots for Modular Interface Cards (MICs). MPC4E is available in two models: MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE.

MPC4E, like MPC3E, requires the Enhanced MX Switch Control Board (SCBE) for fabric redundancy. MPC4E does not support legacy SCBs. MPC4E interoperates with existing MX Series line cards, including Dense Port Concentrators (DPCs) and Modular Port Concentrators (MPCs).

MPC4E contains two Packet Forwarding Engines—**PFE0** hosts **PIC0** and **PIC1** while **PFE1** hosts **PIC2** and **PIC3**.

MPC4E supports:

- Forwarding capability of up to 130 Gbps per Packet Forwarding Engine. On MX240, MX480, and MX960 routers with Enhanced MX Switch Control Boards (SCBEs), each Packet Forwarding Engine can forward up to 117 Gbps because of the Packet Forwarding Engine and fabric limitations. On M2020 routers, each Packet Forwarding Engine can forward up to 130 Gbps.
- Both 10-Gigabit Ethernet interfaces and 100-Gigabit Ethernet interfaces.
- Small form-factor pluggable (SFP) and C form-factor pluggable (CFP) transceivers for connectivity.
- Up to 240 Gbps of full-duplex traffic.
- Intelligent oversubscription services.
- WAN-PHY mode on 10-Gigabit Ethernet Interfaces on a per-port basis.
- Up to four full-duplex tunnel interfaces on each MPC4E.
- Effective line rate of 200 Gbps for packets larger than 300 bytes.

MPC4E supports feature parity with the Junos OS Release 12.3 software features:

- Basic Layer 2 features and virtual private LAN service (VPLS) functionality, including Operation, Administration, and Maintenance (OAM)
- Class-of-service (CoS) support
- Firewall filters and policers
- Interoperability with existing DPCs and MPCs
- Internet Group Management Protocol (IGMP) snooping with bridging, integrated routing and bridging (IRB), and VPLS
- Layer 3 routing protocols
- J-Flow monitoring and services
- MPLS
- Multicast forwarding

- Precision Time Protocol (IEEE 1588)
- Tunnel Interfaces support

The following features are not supported on the MPC4E:

- Fine-grained queuing and input queuing
- Intelligent hierarchical policers
- Layer 2 trunk port
- MPLS fast reroute (FRR) VPLS instance prioritization
- Multilink services
- Virtual Chassis support

For more information about the supported and unsupported Junos OS software features for this MPC, see *Protocols and Applications Supported by the MX240, MX480, MX960, and MX2000 MPC4E* in the *MX Series Line Card Guide*.

- **Support for Ethernet synthetic loss measurement**—You can trigger on-demand and proactive Operations, Administration, and Maintenance (OAM) for measurement of statistical counter values corresponding to ingress and egress synthetic frames. Frame loss is calculated using synthetic frames instead of data traffic. These counters maintain a count of transmitted and received synthetic frames and frame loss between a pair of maintenance association end points (MEPs).

The Junos OS implementation of Ethernet synthetic loss measurement (ETH-SLM) is fully compliant with the ITU-T Recommendation Y.1731. Junos OS maintains various counters for ETH-SLM PDUs, which can be retrieved at any time for sessions that are initiated by a certain MEP. You can clear all the ETH-SLM statistics and PDU counters.

The ETH-SLM feature provides the option to perform ETH-SLM for a given 802.1p priority; to set the size of the ETM-SLM protocol data unit (PDU); and to generate XML output.

You can perform ETH-SLM in on-demand ETH-SLM mode (triggered through the CLI) or in proactive ETH-SLM mode (triggered by the iterator application). To trigger synthetic frame loss measurement (on-demand mode) and provide a run-time display of the measurement values, use the **monitor ethernet synthetic-loss-measurement (remote-mac-address | mep mep-id) maintenance-domain md-name maintenance-association ma-name count frame-count wait time priority 802.1p value size xml** operational mode command.

To display the archived on-demand synthetic frame loss measurement values, use the **show oam ethernet connectivity-fault-management synthetic-loss-statistics maintenance-domain md-name maintenance-association ma-name local-mep local-mep-id remote-mep remote-mep-id count entry-count** operational mode command. To display the cumulative on-demand synthetic frame loss measurement values, use the **show oam ethernet connectivity-fault-management interfaces detail** operational mode command.

To perform proactive ETH-SLM, you need to create an SLA iterator profile and associate the profile with a remote MEP. To create an SLA iterator profile for ETH-SLM, include

the **measurement-type slm** statement at the **[edit protocols oam ethernet connectivity-fault-management performance-monitoring sla-iterator-profiles profile-name]** hierarchy level. To display proactive synthetic loss measurement values, use the **show oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-domain md-name maintenance-association ma-name local-mep local-mep-id remote-mep remote-ip-id sla-iterator identifier** operational mode command.

You can reset the SLM statistics by clearing the currently measured ETH-SLM statistical counters. To clear the existing on-demand Ethernet loss statistics measured for a specific maintenance domain and maintenance association local and remote MEP and restart the counter, use the **clear oam ethernet connectivity-fault-management synthetic-loss-measurement maintenance-domain md-name maintenance-association ma-name local-mep local-mep remote-mep remote-mep** operational mode command. To clear the existing proactive ETH-SLM counters for a specific maintenance domain, maintenance association, local MEP, remote MEP, and an SLA iterator, use the **clear oam ethernet connectivity-fault-management sla-iterator-statistics maintenance-domain md-name maintenance-association ma-name local-mep local-mep-id remote-mep remote-mep-id sla-iterator identifier** operational mode command.

The following list consists of the connectivity fault management (CFM)-related operational mode commands that display ETH-SLM statistics:

- The **show oam ethernet connectivity-fault-management interfaces detail** command is enhanced to display on-demand ETH-SLM statistics for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.
- The **show oam ethernet connectivity-fault-management mep-statistics** command is enhanced to display on-demand ETH-SLM statistics and frame counts for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.
- The **show oam ethernet connectivity-fault-management mep-database** command is enhanced to display on-demand ETH-SLM frame counters for MEPs in the specified CFM maintenance association within the specified CFM maintenance domain.
- The **show oam ethernet connectivity-fault-management sla-iterator-statistics** command is enhanced to display service-level agreement (SLA) iterator statistics for ETH-SLM.

[Release Notes]

- **Support for OSS mapping to represent a T4000 chassis as a T1600 or a T640 chassis (T4000 routers)**—Starting with Junos OS Release 12.3R3, you can map a T4000 chassis to a T1600 chassis or a T640 chassis, so that the T4000 chassis is represented as a T1600 chassis or a T640 chassis, respectively, without changing the operations support systems (OSS) qualification. Therefore, you can avoid changes to the OSS when a T1600 chassis or a T640 chassis is upgraded to a T4000 chassis.

You can configure the OSS mapping feature with the **set oss-map model-name t640|t1600** configuration command at the **[edit chassis]** hierarchy level. This command changes the **chassis** field to the **known chassis** field in the output of the **show chassis hardware** and the **show chassis oss-map** operational mode commands. You can verify

the change with the **show snmp mib walk system** and **show snmp mib walk jnxBoxAnatomy** operational commands as well.

You can delete the OSS mapping feature by using the **delete chassis oss-map model-name t640|t1600** configuration command.

- **Enhanced load balancing for MIC and MPC interfaces (MX Series)** — Starting with Junos OS Release 12.3R4, the following load-balancing solutions are supported on an aggregated Ethernet bundle to correct genuine traffic imbalance among the member links:
 - Adaptive — Uses a real-time feedback and control mechanism to monitor and manage traffic imbalances.
 - Per-packet random spray — Randomly sprays the packets to the aggregate next hops to ensure that the next hops are equally loaded resulting in packet reordering.

The aggregated Ethernet load-balancing solutions are mutually exclusive. To configure these solutions, include the **adaptive** or **per-packet** statement at the **[edit interfaces aex aggregated-ether-options load-balance]** hierarchy level.

- **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 module supports 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 Networks 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 below. 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](#).]

Junos OS XML API and Scripting

- **Support for service template automation**—Starting with Junos OS Release 12.3, you can use service template automation to provision services such as VPLS VLAN, Layer 2 and Layer 3 VPNs, and IPsec across similar platforms running Junos OS. Service template automation uses the **service-builder.slax** op script to transform a user-defined service template definition into a uniform API, which you can then use to configure and provision services on similar platforms running Junos OS. This permits you to create a service template on one device, generalize the parameters, and then quickly and uniformly provision that service on other devices. This decreases the time required to configure the same service on multiple devices, and reduces configuration errors associated with manually configuring each device.

[See [Service Template Automation](#).]

- **Support for configuring limits on concurrently running event policies and memory allocation for scripts**—Junos OS Release 12.3 supports configuring limits on the maximum number of concurrently running event policies and the maximum amount of memory allocated for the data segment for scripts of a given type. By default, the maximum number of event policies that can run concurrently in the system is 15, and the maximum amount of memory allocated for the data segment portion of an executed script is half of the total available memory of the system, up to a maximum value of 128 MB.

To set the maximum number of event policies that can run concurrently on a device, configure the **max-policies** *policies* statement at the **[edit event-options]** hierarchy level. You can configure a maximum of 0 through 20 policies. To set the maximum memory allocated to the data segment for scripts of a given type, configure the **max-datasize** *size* statement under the hierarchy appropriate for that script type, where *size* is the memory in bytes. To specify the memory in kilobytes, megabytes, or gigabytes, append **k**, **m**, or **g**, respectively, to the size. You can configure the memory in the range from 2,3068,672 bytes (22 MB) through 1,073,741,824 bytes (1 GB).

[See [Configuring Limits on Executed Event Policies and Memory Allocation for Scripts.](#)]

Layer 2 Features

- **Support for Synchronous Ethernet and Precision Time Protocol on MX Series routers with 16-port Channelized E1/T1 Circuit Emulation MIC (MX Series Routers)**—Starting with Junos OS Release 12.3, Synchronous Ethernet and Precision Time Protocol (PTP) are supported on MX Series routers with the 16-port Channelized E1/T1 Circuit Emulation MIC (MIC-3D-CH-16E1T1-CE). The clock derived by Synchronous Ethernet, PTP, or an internal oscillator is used to drive the T1/E1 interfaces on the 16-port Channelized E1/T1 Circuit Emulation MIC.
- **E-TREE with remote VSI support on Juniper NSN Carrier Ethernet Transport**—The Juniper NSN Carrier Ethernet Transport solution supports Metro Ethernet Forum (MEF) Ethernet Tree (E-TREE) services using centralized virtual switch instances (VSIs). E-TREE is a rooted multipoint service, where end points are classified as Roots and Leaves. Root end points can communicate with both Root and Leaf end points, but Leaf end points can only communicate with the Root end points.

Juniper's NSN CET solution employs E-TREE services using a centralized VSI model. This means that VSIs are only provisioned on certain selected PEs. End points are connected to these central VSIs using spoke pseudowires. The centralized VSI model uses a lower number of pseudowires and less bandwidth than the distributed VSI model.

- **Adds support to send and receive untagged RSTP BPDUs on Ethernet interfaces (MX Series platforms)**—VLAN Spanning Tree Protocol (VSTP) can send and receive untagged Rapid Spanning Tree Protocol (RSTP) bridge protocol data units (BPDUs) on Gigabit Ethernet (ge), 10 Gigabit Ethernet (xe), and aggregated Ethernet (ae) interfaces.

To configure this feature, include the **access-trunk** statement at the following hierarchy levels:

```
[edit protocols vstp vlan vlan-identifier interface interface-name]
[edit routing-instances routing-instance-name instance-type (layer2-control |
virtual-switch)]
[edit logical-systems logical-system-name protocols vstp]
[edit logical-systems logical-system-name routing-instances routing-instance-name
protocols vstp]
```

[See [access-trunk.](#)]

- **Extends support for multilink-based protocols on T4000 and TX Matrix Plus routers**—Starting with Junos OS Release 12.3R3, multilink-based protocols are supported on the T4000 and TX Matrix Plus routers with Multiservices PICs.
 - Multilink Point-to-Point Protocol (MLPPP)—Supports Priority-based Flow Control (PFC) for data packets and Link Control Protocol (LCP) for control packets. Compressed Real-Time Transport Protocol (CRTP) and Multiclass MLPPP are supported for both data and control packets.
 - Multilink Frame Relay (MLFR) end-to-end (FRF.15)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
 - Multilink Frame Relay (MFR) UNI NNI (FRF.16)—Supports Ethernet Local Management Interface (LMI), Consortium LMI (C-LMI), and Link Integrity Protocol (LIP) for data and control packets.
 - Link fragmentation and interleaving (LFI) non multilink MLPPP and MLFR packets.
 - Communications Assistance for Law Enforcement Act (CALEA)--Defines electronic surveillance guidelines for telecommunications companies.
 - Two-Way Active Measurement Protocol (TWAMP)-- Adds two-way or round-trip measurement capabilities

[Interfaces Command Reference]

- **Extends support of IPv6 statistics for MLPPP bundles on T4000 and TX Matrix Plus routers**—Starting with Junos OS Release 12.3R3, the **show interfaces lsq-fpc/pic/port** command displays the packet and byte counters for IPv6 data for Multilink Point-to-Point Protocol (MLPPP) bundles on link services intelligent queuing (LSQ) interfaces.

[Interfaces Command Reference]

- **Link Layer Discovery Protocol (LLDP) support (MX240, MX480, and MX960)**—You can configure the LLDP protocol on MX Series routers with MPC3E and MPC4E. To configure and adjust default parameters, include the **lldp** statement at the **[edit protocols]** hierarchy level.

LLDP is disabled by default. At the **[edit protocols lldp]** hierarchy level, use the **enable** statement to enable LLDP, and the **interfaces** statement to enable LLDP on all or some interfaces. Use the following statements at the **[edit protocols lldp]** hierarchy level to configure or adjust the default LLDP parameters:

- **advertisement-interval**
- **transmit-delay**
- **hold-multiplier**
- **ptopo-configuration-trap-interval**
- **ptopo-configuration-maximum-hold-time**
- **lldp-configuration-notification-interval.**

- **Configuration support for manual and automatic link switchover mechanism on multichassis link aggregation interface**—You can configure a multichassis link aggregation (MC-LAG) interface in active-standby mode to automatically revert to a preferred node. In an MC-LAG topology with active-standby mode, a link switchover happens only if the active node goes down. With this configuration, you can trigger a link switchover to a preferred node even when the active node is available.

To enable automatic link switchover for an multichassis link aggregation (mc-ae) interface, you must configure the **switchover-mode revertive** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. You can also specify the revert time for the switchover by using the **revert-time** statement. To continue using the manual switchover mechanism, you must configure the **switchover-mode non-revertive** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. For nonrevertive mode, you can configure manual switchover to the preferred node by using the **switchover immediate** and **mc-ae-id** statements at the **[request interface mc-ae]** hierarchy level.

With this feature, you can use the **show interfaces mc-ae revertive-info** command to view the switchover configuration information.

- **Uniform Enhanced Layer 2 Software CLI configuration statements and operational commands**—Enhanced Layer 2 Software (ELS) provides a uniform CLI for configuring and monitoring Layer 2 features on MX Series routers in LAN mode (MX-ELM). With ELS, for example, you can configure a VLAN and other Layer 2 features on an MX-ELM router by using the same configuration commands.

[See the ELS CLI documentation for MX Series routers: [Junos OS for EX9200 Switches, Release 12.3.](#)]

- When changing modes the user must delete any unsupported configurations.

[See [Configuring MX Enhanced LAN Mode.](#)]

- The web-based ELS Translator tool is available for registered customers to help them become familiar with the ELS CLI and to quickly translate existing MX Series router CLI configurations into ELS CLI configurations.

[See [ELS Translator.](#)]

Layer 2 Tunneling Protocol

- **Support for filtering trace results by subscriber or domain for AAA, L2TP, and PPP (MX Series routers)**—You can now filter trace results for AAA (authd), L2TP (l2tpd), and PPP (pppd) by subscribers or domains. Specify the **filter user username** option at the appropriate hierarchy level:

- AAA—**[edit system processes general-authentication-service traceoptions filter]**
- L2TP—**[edit services l2tp traceoptions filter]**
- PPP—**[edit protocols ppp-service traceoptions filter]**

For subscriber usernames that have the expected form of **user@domain**, you can filter on either the user or the domain. The filter supports the use of a wildcard (*) at the beginning or end of the user, the domain, or both. For example, the following are all

acceptable uses of the wildcard: tom@example.com, tom*, *tom, *ample.com, tom@ex*, tom*.*example.com.

You cannot filter results using a wildcard in the middle of the user or domain. For example, the following uses of the wildcard are not supported: tom*25@example.com, tom125@ex*.com.

When you enable filtering by username, traces that have insufficient information to determine the username are automatically excluded.

MPLS

- **Link protection for MLDP**—MLDP link protection enables 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.

To configure MLDP link protection, use the **make-before-break**, and **link-protection-timeout** statements at the **[edit protocols ldp]** hierarchy level. To view MBB capabilities, use the **show ldp session detail** command. To verify that link protection is active, use the **show ldp interface extensive** command. To view adjacency type, use the **show ldp neighbor extensive** command. To view MBB interval, use the **show ldp overview** command.

[See [Example: Configuring LDP Link Protection](#).]

- A new ultimate-hop popping feature is now available for LSPs configured on M Series, MX Series, and T Series platforms. An ultimate-hop popping LSP pops the MPLS label at the LSP egress. The default behavior for an LSP on a Juniper Networks device is to pop the MPLS label at the penultimate-hop router (the router before the egress router). Ultimate-hop popping is available on RSVP-signaled LSPs and static LSPs.

The following network applications could require that you configure UHP LSPs:

- MPLS-TP for performance monitoring and in-band OAM
- Edge protection virtual circuits
- UHP static LSPs

To enable ultimate-hop popping on an LSP, include the **ultimate-hop-popping** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level to enable ultimate-hop popping on a specific LSP or at the **[edit protocols mpls]** hierarchy level to enable ultimate-hop popping on all of the ingress LSPs configured on the router. When you enable ultimate-hop popping, RSVP attempts to resignal existing LSPs as ultimate-hop popping LSPs in a make-before-break fashion. If an egress router does not support ultimate-hop popping, the existing LSP is torn down. If you disable

ultimate-hop popping, RSVP resignals existing LSPs as penultimate-hop popping LSPs in a make-before-break fashion.

[See [Configuring Ultimate-Hop Popping for LSPs](#).]

- **Enable local receivers on the ingress of a point-to-multipoint circuit cross-connect (CCC)**—This feature enables you to switch the traffic entering a P2MP LSP to local interfaces. On the ingress PE router, CCC can be used to switch an incoming CCC interface to one or more outgoing CCC interfaces. To configure the output interface, include the **output-interface** statement at the **[edit protocols connections p2mp-transmit-switch <p2mp-lsp-name-on-which-to-transmit>]** hierarchy level. One or more output interfaces can be configured as local receivers on the ingress PE router using this statement. Use the **show connections p2mp-transmit-switch (extensive | history | status)**, **show route ccc <interface-name> (detail | extensive)**, and **show route forwarding-table ccc <interface-name> (detail | extensive)** commands to view details of the local receiving interfaces at ingress. *[MPLS]*
- **Support for Bidirectional Forwarding Detection protocol, LSP traceroute, and LSP ping on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX Series Routers)**—Starting with Junos OS 12.3, support for Bidirectional Forwarding Detection (BFD) protocol, LSP traceroute, and LSP ping is extended to Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE).

The BFD protocol is a simple hello mechanism that detects failures in a network. You can configure Bidirectional Forwarding Detection (BFD) for LDP LSPs. You can also use the LSP ping commands to detect LSP data plane faults. You can trace the route followed by an LDP-sigaled LSP.

LDP LSP traceroute is based on RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. This feature allows you to periodically trace all paths in a Forwarding Equivalence Class (FEC). The FEC topology information is stored in a database accessible from the CLI.

- **Host fast reroute (HFRR)**—Adds a precomputed protection path into the Packet Forwarding Engine, such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the Packet Forwarding Engine can use another path without having to wait for the router or the protocols to provide updated forwarding information. HFRR is a technology that protects IP endpoints on multipoint interfaces, such as Ethernet. This technology is important in data centers where fast service restoration for server endpoints is critical. After an interface or a link goes down, HFRR enables the local repair time to be approximately 50 milliseconds. You can configure HFRR by adding the **link-protection** statement to the interface configuration in the routing instance. We recommend that you include this statement on all provider edge (PE) devices that are connected to server farms through multipoint interfaces.

[See [Example: Configuring Host Fast Reroute](#).]

- **Support of Path Computation Element Protocol for RSVP-TE**—Starting with Junos OS Release 12.3, the MPLS RSVP-TE functionality is extended to provide a partial client-side implementation of the stateful Path Computation Element (PCE) architecture (draft-ietf-pce-stateful-pce). The PCE computes path for the traffic engineered LSPs (TE LSPs) of ingress routers that have been configured for external control. The ingress router that connects to a PCE is called a Path Computation Client

(PCC). The PCC is configured with the Path Computation Client Protocol (PCEP) (defined in RFC 5440, but limited to the functionality supported on a stateful PCE only) to facilitate external path computing by a PCE.

In this new functionality, the active stateful PCE sets parameters for the PCC's TE LSPs, such as bandwidth, path (ERO), and priority. The TE LSP parameters configured from the PCC's CLI are overridden by the PCE-provided parameters. The PCC re-signals the TE LSPs based on the path specified by the PCE. Since the PCE has a global view of the bandwidth demand in the network and performs external path computations after looking up the traffic engineering database, this feature provides a mechanism for offline control of TE LSPs in the MPLS RSVP TE enabled network.

To enable external path computing by a PCE, include the **lsp-external-controller** statement on the PCC at the **[edit mpls]** and **[edit mpls lsp lsp-name]** hierarchy levels. To enable PCE to PCC communication, configure **pcep** on the PCC at the **[edit protocols]** hierarchy level.

[See [PCEP Configuration Guide](#).]

Multicast

- **Redundant virtual tunnel (VT) interfaces in Multiprotocol BGP (MBGP) multicast VPNs (MVPNs)**—VT interfaces are needed for multicast traffic on routing devices that function as combined provider edge (PE) and provider core (P) routers to optimize bandwidth usage on core links. VT interfaces prevent traffic replication when a P router also acts as a PE router (an exit point for multicast traffic). You can configure up to eight VT interfaces in a routing instance, thus providing Tunnel PIC redundancy inside the same multicast VPN routing instance. When the active VT interface fails, the secondary one takes over, and you can continue managing multicast traffic with no duplication. To configure, include multiple VT interfaces in the routing instance and, optionally, apply the **primary** statement to one of the VT interfaces.

[See [Example: Configuring Redundant Virtual Tunnel Interfaces in MBGP MVPNs](#).]

- **Enhancements to RPD_MC_OIF_REJECT and RPD_MC_OIF_RE_ADMIT system log messages**—When multicast call admission control (CAC) is enabled on an interface, the routing software cannot add a multicast flow to that interface if doing so exceeds the maximum configured bandwidth for that interface. Consequently, the interface is rejected for that flow due to insufficient bandwidth, and the router writes the RPD_MC_OIF_REJECT system log message to the log file at the **info** severity level. When bandwidth again becomes available on the interface, interfaces previously rejected for a flow are readmitted. In that case, the router writes the RPD_MC_OIF_RE_ADMIT system log message to the log file at the **info** severity level.

Both the RPD_MC_OIF_REJECT system log message and the RPD_MC_OIF_RE_ADMIT system log message include the **interface-name**. In RPD_MC_OIF_REJECT messages, **interface-name** identifies the interface that was rejected for a multicast flow due to insufficient bandwidth. In RPD_MC_OIF_RE_ADMIT messages, **interface-name** identifies the interface that was re-admitted for a multicast flow due to newly available bandwidth on the interface.

The RPD_MC_OIF_REJECT and RPD_MC_OIF_RE_ADMIT system log messages have been enhanced in this release to include the following information in addition to the *interface-name*:

- **group-address**—IP address of the multicast group
- **source-address**—Source IP address of the multicast flow
- **flow-rate**—Bandwidth of the multicast flow, in bits per second (bps)
- **maximum-flow-rate**—Maximum bandwidth that is the sum of all multicast flows on the interface, in bps
- **admitted-flow-rate**—Admitted bandwidth that is the sum of all multicast flows on the interface, in bps

When the maximum allowable bandwidth is exceeded on the logical interface (also known as the *map-to interface*) to which an outgoing interface (OIF) map directs (maps) multicast traffic, the RPD_MC_OIF_REJECT and RPD_MC_OIF_RE_ADMIT system log messages also display the *oif-map-interface-name* string. The *oif-map-interface-name* is an optional string that identifies one or more subscriber interfaces that requested the multicast traffic and are associated with the OIF map.

The *oif-map-interface-name* string appears in the RPD_MC_OIF_REJECT and RPD_MC_OIF_RE_ADMIT system log messages when all of the following conditions are met:

- The subscriber interface has CAC enabled, and is associated with an OIF map.
- The map-to interface (also known as the multicast VLAN, or M-VLAN) has CAC enabled.
- The subscriber interface that maps traffic to the M-VLAN receives an IGMP or MLD join message.
- The M-VLAN is not already sending the group and source of the multicast flow.
- Adding a multicast flow to the M-VLAN exceeds the maximum bandwidth configured on the M-VLAN interface.
- The group and source is source-specific multicast (SSM), or multicast data traffic is flowing.

Being able to view all of this information in a single system log message makes it easier for you to identify, troubleshoot, and resolve problems when using multicast protocols in your network. In earlier Junos OS releases, the RPD_MC_OIF_REJECT and RPD_MC_OIF_RE_ADMIT system log messages included only *interface-name*, but did not include *group-address*, *source-address*, or *oif-map-interface-name*.

The following example shows an RPD_MC_OIF_REJECT system log message for an oversubscribed interface. Because the interface is not configured with an OIF map, the *oif-map-interface-name* string does not appear.

```
Oct 26 08:09:51 wfpro-mx1-c r1:rp[12955]: RPD_MC_OIF_REJECT: 225.1.0.2 193.0.1.2
(5000000 bps) rejected due to lack of bandwidth on ge-4/1/0.1 (maximum 12000000
bps, admitted 10000000 bps)
```

This example includes the following information:

- The **group-address** is 225.1.0.2
- The **source-address** is 193.0.1.2
- The **flow-rate** is 5000000 bps
- The **interface-name** is ge-4/1/0.1
- The **maximum-flow-rate** is 12000000 bps
- The **admitted-flow-rate** is 10000000 bps

The following example shows the same RPD_MC_OIF_REJECT system log message for an interface configured with an OIF map. All of the values are the same as in the preceding RPD_MC_OIF_REJECT example except for the addition of the **oif-map-interface-name** string, which is requested from ge-4/1/0.1 ge-4/1/0.2 ge-4/1/0.3.

Oct 26 08:17:05 wfpro-mx1-c r1:rpdd[15133]: RPD_MC_OIF_REJECT: 225.1.0.2 193.0.1.2 (5000000 bps) rejected due to lack of bandwidth on ge-4/1/0.4 (maximum 12000000 bps, admitted 10000000 bps) requested from ge-4/1/0.1 ge-4/1/0.2 ge-4/1/0.3

The enhancements to the RPD_MC_OIF_REJECT and RPD_MC_OIF_RE_ADMIT system log messages make no changes to how bandwidth is managed on the router for multicast configurations.

[*Multicast Protocols Configuration Guide*]

- **Static ARP with multicast MAC address for an IRB interface**—Enables you to configure a static ARP entry with a multicast MAC address for an IRB interface which acts as the gateway to the network load balancing (NLB) servers. Earlier, the NLB servers dropped packets with a unicast IP address and a multicast MAC address. Junos OS 12.3 supports the configuration of a static ARP with a multicast MAC address.

To configure a static ARP entry with a multicast MAC address for an IRB interface, configure the ARP entry at the **[edit interfaces irb unit logical-unit-number family inet address address]** hierarchy level.

```
irb {
  unit logical-unit-number {
    family inet {
      address address {
        arp address multicast-mac mac-add;
      }
    }
  }
}
```

Power Management

- **Power management support on T4000 routers with six-input DC power supply**
—Starting with Junos OS Release 12.3, the power management feature is enabled on a Juniper Networks T4000 Core Router. This feature enables you to limit the overall chassis output power consumption. That is, power management enables you to limit the router from powering on a Flexible PIC Concentrator (FPC) when sufficient output power is not available to power on the FPC.

The power management feature is enabled only when six input feeds with 40 amperes (A) each or four input feeds with 60 A each are configured on the router. The power management feature is *not* enabled for any other input feed–current combination. When the power management feature is *not* enabled, Junos OS tries to power on all the FPCs connected to the router.



CAUTION: If you do not configure the power management feature and the maximum power draw is exceeded by the router, FPCs' states might change from Online to Offline or Present, some traffic might drop, or the interfaces might flap.

After you connect the input feeds to the router, you must configure the number of input feeds connected to the router and the amount of current received at the input feeds. Use the **feeds** statement and the **input current** statement at the **[edit chassis pem]** hierarchy level to configure the number of input feeds and the amount of current received at the input feeds, respectively.



NOTE: You can connect three 80 A DC power cables to the six-input DC power supply by using terminal jumpers. When you do this, ensure that you set the value of feeds statement to 6 and that of the input current statement to 40. If these configurations are not set, the power management feature is *not* enabled and, therefore, Junos OS tries to power on all the FPCs connected to the router.

When the power management feature is enabled, FPCs connected to the router are powered on based on the power received by the router. If the router receives sufficient power to power on all the FPCs connected to the router, all the FPCs are powered on. If sufficient power is not available, Junos OS limits the number of FPCs brought online. That is, Junos OS uses the total available chassis output power as a factor to decide whether or not to power on an FPC connected to the router.

[See [T4000 Power Management Overview](#) and [T4000 Core Router Hardware Guide](#).]

Routing Policy and Firewall Filters

- **Source checking for forwarding filter tables**—On MX Series 3D Universal Edge Routers, you can apply a forwarding table filter by using the **source-checking** statement at the **[edit forwarding-options family inet6]** hierarchy level. This discards IPv6 packets when the source address type is unspecified, loopback, multicast, or link-local. RFC 4291, *IP*

Version 6 Addressing Architecture, refers to four address types that require special treatment when they are used as source addresses. The four address types are: Unspecified, Loopback, Multicast, and Link-Local Unicast. The loopback and multicast addresses must never be used as a source address in IPv6 packets. The unspecified and link-local addresses can be used as source addresses but routers must never forward packets that have these addresses as source addresses. Typically, packets that contain unspecified or link-local addresses as source addresses are delivered to the local host. If the destination is not the local host, then the packet must not be forwarded. Configuring this statement filters or discards IPv6 packets of these four address types.

[See [Applying Filters to Forwarding Tables](#).]

- **Unidirectional GRE tunnels across IPv4 without tunnel interfaces**—For Junos OS Release 12.3R2 and later, you can configure a tunnel that transports IPv4, IPv6, protocol-independent, or MPLS traffic across an IPv4 network without having to create tunnel interfaces on services PICs. This type of GRE tunnel is unidirectional and transports unicast or multicast transit traffic as clear text. Encapsulation, de-encapsulation, and forwarding of payloads is executed by Packet Forwarding Engine processes for logical Ethernet interfaces or aggregated Ethernet interfaces hosted on MICs and MPCs in MX Series routers. Two MX Series routers installed as PE routers provide network connectivity to two CE routers that lack a native routing path between them. This feature is also supported in logical systems.

Specify tunnel characteristics by configuring the **tunnel-end-point** statement on the ingress PE router:

```
firewall {  
  tunnel-end-point tunnel-name {  
    ipv4 {  
      source-address source-host-address;  
      destination-address destination-host-address;  
    }  
    gre [key number];  
  }  
}
```

To configure the ingress PE router to encapsulate passenger protocol packets, attach a passenger protocol family firewall filter at the input of a supported interface. The following terminating firewall filter action refers to the specified tunnel and initiates encapsulation of matched packets:

```
encapsulate tunnel-name
```

To configure the egress PE router to de-encapsulate GRE packets and forward the original passenger protocol packets, attach an IPv4 firewall filter at the input of all interfaces that are advertised addresses for the router. The following terminating firewall filter action initiates de-encapsulation of matched packets:

```
decapsulate [routing-instance instance-name]
```

By default, the Packet Forwarding Engine uses the default routing instance to forward payload packets to the destination network. If the payload is MPLS, the Packet Forwarding Engine performs route lookup on the MPLS path routing table using the route label in the MPLS header.

If you specify the **decapsulate** action with an optional routing instance name, the Packet Forwarding Engine performs route lookup on the routing instance, and the instance must be configured.

[*Firewall Filters Configuration Guide*]

Routing Protocols

- **Expanded support for advertising multiple paths to a destination in BGP**—This feature now supports graceful restart and additional address families. Previously, graceful restart was not supported and only the IPv4 address family was supported with the BGP **add-path** feature. Now the following address families are supported:

- IPv4 unicast (**net unicast**)
- IPv6 unicast (**inet6 unicast**)
- IPv4 labeled unicast (**inet labeled-unicast**)
- IPv6 labeled unicast (**inet6 labeled-unicast**)

To configure these address families, include the **family <address-family> add-path** statement at the [**edit protocols bgp**] hierarchy level.

To configure graceful restart, include the **graceful-restart** statement at the [**edit routing-options**] hierarchy level.

[See [Example: Advertising Multiple BGP Paths to a Destination](#).]

- **Support for multihop BFD session**— One desirable application of BFD is to detect connectivity to routing devices that span multiple network hops and follow unpredictable paths. This is known as a multihop session. Until Junos OS Release 12.3, multihop BFD was non-distributed and ran on the Routing Engine. Starting in Junos OS 12.3, multihop BFD is distributed, meaning that it runs on the Packet Forwarding Engine. This change provides multiple scalability improvements.

Security

- **DDoS Protection Flow Detection (MX Series Routers)**—Flow detection is an enhancement to DDoS protection that supplements the DDoS policer hierarchies. When you enable flow detection by including the **flow-detection** statement at the **[edit system ddos-protection global]** hierarchy level, a limited amount of hardware resources are used to monitor the arrival rate of host-bound flows of control traffic. This behavior makes flow detection highly scalable compared to filter policers, which track all flows and therefore consume a considerable amount of resources.

Flows that violate a DDoS protection policer are tracked as suspicious flows; they become culprit flows when they violate the policer bandwidth for the duration of a configurable detection period. Culprit flows are dropped, kept, or policed to below the allowed bandwidth level. Suspicious flow tracking stops if the violation stops before the detection period expires.

Most flow detection attributes are configured at the packet level or flow aggregation level. [Table 2 on page 142](#) lists these statements, which you can include at the **[edit system ddos-protection protocols protocol-group packet-type]** hierarchy level. You can disable flow detection, configure the action taken for culprit flows, specify a bandwidth different than the policer bandwidth, configure flows to be monitored even when a policer is not in violation, disable automatic event reporting, or enable a timeout period that automatically removes flows as culprit flows after the timeout has expired.

Table 2: Flow Detection Packet-Level Statements

flow-detection-mode	flow-level-detection	no-flow-logging
flow-detect-time	flow-recover-time	physical-interface
flow-level-bandwidth	flow-timeout-time	subscriber
flow-level-control	logical-interface	timeout-active-flows

By default, flow detection automatically generates reports for events associated with the identification and tracking of culprit flows and bandwidth violations. You can include the **flow-report-rate** and **violation-report-rate** statements at the **[edit system ddos-protection global]** hierarchy level to configure the event reporting rate.

Use the **show ddos-protection protocols flow-detection** command to display flow detection information for all protocol groups or for a particular protocol group. Use the **show ddos-protection protocols culprit-flows** command to display information about culprit flows for all packet types, including the number of culprit flows discovered, the protocol group and packet type, the interface on which the flow arrived, and the source address for the flow. The **show ddos-protection statistics** command now provides a global count of discovered and currently tracked culprit flows. You can use the **clear ddos-protection protocols culprit-flows** command to clear all culprit flows, or just those for a protocol group or individual packet type.

[DDoS Configuration]

Subscriber Access Management

- **Support for PPP subscriber services over ATM networks (MX Series routers with MPCs and ATM MICs with SFP)**—Enables you to create PPP-over-ATM (PPPoA) configurations on an MX Series router that has an ATM MIC with SFP (model number MIC-3D-80C3-20C12-ATM) and a supported MPC installed. PPPoA configurations support statically created PPP logical subscriber interfaces over static ATM underlying interfaces. (Dynamic creation of the PPP interfaces is not supported.) Most features supported for PPPoE configurations are also supported for PPPoA configurations on an MX Series router. You can dynamically apply subscriber services such as CoS and firewall filters to the static PPP logical subscriber interface by configuring the services in the dynamic profile that creates the PPP logical interface.

PPPoA configurations on an MX Series router support two types of encapsulation on the ATM underlying interface:

- To configure PPPoA encapsulation that uses LLC, you must configure the ATM underlying interface with PPP-over-AAL5 LLC encapsulation. To do so, include the **encapsulation atm-ppp-llc** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level.
- To configure PPPoA encapsulation that uses VC multiplexing, you must configure the ATM underlying interface with PPP-over-ATM AAL5 multiplex encapsulation. To do so, include the **encapsulation atm-ppp-vc-mux** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level.

PPPoA configurations enable the delivery of subscriber-based services, such as CoS and firewall filters, for PPP subscribers accessing the router over an ATM network. You use the same basic statements, commands, and procedures to create, verify, and manage PPPoA configurations as you use for PPPoA configurations on M Series routers and T Series routers.

[Subscriber Access, Network Interfaces]

- **Support for adjusting shaping rate and overhead accounting attributes based on PPPoE access line parameters for agent circuit identifier interface sets (MX Series routers with MPCs/MICs)**—Extends the functionality available in earlier Junos OS releases to enable you to configure the router to use the Actual-Data-Rate-Downstream [26-130] and Access-Loop-Encapsulation [26-144] DSL Forum vendor-specific attributes (VSAs) found in PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets to adjust the shaping-rate and overhead-accounting class of service (CoS) attributes, respectively, for dynamic agent circuit identifier (ACI) interface sets. In earlier Junos OS releases, you used this feature to adjust the shaping-rate and overhead-accounting attributes only for dynamic subscriber interfaces not associated with ACI interface sets.

The shaping-rate attribute is based on the value of the Actual-Data-Rate-Downstream VSA. The overhead-accounting attribute is based on the value of the Access-Loop-Encapsulation VSA, and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode) encapsulation. In subscriber access networks where the router passes downstream ATM traffic to Ethernet interfaces, the different Layer 2 encapsulations between the router and the PPPoE Intermediate Agent on the digital subscriber line access multiplexer (DSLAM) make managing the bandwidth of

downstream ATM traffic difficult. Using the Access-Loop-Encapsulation VSA to shape traffic based on frames or cells enables the router to adjust the shaping-rate and overhead-accounting attributes in order to apply the correct downstream rate for the subscriber.

You can enable this feature in either the dynamic profile that defines the ACI interface set, or in the dynamic profile for the dynamic PPPoE (**pp0**) subscriber interface associated with the ACI interface set, as follows:

- To configure the router to use the Actual-Data-Rate-Downstream VSA to adjust the shaping-rate CoS attribute, include the **vendor-specific-tags actual-data-rate-downstream** statement at the **[edit dynamic-profiles profile-name class-of-service dynamic-class-of-service-options]** hierarchy level.
- To configure the router to use the Access-Loop-Encapsulation VSA to adjust the overhead-accounting CoS attribute, include the **vendor-specific-tags access-loop-encapsulation** statement at the **[edit dynamic-profiles profile-name class-of-service dynamic-class-of-service-options]** hierarchy level.

When you enable this feature, the router adjusts the shaping-rate and overhead-accounting attributes when the dynamic ACI interface set is created and the router receives the PADI and PADR packets from the first subscriber interface member of the ACI interface set. The value of the Actual-Data-Rate-Downstream VSA in the PADI and PADR control packets overrides the **shaping-rate** value configured at the **[edit dynamic-profiles profile-name class-of-service traffic-control-profiles]** hierarchy level only if the Actual-Data-Rate-Downstream value is less than the **shaping-rate** value configured with the CLI. The value of the Access-Loop-Encapsulation VSA always overrides the **overhead-accounting** value configured at the **[edit dynamic-profiles profile-name class-of-service traffic-control-profiles]** hierarchy level.

As part of this feature, the output of the following operational commands has been enhanced to display the adjustment value (frame mode or cell mode) for the overhead-accounting attribute:

- **show class-of-service interface**
- **show class-of-service interface-set**
- **show class-of-service traffic-control-profile**

[Subscriber Access]

- **DHCP relay agent selective traffic processing based on DHCP options (MX Series routers)**—Subscriber management enables you to configure DHCP relay agent to provide subscriber support based on information in DHCP options. For DHCPv4 relay agent, you use DHCP option 60 and option 77 to identify the client traffic. For DHCPv6 relay agent, you use DHCPv6 option 15 and option 16.

You can use the DHCP option information to specify the action DHCP relay agent takes on client traffic that meets the specified match criteria, such as forwarding traffic to a specific DHCP server, or dropping the traffic. You can also specify a default action, which DHCP relay agent uses when the option string in the client traffic does not satisfy any match criteria or when no other action is configured.

To configure DHCP relay agent selective processing, you use the **relay-option** statement at the **[edit forwarding-options dhcp-relay]** or **[edit forwarding-options dhcp-relay dhcpv6]** hierarchy level. To display statistics for the number of forwarded packets, use the **show dhcp relay statistics** and **show dhcpv6 relay statistics** commands.

[Subscriber Access]

- **Ensuring that RADIUS clears existing session state before performing authentication and accounting for new sessions (MX Series routers)**—At subscriber session startup, the Junos OS authd process sends an Acct-On message to RADIUS servers. In some service provider environments, upon receipt of the Acct-On message, the RADIUS server cleans up the previous session state and removes accounting statistics. However, authentication or accounting for the new session can start before the RADIUS cleanup of the previous session—this can result in RADIUS deleting the new session's authentication and accounting information (which might include billing information).

To ensure that the new session's authentication and accounting information is not deleted, you can optionally configure authd to wait for an Acct-On-Ack response message from RADIUS before sending the new authentication and accounting updates to the RADIUS server. When this feature is enabled, all authentication requests fail until the router receives the Acct-On-Ack response from at least one configured RADIUS server.

To enable this feature, you configure the **wait-for-acct-on-ack** statement at the **[edit access profile profile-name accounting]** hierarchy level. To display the response status of the Acct-On messages (for example, Ack, Pending, None), use the **show network-access aaa accounting** command.

[Subscriber Access]

- **Enhanced local configuration of DNS name server addresses (MX Series Routers)**—You can now configure the DNS name server addresses locally per routing instance or per access profile. The new configuration applies to both terminated and tunneled PPP subscribers (IPv4 and IPv6), DHCP subscribers (DHCPv4 and DHCPv6), and IP-over-Ethernet (VLAN) subscribers. In earlier releases, the local configuration for the DNS server address applied only to DHCP subscribers (configured as a DHCP attribute), and only at the more granular level of the address pool.

As with the address-pool configuration, the new statements enable you to configure multiple DNS name server addresses per routing instance and access profile by issuing the statement for each address.

Because you can both configure name server addresses at more than one level and configure more than one address within a level, a preference order for the configurations determines which address is returned to the client.

- Within a configuration level, the preference order for the address matches the order in which the address is configured. For example, the first address configured within an access profile is preferred to the second address configured in that profile.
- Among configuration levels, the preference order depends on the client type:
 - For DHCP subscribers, the preference in descending order is:
RADIUS > DHCP address pool > access profile > global

- For non-DHCP subscribers, the preference in descending order is:

RADIUS > access profile > global

- Accordingly, all subscriber types prefer a name server address configured in RADIUS to the address configured anywhere else. When a name server address is configured only in a DHCP address pool, then no address is available to non-DHCP subscribers. For all subscriber types, the global name server address is used only when no other name server addresses are configured.

To configure a name server address in a routing instance, include the **domain-server-name-inet** or **domain-name-server** statement for IPv4 addresses, or the **domain-name-server-inet6** statement for IPv6 addresses, at the **[edit access]** hierarchy level.

To configure a name server address in an access profile, include any of the same statements at the **[edit access profile]** hierarchy level.



BEST PRACTICE: In practice, choose either the **domain-name-server** statement or the **domain-name-server-inet** statement for IPv4 addresses. They both have the same effect and there is no need to use both statements.

[Subscriber Access]

- **Gx-Plus support for service provisioning (MX Series routers)**—Gx-Plus now supports service (policy rule) provisioning, service activation, threshold notifications, threshold updates, service termination, and recovery. Previously, Gx-Plus supported only notification, termination, and recovery. To request subscriber service provisioning from the Policy Control and Charging Rules Function (PCRF), include the **provisioning-order gx-plus** statement in the subscriber access profile.

By default, Gx-Plus provisioning requests are made only for IPv4 subscribers. To enable requests to be made also for IPv6 subscribers, include the **include-ipv6** statement at the **[edit access gx-plus global]** hierarchy level.

The PCRF can request usage monitoring for the provisioned services for one or more of the following: number of bytes transmitted (CC-Output-Octets), number of bytes received (CC-Input-Octets), number of bytes transmitted and received (CC-Total-Octets), and elapsed time (CC-Time). If the specified threshold is reached, the router sends a usage report back to the PCRF. The PCRF can then return new threshold triggers and request that services be activated or deactivated.

When a subscriber has been provisioned with Gx-Plus, only the PCRF can activate or deactivate services for that subscriber. Accordingly, AAA rejects any RADIUS CoA or CLI service activation or deactivation requests for these subscribers. You can override PCRF control on an individual session, which is useful for session and service troubleshooting. To do so, issue the new **request network-access aaa subscriber set session-id** command. You can then activate and deactivate services with the existing **request network-access aaa subscriber add session-id** and **request network-access aaa subscriber delete session-id** commands, respectively.

[Subscriber Access]

- **Support for maintenance of CoS shaping rates for ANCP subscribers across ANCP restarts (MX Series Routers)**—When ANCP stops due to a process restart or graceful Routing Engine switchover (GRES), CoS now enforces the ANCP downstream shaping-rates until the CoS keepalive timer expires. When the timer expires, CoS reverts to its configured shaping-rate for the interfaces.

You can configure the CoS keepalive timer by including the existing **maximum-helper-restart-time seconds** statement at the **[edit protocols ancp]** hierarchy level. It specifies how much time other daemons such as CoS will wait for ANCP to restart and is used to configure the CoS rate update keepalive timer.

ANCP does not maintain TCP sessions from neighbors across the restart or graceful Routing Engine switchover (GRES). When it restarts, it must re-establish sessions with neighbors and subscriber sessions before the timer expires. For all the re-established sessions, ANCP updates CoS with the updated downstream shaping rates and provides DSL line attributes to the session database for AAA.

If CoS stops or restarts while ANCP is up, ANCP retransmits all known subscriber downstream rates to CoS. Any existing adjusted shaping rates that have not been updated revert to the configured CoS shaping rates when the CoS restart timer expires.

[Subscriber Access]

- **MAC address validation in enhanced network services modes**—MAC address validation is now optimized for scaling when the router is configured for Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode. When MAC address validation is enabled, the router compares the IP source and MAC source addresses against trusted addresses, and forwards or drops the packets according to the match and the validation mode. This feature is not available for IPv6.



NOTE: When the router is configured for either of the enhanced network services modes, MAC address validation is supported only on MPCs. If the router has both DPCs and MPCs, or only DPCs, you cannot configure the chassis to be in enhanced mode.

In contrast, when the router is configured for a normal (non-enhanced) network services mode, MAC address validation is supported on both DPCs and MPCs. The router can be populated completely with one or the other type of line card, or have a mix of both types. Normal network services mode is the default.

To configure an enhanced network services mode, include the **network-services service** statement at the **[edit chassis]** hierarchy level, and then configure MAC address validation as usual.



NOTE: In normal network services mode, you can use the `show interfaces statistics interface-name` command to display a per-interface count of the packets that failed validation and were dropped. In enhanced network services modes, this command does not count the dropped packets; you must contact Juniper Networks Customer Support for assistance in collecting this data.

[Subscriber Access]

- **Fail filters for RPF checks in dynamic profiles**—By default, unicast RPF checks prevent DHCP packets from being accepted on interfaces protected by the RPF check. When you enable an RPF check with a dynamic profile, you must configure a fail filter that identifies and passes DHCP packets.

To configure a fail filter, include the `fail-filter filter-name` statement at the `[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family rpf-check]` hierarchy level. To configure the terms of the fail filter, include the `filter filter-name` statement at the `[edit firewall family family]` hierarchy level. Include conditions in a filter term to identify DHCP packets, such as `from destination-port dhcp` and `from destination-address 255.255.255.255/32`. Define another filter term to drop all other packets that fail the RPF check. This feature is available for both IPv4 and IPv6 address families.

To confirm that the fail filter is active, you can issue the `show subscribers extensive` command, which displays the name of active filters.

[Subscriber Access]

- **Filtering traffic that is mirrored using DTCP-initiated subscriber secure policy**—You can now filter mirrored traffic before it is sent to a mediation device. This feature allows service providers to reduce the volume of traffic sent to a mediation device. For some types of traffic, such as IPTV or video on demand, it is not necessary to mirror the entire content of the traffic because the content might already be known or controlled by the service provider.

To configure, create a policy at the `[edit services radius-flow-tap policy policy-name]` hierarchy level. You can set up the policy to filter IPv4 or IPv6 traffic by source or destination address or port, protocol, or DSCP value. You then apply the policy by using the new DTCP attribute `X-Drop-Policy`. You can use the `X-Drop-Policy` attribute with the `ADD DTCP` command to begin filtering traffic when mirroring is triggered using the `ADD DTCP` command. To begin filtering traffic that is currently being mirrored, use the `X-Drop-Policy` attribute with the new `ENABLE DTCP` command. To stop filtering traffic that is currently being mirrored, use the `X-Drop-Policy` attribute with the new `DISABLE DTCP` command.

[Subscriber Access Configuration Guide]

- **Enhancements to Multicast Subscriber Flow Distribution in an Aggregated Ethernet Bundle (MX Series Routers)**—Enables you to both target and separate the distribution of multicast subscriber traffic using enhanced IP chassis network services mode in an aggregated Ethernet bundle that is configured without link protection.

This feature enhances already released scheduling and scaling improvements made for subscribers in an aggregated Ethernet bundle and includes support for the following:

- IP demux subscriber interfaces on the EQ DPC and MPC/MIC modules and VLAN demux subscriber interfaces on MPC/MIC modules.



NOTE: This feature is not supported for VLAN subscriber interfaces.

- Multicast using the **enhanced-ip** mode setting at the **[edit chassis network-services]** hierarchy level.
- Multicast traffic to egress in parallel with unicast traffic, sharing the CoS hierarchy and aggregated Ethernet flow distribution.
- Targeted multicast flow distribution over inter-chassis redundancy (ICR) configurations where multicast traffic flows toward the subscriber primary interface even if that interface resides on a remote chassis within the virtual system.
- The ability to separate unicast and multicast subscriber traffic on a per VLAN basis using OIF mapping.

Targeted distribution enables you to target egress traffic for subscribers on a link. The system distributes subscriber interfaces equally among the links. For multicast traffic to egress in parallel with unicast traffic, share the CoS hierarchy and aggregated Ethernet flow distribution:

- Configure subscriber distribution. See [Distribution of Demux Subscribers in an Aggregated Ethernet Interface](#).
- Configure the **network-services** statement at the **[edit chassis]** hierarchy level to use **enhanced-ip** mode to take advantage of using the EQ DPC and MPC/MIC modules.

Separated target distribution enables you to target multicast traffic to use a specific VLAN over the aggregated Ethernet interface instead of flowing over the same interface in parallel. To configure separated targeted distribution for a multicast link:

- Configure an interior gateway protocol. See the [Junos OS Routing Protocols Configuration Guide](#).
- Configure IGMP or MLD on the interfaces. See the [Junos OS Multicast Protocols Configuration Guide](#) for static configuration. See the [Junos OS Subscriber Access Configuration Guide](#) for dynamic configuration.
- Configure the **network-services** statement at the **[edit chassis]** hierarchy level to use **enhanced-ip** mode to take advantage of using the EQ DPC and MPC/MIC modules.
- Configure an OIF mapping for any subscriber VLAN interfaces. See [Example: Configuring Multicast with Subscriber VLANs](#) in the [Junos OS Multicast Protocols Configuration Guide](#).
- Configure the distribution type for demux subscribers on an aggregated Ethernet interface by including the **targeted-distribution** statement at the **[edit dynamic-profiles]**

profile-name interfaces demux0 unit unit-name or ***[edit interfaces demux0 unit unit-name]*** hierarchy level.

When links are removed, affected flows are redistributed among the remaining active backup links. When links are added to the system, no automatic redistribution occurs. New subscriber and multicast flows are assigned to the links with the least number of subscribers (typically, the new links). You can configure the system to periodically rebalance the distribution of subscribers on the links by including the ***rebalance-periodic time hours:minutes interval hours*** statement at the ***[edit interfaces ae0 aggregated-ether-options]*** hierarchy level. To manually rebalance the subscribers on the interface, issue the ***request interface rebalance interface interface-name*** command.

To display a summary of the targeted distribution on a logical interface, issue the ***show interface interface-name extensive*** command. To display the targeted distribution on a specific aggregated Ethernet bundle, issue the ***show interface targeting aex*** command.

[Subscriber Access, Network Interfaces]

- **Layer-2 Control Packets**—The forwarding path supports the following types of Layer-2 control packets (excluding Operation, Administration, and Maintenance (OAM) packets) in both directions, receiving and forwarding:
 - Ethernet control packets—ARP, IS-IS, 1588v2, Ethernet Synchronization Messaging Channel (ESMC).
- **Host Path**—The host path to and from the CPU is supported in the following ways:
 - Host-bound traffic, prioritized into multiple queues, to support various levels of traffic.
 - Hardware-based policing used to limit denial of service attacks.
 - Protocol and flow-based policing.
 - Code point-based classification and prioritization of packets from the host to the external world.
- **Counters and statistics**—Most packet and byte-level statistics for various entities in the forwarding path available in Junos OS are supported. The following counters and statistics are supported:
 - Ingress and egress packet and byte counters for logical interfaces, Ethernet pseudowires, and MPLS transit label-switched paths.
 - Discard packets counter for system-wide global Packet Forwarding Engine statistics.
- **Statistics collection and reporting for Gigabit Ethernet interfaces**—For Gigabit Ethernet interfaces, Packet Forwarding Engine statistics are disabled by default. To enable Gigabit Ethernet interface statistics, you must specifically configure them. To configure Gigabit Ethernet interface statistics, include the new ***statistics*** statement at the ***[edit interfaces interface-name unit logical-unit-number]*** hierarchy level. To display statistics, issue the ***show interfaces interface-name (brief-|-extensive)*** operational mode command.
- **Address Resolution Protocol (ARP) parameters**—The maximum number of ARP entries is 7,000.

- **Support for configuring NAS-Port and NAS-Port-Type RADIUS attributes per physical interface, VLAN, or S-VLAN (MX Series routers with MPCs/MICs)**—Enables you to configure the NAS-Port-Type (61) RADIUS IETF attribute, and an extended format for the NAS-Port (5) RADIUS IETF attribute, on a per-physical interface, per-static VLAN, or per-static stacked VLAN (S-VLAN) basis. The router passes the NAS-Port and NAS-Port-Type attributes to the RADIUS server during the authentication, authorization, and accounting (AAA) process.

The NAS-Port-Type attribute specifies the type of physical port that the network access server (NAS) uses to authenticate the subscriber. The NAS-Port attribute specifies the physical port number of the NAS that is authenticating the user, and is formed by a combination of the physical port's slot number, port number, adapter number, VLAN ID, and S-VLAN ID. The NAS-Port extended format configures the number of bits (bit width) for each field in the NAS-Port attribute: slot, adapter, port, VLAN, and S-VLAN.

Configuring the NAS-Port-Type and the extended format for NAS-Port on a per-VLAN, per-S-VLAN, or per-physical interface basis is useful in the following network configurations:

- **1:1 access model (per-VLAN basis)**—In a 1:1 access model, dedicated customer VLANs (C-VLANs) provide a one-to-one correspondence between an individual subscriber and the VLAN encapsulation.
- **N:1 access model (per-S-VLAN basis)**—In an N:1 access model, service VLANs are dedicated to a particular service, such as video, voice, or data, instead of to a particular subscriber. Because a service VLAN is typically shared by many subscribers within the same household or in different households, the N:1 access model provides a many-to-one correspondence between individual subscribers and the VLAN encapsulation.
- **1:1 or N:1 access model (per-physical interface basis)**—You can configure the NAS-Port-Type and NAS-Port format on a per-physical interface basis for both the 1:1 access model and the N:1 access model.

To configure the NAS-Port-Type and the format for NAS-Port on a per-VLAN, or per-S-VLAN, or per-physical interface basis, you must create a NAS-Port options definition. The NAS-Port options definition includes the NAS-Port extended format, the NAS-Port-Type, and either the VLAN range of subscribers or the S-VLAN range of subscribers to which the definition applies.

The basic tasks for configuring a NAS-Port options definition are as follows:

- To create a named NAS-Port options definition, include the **nas-port-options *nas-port-options-name*** statement at the **[edit interfaces *interface-name* radius-options]** hierarchy level.
- To configure the extended format for the NAS-Port, include the **nas-port-extended-format** statement and appropriate options at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]** hierarchy level. To include S-VLAN IDs, in addition to VLAN IDs, in the extended format, include the **stacked** statement at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name* nas-port-extended-format]** hierarchy level.

- To configure the NAS-Port-Type, include the **nas-port-type** *port-type* statement at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]** hierarchy level.
- To configure the VLAN range of subscribers to which the NAS-Port options definition applies, include the **vlan-ranges** statement at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]** hierarchy level. To specify all VLANs in the VLAN range, include the **any** statement at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name* vlan-ranges]** hierarchy level.
- To configure the S-VLAN range of subscribers to which the NAS-Port options definition applies, include the **stacked-vlan-ranges** statement at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name*]** hierarchy level. To specify all VLAN IDs in the outer tag of the S-VLAN range, include the **any** statement at the **[edit interfaces *interface-name* radius-options nas-port-options *nas-port-options-name* stacked-vlan-ranges]** hierarchy level. You cannot configure the inner tag (S-VLAN ID) of the S-VLAN range; the inner tag is always specified as **any** to represent all S-VLAN IDs.



NOTE: You can create a maximum of 16 NAS-Port options definitions per physical interface. Each definition can include a maximum of 32 VLAN ranges or 32 S-VLAN ranges, but cannot include a combination of VLAN ranges and S-VLAN ranges.

[Subscriber Access]

- **Support for one dynamic profile for both single-stack and dual-stack subscribers**—On PPP access networks, you can use one dynamic profile to support the following address combinations: IPv4 only, IPv6 only, and IPv4 and IPv6 dual stack.

[*Designing an IPv6 Architecture and Implementing IPv4 and IPv6 Dual Stack for Broadband Edge*]

- **Support for DHCPv6 requests that include a request for both DHCPv6 IA_NA and DHCPv6 prefix delegation**—For DHCPv6 subscribers on DHCP access networks, a client can solicit both an IA_NA address and a prefix for DHCP prefix delegation, and the session comes up even if either the address or the prefix is not allocated. In earlier releases, an error was returned if the BNG did not return both an address for DHCPv6 IA_NA and a prefix for DHCPv6 prefix delegation.

[*Designing an IPv6 Architecture and Implementing IPv4 and IPv6 Dual Stack for Broadband Edge*]

- **Support for new Juniper Networks Diameter AVP (MX Series routers)**—Junos OS supports a new Juniper Networks Diameter AVP, **Juniper-State-ID** (AVP code 2058). **Juniper-State-ID** specifies the value assigned to each synchronization cycle for the purpose of identifying which messages to discard. The **Juniper-State-ID** AVP can be included in Diameter messages and used by supported Diameter applications such as JSRC and PTSP.

[*Subscriber Access Configuration Guide*]

- **Subscriber management and services feature and scaling parity (MX2010 and MX2020)**—Starting in Junos OS Release 12.3R4, the MX2010 router and the MX2020 router 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 router and the MX2020 router match those of MX960 routers.

System Logging

New and deprecated system log tags—The following set of system log message is new in this release:

- **LLDP**—This section describes messages with the **LLDP** prefix. They are generated by the link layer discovery protocol process (lldpd) which is used by EX Series switches to learn and distribute device information on network links. The information allows the switch to quickly identify a variety of devices, including IP telephones, resulting in a LAN that interoperates smoothly and efficiently.

The following system log messages are new in this release:

- ASP_NAT_PORT_BLOCK_ACTIVE
- ASP_PCP_NAT_MAP_CREATE
- ASP_PCP_NAT_MAP_DELETE
- ASP_PCP_TPC_ALLOC_ERR
- ASP_PCP_TPC_NOT_FOUND
- AUTHD_ACCT_ON_ACK_NOT_RECEIVED
- CHASSISD_FPC_OPTICS_HOT_NOTICE
- CHASSISD_MAC_ADDRESS_VIRB_ERROR
- CHASSISD_RE_CONSOLE_ME_STORM
- COSD_CLASS_NO_SUPPORT_IFD
- COSD_CLASS_NO_SUPPORT_L3_IFL
- COSD_MAX_FORWARDING_CLASSES_ABC
- DDOS_SCFD_FLOW_AGGREGATED
- DDOS_SCFD_FLOW_CLEARED
- DDOS_SCFD_FLOW_DEAGGREGATED
- DDOS_SCFD_FLOW_FOUND
- DDOS_SCFD_FLOW_RETURN_NORMAL
- DDOS_SCFD_FLOW_TIMEOUT
- ESWD_VMEMBER_MAC_LIMIT_DROP
- FC_PROXY_NP_PORT_RESTORE_FAILED

- LIBJNX_PRIV_RAISE_FAILED
- LLDP_NEIGHBOR_DOWN
- LLDP_NEIGHBOR_UP
- PPMD_MIRROR_ERROR
- RPD_PARSE_BAD_COMMAND
- RPD_PARSE_BAD_FILE
- RPD_PIM_IP_INFINITE_HOLDTIME
- UFDD_LINK_CHANGE
- WEB_CERT_FILE_NOT_FOUND_RETRY
- WEB_DUPLICATE_HTTPD

The following system log messages are no longer documented, either because they indicate internal software errors that are not caused by configuration problems or because they are no longer generated. If these messages appear in your log, contact your technical support representative for assistance:

- FABOAMD_TASK_SOCK_ERR
- JCS_EXT_LINK_STATE
- JCS_RSD_LINK_STATE
- JCS_SWITCH_COMMUNICATION_OK
- LIBJNX_AUDIT_ERROR
- LIBJNX_COMPRESS_EXEC_FAILED
- LIBJNX_INVALID_CHASSIS_ID
- LIBJNX_INVALID_RE_SLOT_ID
- LIBJNX_REPLICATE_RCP_EXEC_FAILED

User Interface and Configuration

- **Support for HTTP Reverse Proxy and HTTP Transparent Proxy on Application Services Modular Line Card (MX240, MX480, MX960 3D Edge Universal Routers)**--The Application Services Modular Line Card with Media Flow Controller software installed enables configuring support for HTTP reverse proxy and HTTP transparent proxy caching.

The Application Services Modular Line Card (AS MLC) has three components:

- Application Services Modular Carrier Card (AS MCC)
- Application Services Modular Processing Card with 64G (AS MXC)
- Application Services Modular Storage Card with 6.4 TB capacity (AS MSC)

The AS MLC for MX Series routers supports high throughput for applications developed with Juniper Networks Media Flow Controller software. A Media Flow Controller

application functions as a web-caching proxy server that processes HTTP traffic. HTTP requests are routed to the Media Flow Controller either explicitly for a domain (reverse proxy) or by redirecting traffic based on a policy (transparent proxy).

Media Flow Controller software can operate in HTTP reverse proxy mode, HTTP transparent proxy mode, or mixed mode.

In HTTP reverse proxy configurations, the service provider provides services to a set of domains (content providers) that buy content caching capability from the service provider. Clients connect to content providers through virtual IP (VIP) addresses. Service providers in the reverse proxy scenario generally deploy the routers with AS MLC hardware to honor service requests (such as caching) from the domain users.

HTTP reverse proxy supports the following features:

- Retrieve and deliver content from content providers in response to client requests as if the content originated at the proxy
- Prevent attacks from the Web when a firewall is included in the reverse proxy configuration
- Load balance client requests among multiple servers
- Lessen load on origin servers by caching both static and dynamic content

In HTTP transparent proxy configurations, the service provider implements the AS MLC to improve its own caching capability and to reduce the load on its own network. Implementing caching on an MX Series router with an AS MLC improves the retrieval speeds for data and optimizes the back-end network utilization. Typically, HTTP transparent proxy retrieves content for clients from the Internet. The client identifies the target of the request, which is commonly a location on the Internet.

HTTP transparent proxy does not enforce local policies: it does not add, delete, or modify information contained in the messages it forwards. HTTP transparent proxy is a cache for data. HTTP transparent proxy satisfies client requests directly because it retains the data that was previously requested by the same or by a different client. HTTP transparent proxy improves the efficiency and performance of network bandwidth within the content provider's data center.

In mixed mode, both reverse proxy and transparent proxy are configured on the same router.

[Junos OS Ethernet Interfaces Configuration Guide]

- **Support for 10-port 10-Gigabit Ethernet MIC with SFPP on MPC3E (MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.3, the MPC3E supports the 10-port 10-Gigabit Ethernet MIC with SFPP (MIC3-3D-10XGE-SFPP). The 10-port 10-Gigabit Ethernet MIC with SFPP uses SFP+ optical transceiver modules for connectivity. The MIC supports up to ten 10-Gigabit Ethernet interfaces and occupies MIC slot 0 or 1 in the MPC3E.

The MIC supports both LAN-PHY and WAN-PHY interface framing modes. You can configure the framing mode on a per-port basis. Use the existing command to switch between LAN-PHY and WAN-PHY modes:

set interfaces *interface-name* framing (lan-phy | wan-phy)

The 10-Gigabit Ethernet MIC with SFPP supports the same features as the other MICs supported on the MPC3E.

[See [MPC3E MIC Overview](#).]

[*MX Series 3D Universal Edge Router Line Card Guide, Ethernet Interfaces Configuration Guide, System Basics.*]

- **Inline flow monitoring support (MX Series routers with MPC3E)**—Junos OS Release 12.3 supports inline flow monitoring and sampling services on MX Series routers with MPC3E. To configure inline flow monitoring, include the **inline-jflow** statement at the **[edit forwarding-options sampling instance *instance-name* family inet output]** hierarchy level. Inline flow monitoring supports a specified sampling output format designated as IP_FIX and uses UDP as the transport protocol. Inline flow monitoring supports both IPv4 and IPv6 formats.

[See [Configuring Inline Sampling](#), and [Protocols and Applications Supported by MX240, MX480, MX960 MPC3E](#).]

- **Enhancements to IPv4 and IPv6 inline-jflow Flow IPFIX Record Templates**—Junos OS Release 12.3 introduces the VLAN ID field in the inline-jflow flow IPFIX record templates for IPv4 and IPv6 traffic. The VLAN ID field is not valid for egress traffic, and returns a value of 0 for egress traffic. Note that the VLAN ID field is updated while creating a new flow record, and any change in VLAN ID after that might not be updated in the record. [*Services Interfaces*]
- **Support for IPv6 Flow Servers on Interfaces Hosted on MICs or MPCs**—Starting with Release 12.3, Junos OS enables you to configure IPv6 flow servers for inline flow monitoring. When you configure an IPv6 address for the **flow-server** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output]** hierarchy level, you must also configure an IPv6 address for the **inline-jflow source-address** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) output]** hierarchy level. You can configure different families that use IPv4 and IPv6 flow servers under the same sampling instance. However, you can configure only one flow server per family. [*Services Interfaces*]
- **Optical transceiver support for MIC3-3D-1X100GE-CFP on MPC3E (MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.3, the 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP) on MPC3E supports the CFP-100GBase-ER4 optical transceiver.

If the ambient temperature exceeds 40° C and the other MIC slot is not empty, the CFP-100GBase-ER4 optical transceiver is put on low power mode, which disables the transmitter and takes the optic modules on the MIC offline. This protects the optical transceiver and also prevents damage to adjacent components.

When the optical transceiver is taken offline, you might see the following system log (syslog) message:

PIC 1 optic modules in Port 0 8 have been disabled since ambient temperature is over threshold.



NOTE: The CFP-100GBase-ER4 optical transceiver is NEBS (Network Equipment Building System) compliant only when plugged into the 100-Gigabit Ethernet MIC with CFP and when the other MIC slot is empty.

To reactivate the optical transceiver, use the **request chassis optics fpc-slot *fpc-slot-number* reactivate** operational mode command.

[*System Basics Configuration Guide*]

- **Optical transceiver support for MIC3-3D-10XGE-SFPP on MPC3E (MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.3, the 10-port 10-Gigabit Ethernet MIC with SFPP (MIC3-3D-10XGE-SFPP) on MPC3E supports the SFPP-10GE-ZR optical transceiver.

If the ambient temperature exceeds 40° C, the transmitter on the SFPP-10GE-ZR optical transceiver is disabled, which takes the optic modules on the MIC offline. This protects the optical transceiver and also prevents damage to adjacent components.

When the optical transceiver is taken offline, you might see the following system log (syslog) message:

PIC 1 optic modules in Port 0 8 have been disabled since ambient temperature is over threshold.



NOTE: The SFPP-10GE-ZR optical transceiver is not NEBS (Network Equipment Building System) compliant when plugged into the 10-port 10-Gigabit Ethernet MIC with SFPP. If other optical transceivers have been added, they can continue to operate.

To reactivate the optical transceiver, use the **request chassis optics fpc-slot *fpc-slot-number* reactivate** operational mode command.

[*System Basics Configuration Guide*]

VPLS

- **PIM Snooping for VPLS**—PIM snooping is introduced to restrict multicast traffic to interested devices in a VPLS. A new statement, **pim-snooping**, is introduced at the [**edit routing-instances *instance-name* protocols**] hierarchy level to configure PIM snooping on the PE device. PIM snooping configures a device to examine and operate only on PIM hello and join/prune packets.

A PIM snooping device snoops PIM hello and join/prune packets on each interface to find interested multicast receivers and populates the multicast forwarding tree with this information. PIM snooping can also be configured on PE routers connected as pseudowires, which ensures that no new PIM packets are generated in the VPLS, with the exception of PIM messages sent through LDP on the pseudowire.

PIM snooping improves IP multicast bandwidth in the VPLS core. Only devices that are members of a multicast group receive the multicast traffic meant for the group. This ensures network integrity and reliability, and multicast data transmission is secured.

[See [Example: Configuring PIM Snooping for VPLS](#).]

- **Improved VPLS MAC address learning on T4000 routers with Type 5 FPCs**—Junos OS Release 12.3 enables improved virtual private LAN service (VPLS) MAC address learning on T4000 routers with Type 5 FPCs by supporting up to 262,143 MAC addresses per VPLS routing instance. In Junos OS releases before Release 12.3, T4000 routers with Type 5 FPCs support only 65,535 MAC addresses per VPLS routing instance.

To enable the improved VPLS MAC address learning on T4000 routers with Type 5 FPCs:

- Include the **enhanced-mode** statement at the **[edit chassis network-services]** hierarchy level and perform a system reboot. By default, the **enhanced-mode** statement is not configured.
- Include the **mac-table-size** statement at the **[edit routing-instances vpls protocols vpls]** hierarchy level.



NOTE:

- You can configure the **enhanced-mode** statement only on T4000 routers with Type 5 FPCs.
 - The **enhanced-mode** statement supports up to 262,143 MAC addresses per VPLS routing instance. However, the MAC address learning limit for each interface remains the same (that is, 65,535 MAC addresses).
 - You must reboot the system after configuring the **enhanced-mode** statement. Otherwise, the improved VPLS MAC address learning does not take effect.
 - When the T4000 router reboots after the **enhanced-mode** statement has been configured, all Type 4 FPCs go offline.
-

[See [Configuring Improved VPLS MAC Address Learning on T4000 Routers with Type 5 FPCs](#).]

- **VPLS Multihoming (support extended to FEC 129)**—Enables you to connect a customer site to two or more PE routers to provide redundant connectivity. A redundant PE router can provide network service to the customer site as soon as a failure is detected. VPLS multihoming helps to maintain VPLS service and traffic forwarding to and from the multihomed site in the event of network failures. BGP-based VPLS autodiscovery (FEC 129) enables each VPLS PE router to discover the other PE routers that are in the same VPLS domain. VPLS autodiscovery also automatically detects when PE routers are added or removed from the VPLS domain. You do not need to manually configure the VPLS and maintain the configuration when a PE router is added or deleted. VPLS autodiscovery uses BGP to discover the VPLS members and to set up and tear down pseudowires in the VPLS. To configure, include the **multi-homing** statement at the **[edit routing-instances *instance-name*]** hierarchy level.

[See [Example: Configuring VPLS Multihoming \(FEC 129\)](#).]

- **BGP Path Selection for Layer 2 VPNs and VPLS**—By default, Juniper Networks routers use just the designated forwarder path selection algorithm to select the best path to reach each Layer 2 VPN or VPLS routing instance destination. However, you can now configure the routers in your network to use both the BGP path selection algorithm and the designated forwarder path selection algorithm. The Provider routers within the network can use the standard BGP path selection algorithm. Using the standard BGP path selection for Layer 2 VPN and VPLS routes allows a service provider to leverage the existing Layer 3 VPN network infrastructure to also support Layer 2 VPNs and VPLS. The BGP path selection algorithm also helps to ensure that the service provider's network behaves predictably with regard to Layer 2 VPN and VPLS path selection. This is particularly important in networks employing route reflectors and multihoming.

The PE routers continue to use the designated forwarder path selection algorithm to select the preferred path to reach each CE device. The VPLS designated forwarder algorithm uses the D-bit, preference, and PE router identifier to determine which path to use to reach each CE device in the Layer 2 VPN or VPLS routing instance.

To enable the BGP path selection algorithm for Layer 2 VPN and VPLS routing instances, do the following:

- Specify a unique route distinguisher on each PE router participating in a Layer 2 VPN or VPLS routing instance.
- Configure the **l2vpn-use-bgp-rules** statement on all of the PE and Provider routers participating in Layer 2 VPN or VPLS routing instances. You can configure this statement at the **[edit protocols bgp path-selection]** hierarchy level to apply this behavior to all of the routing instances on the router or at the **[edit routing-instances *routing-instance-name* protocols bgp path-selection]** hierarchy level to apply this behavior to a specific routing instance.

On all of the PE and Provider routers participating in Layer 2 VPN or VPLS routing instances, run Junos OS Release 12.3 or later. Attempting to enable this functionality on a network with a mix of routers that both do and do not support this feature can

result in anomalous behavior. [See [Enabling BGP Path Selection for Layer 2 VPNs and VPLS.](#)]

VPNs

- **Provider Edge Link Protection in Layer 3 VPNs**—A precomputed protection path can be configured in a Layer 3 VPN such that if a link between a CE router and a PE router goes down, the protection path (also known as the backup path) between the CE router and an alternate PE router can be used. This is useful in an MPLS service provider network, where a customer can have dual-homed CE routers that are connected to the service provider through different PE routers. In this case, the protection path avoids disruption of service if a PE-CE link goes down.

The protection path can be configured on a PE router in a Layer 3 VPN by configuring the **protection** statement at the **[edit routing-instances *instance-name* protocols bgp family inet unicast]** or **[edit routing-instances *instance-name* protocols bgp family inet6 unicast]** hierarchy level.

The **protection** statement indicates that protection is required on prefixes received from a particular neighbor or family. After protection is enabled for a given family, group, or neighbor, protection entries are added for prefixes or next hops received from the respective peer.

A protection path can be selected only if the best path has already been installed by BGP in the forwarding table. This is because a protection path cannot be used as the best path. There are two conditions under which the protection path will not work:

- When configured for an internal BGP peer.
- When configured with external and internal BGP multipath.

[See [Example: Configuring Provider Edge Link Protection in Layer 3 VPNs.](#)]

- **Edge node failure protection for LDP-signaled pseudowires**—This feature provides a fast protection mechanism against egress PE router failure when transport LSPs are RSVP-TE LSPs. This is achieved by using multihomed CEs, upstream assigned labels, context-specific label switching (**egress-protection** and **context-identifier** statements), and by extending RSVP facility backup fast reroute (FRR) to enable node protection at the penultimate hop router of the LSP. With node protection capability, the penultimate hop router can perform local repair upon an egress PE failure and redirect pseudowire traffic very quickly to a protector PE through a bypass LSP. You must configure a Layer 2 circuit and transport LSP to enable this feature. Use the **show rsvp session** and **show mpls lsp** commands to view bypass LSP and backup LSP information on the penultimate hop router and a protector PE router.

[VPNs]

- **Support for Configuring More Than One Million Layer 3 VPN Labels**—For Layer 3 VPNs configured on Juniper Networks routers, Junos OS normally allocates one inner VPN label for each customer edge (CE)-facing virtual routing and forwarding (VRF) interface of a provider edge (PE) router. However, other vendors allocate one VPN label for each route learned over the CE-facing interfaces of a PE router. This practice increases the number of VPN labels exponentially, which leads to slow system processing and slow convergence time.

For Juniper Networks routers participating in a mixed vendor network with more than one million Layer 3 VPN labels, include the **extended-space** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]** hierarchy level. The **extended-space** statement is disabled by default.

We recommend that you configure the **extended-space** statement in mixed vendor networks containing more than one million BGP routes to support Layer 3 VPNs. However, because using this statement can also enhance the Layer 3 VPN performance of Juniper Networks routers in networks where only Juniper Networks routers are deployed, we recommend configuring the statement in these networks as well. [See [Accepting BGP Updates with Unique Inner VPN Labels in Layer 3 VPNs.](#)]

- **Layer 2 circuit switching protection**—Provides traffic protection for the Layer 2 circuit paths configured between PE routers. In the event the path (working path) used by a Layer 2 circuit fails, traffic can be switched to an alternate path (protection path). Switching protection is supported for locally switched Layer 2 circuits and provides 1 to 1 protection for each Layer 2 circuit interface.

Each working path can be configured to have either a protection path routed directly to the neighboring PE router or indirectly using a pseudowire configured through an intermediate PE router. The protection path provides failure protection for the traffic flowing between the PE routers. Ethernet OAM monitors the status of these paths. When OAM detects a failure, it reroutes the traffic from the failed working path to the protection path. You can configure OAM to revert the traffic automatically to the working path when it is restored. You can also manually switch traffic between the working path, the protection path, and back.

Layer 2 circuit switching protection is supported on MX Series routers only. Nonstop routing (NSR) and graceful Routing Engine switchover (GRES) are not supported.

To enable Layer 2 circuit switching protection, include the **connection-protection** statement at the **[edit protocols l2circuit local switching interface *interface-name* end-interface]** hierarchy level. You also need to configure OAM for the working path and the protection path by configuring the **maintenance-association** statement and sub-statements at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain *maintenance-domain-name*]** hierarchy level. [See [Example: Configuring Layer 2 Circuit Switching Protection](#)]

- **History enhancements for Layer 2 circuit, Layer 2 VPN, and FEC 129-based pseudowires**—Adds the instance-history option to the **show vlpls connections** and **show l2vpn connections** commands. Also adds instance-level logs for the following events:
 - Catastrophic events
 - Route withdrawals
 - Pseudowire switchovers
 - Connect protect switchovers
 - Protect interface swaps
 - Interface flaps (interface down events)
 - Label block changes

These logs are maintained until the instance is deleted from the configuration.

**Related
Documentation**

- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 357](#)

Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

- [Changes in Default Behavior and Syntax on page 162](#)
- [Changes Planned for Future Releases on page 180](#)

Changes in Default Behavior and Syntax

The following are changes made to Junos OS default behavior and syntax.

- [High Availability \(HA\) and Resiliency on page 163](#)
- [Interfaces and Chassis on page 163](#)
- [IPv6 on page 167](#)
- [J-Web on page 167](#)
- [Junos OS XML API and Scripting on page 167](#)
- [MPLS on page 167](#)
- [Multicast on page 168](#)
- [Network Address Translation \(NAT\) on page 168](#)
- [Network Management and Monitoring on page 169](#)
- [Routing Policy and Firewall Filters on page 170](#)
- [Routing Protocols on page 171](#)
- [Security on page 172](#)
- [Services Applications on page 172](#)
- [Subscriber Access Management on page 174](#)
- [User Interface and Configuration on page 180](#)
- [VPNs on page 180](#)

High Availability (HA) and Resiliency

- **Configuration support to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure**—You can now configure the **prefer-status-control-active** statement with the **status-control standby** configuration at the **[edit interfaces aeX aggregated-ether-options mc-ae]** hierarchy level to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP does not go down unless the router is down. You must also configure the **hold-time down** value (at the **[edit interfaces interface-name]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents traffic loss by ensuring that when the router with the **status-control active** configuration goes down, the router with the **status-control standby** configuration does not go into standby mode.
- **Change in behavior of request system reboot command for MX Series Virtual Chassis (MX Series routers with MPC/MIC interfaces)**—Starting in Junos OS Release 12.3R3, the behavior of the **request system reboot** command has been changed when used with an MX Series Virtual Chassis. To reboot both Routing Engines in each member router of the Virtual Chassis, you can now use any of the following commands:
 - **request system reboot**
 - **request system reboot all-members**
 - **request system reboot all-members both-routing-engines**

In Junos OS Release 12.2R2 and earlier releases, the **request system reboot** command rebooted only the master Routing Engine in each member router in the MX Series Virtual Chassis.

[See [request system reboot](#).]

Interfaces and Chassis

- On the Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP (Model number PB-1CHOC48-STM16-IQE), in the presence of line remote defect indication (LRDI) and line alarm indication signal (LAIS), the 3 LSBs of K2 byte cannot be monitored or viewed through the **show interfaces coc48-x/y/z extensive** command.
- **Multichassis Link Aggregation (MC-LAG)**—When you configure the **prefer-status-control-active** statement at the **[edit interfaces aex aggregated-ether-options mc-ae events iccp-peer-down]** hierarchy level, you must also configure the **status-control active** statement at the **[edit interfaces aex aggregated-ether-options mc-ae]** hierarchy level. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning. [*Junos OS Ethernet Interfaces Configuration Guide*]
- Starting with Junos OS Release 12.3, the output of the **show chassis fabric topology** operational command for a TX Matrix Plus Router has been changed. The string that identifies a cross-chassis serial link for an F13 SIB now includes an additional character to identify the SF chip to which the link connects.

- **New fast-failover option for LACP**—You can now configure the Link Aggregation Control Protocol for aggregated Ethernet interfaces to facilitate subsecond failover. To override the default behavior for the IEEE 802.3ad standard and allow the standby link always to receive traffic, include the **fast-failover** statement at the **[edit interfaces aeX aggregated-ether-options lacp]** hierarchy level. [*Junos OS Ethernet Interfaces Configuration Guide*]
- **New options for Multichassis Link Aggregation (MC-LAG)**—For MC-LAG, you can now specify one of two actions to take if the Inter-Chassis Communication Protocol (ICCP) peer if the switch or router goes down. To bring down the interchassis link logical interface if the peer goes down, include the **force-icl-down** statement at the **[edit interfaces aeX aggregated-ether-options events iccp-peer-down]** hierarchy level. To have the router or switch become the active node when a peer goes down, include the **prefer-status-control-active** statement at the **[edit interfaces aeX aggregated-ether-options mc-ae events iccp-peer-down]** hierarchy level. When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement at the **[edit interfaces aeX aggregated-ether-options-mc-ae]** hierarchy level. If you do not configure the **status-control** as **active** with the **prefer-status-control-active** statement, the router or switch does not become the active node if a peer goes down. [*Junos OS Ethernet Interfaces Configuration Guide*]
- **Enhancement to show interfaces queue command**—The output for the **show interfaces queue** command now displays rate-limit statistics for class-of-service schedulers for all IQ and Enhanced IQ (IQ2E) PICs when rate-limiting is configured, even when no traffic is dropped. When rate limiting is configured but no traffic is dropped, the output for the **RL-dropped packets** and **RL-dropped-bytes** fields display the value zero (0). Previously, these fields were not displayed when no traffic was dropped and rate-limiting was configured. To configure rate-limiting for queues before packets are queued for output, you include the **rate-limit** statement at the **[edit class-of-service schedulers transmit-rate rate]** hierarchy level. [*Interfaces Command Reference*]
- **New Link Aggregation Control Protocol (LACP) Commands and SNMP MIB**—You can now view and clear LACP timeout entries. To display information about LACP timeout entries, use the **show lacp timeouts** command. Include the **interfaces interface-name** option to view timeout information about a specific interface only. To clear LACP timeout entries, use the **clear lacp timeouts** command. Include the **interfaces interface-name** option to clear timeout information for a specific interface only. A new SNMP MIB is now also available. The **jnxLacpAggTimeout** MIB lists all interfaces where the **jnxLacpTimeOut** trap is sent. [*Interfaces Command Reference*]
- **Connectivity Fault Management MEPs on Layer 2 Circuits and Layer 2 VPNs (MX Series 3D Universal Edge Routers)**—On interfaces configured on Modular Port Concentrators (MPCs) only, you no longer need to configure the **no-control-word** statement for Layer 2 circuits and Layer 2 VPNs over which you are running CFM maintenance endpoints (MEPs). The control word is enabled by default. For all interfaces not configured on MPCs, you need to continue to include the **no-control-word** statement at either the **[edit protocols l2circuit neighbor neighbor-id interface interface-name]** or the **[edit routing-instances routing-instance-name protocols l2vpn]** hierarchy level when you configure CFM MEPs. [*Ethernet Interfaces Configuration Guide*]

- The OID **jnxBfdSessIntfName** has been added to the BFD SNMP MIB to associate the BFD session and the interface it uses.

[SNMP MIBs and Traps Guide]

- On the Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP (model number PB-1CHOC48-STM16-IQE), in the presence of line remote defect indication and line alarm indication signal, the 3 least significant bits of the K2 byte cannot be monitored or viewed through the **show interfaces coc48-x/y/z extensive** command.
- Starting with Junos OS Release 12.2R1, the quality level parameter for a Synchronous Ethernet interface is optional when the **quality-mode** option is enabled and the **selection-mode** option is set to **receive-quality**. The default quality level for a Synchronous Ethernet interface is SEC for the **option-1** network type and ST3 for the **option-2** network type.
- Starting with Junos OS Release 12.3, the output of the **show chassis fabric topology** operational mode command for a TX Matrix Plus router has been changed. The string that identifies a cross-chassis serial link for an F13 SIB now includes an additional character to identify the SF chip to which the link connects.
- **Version Compatibility for Junos SDK**—As of Junos OS Release 12.3, Junos applications will install on Junos only if the application is built with the same release as the Junos OS release on which the application is being installed. For example, an application built with Release 12.3R2 will only install on Junos OS Release 12.3R2 and will not install on Junos OS Release 12.3R1 or Junos OS Release 12.3R3 or Junos OS Release 13.1R1.
- **Enhancement to Link Layer Discovery Protocol (LLDP) (MX Series and T Series routers)**—You can now configure LLDP to generate the interface name as the port ID Type, Length, and Value (TLV). To generate the interface name as the port ID TLV, include the **interface-name** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level. The default behavior is to generate the SNMP Index of the interface as the port ID TLV. If you have changed the default behavior, include the **locally-assigned** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level to reenact the default behavior of generating the SNMP Index of the interface as the port ID TLV. When you configure LLDP to generate the interface name as the port ID TLV, the **show lldp neighbors** command displays the interface name in the **Port ID** field. The default behavior is for the command to display the SNMP index of the interface in the **Port ID** field. *[Ethernet Interfaces Configuration Guide, Interfaces Command Reference]*
- **Configuring the flow-tap service for IPv6 traffic:** The **family inet | inet6** statement at the **[edit services flow-tap]** hierarchy enables you to specify the type of traffic for which you want to apply the flow-tap service. If the family statement is not included, the flow-tap service is, by default, applied to the IPv4 traffic. To apply the flow-tap service to IPv6 traffic, you must include the **family inet6** statement in the configuration. To enable the flow-tap service for IPv4 and IPv6 traffic, you must explicitly configure the family statement for both inet and inet6 families.

However, you cannot configure the flow-tap service for IPv6 along with port mirroring or sampling of IPv6 traffic on routers that support LMNR-based FPCs. This restriction is in effect even if the router does not have any LMNR-based FPC installed on it. There is no restriction on configuring the flow-tap service on routers that are configured for port mirroring or sampling of IPv4 traffic. *[Services Interfaces]*

- Prior to Junos OS Release 12.2, when you issue the **show system memory** command on MX80 routers, the **unable to load pmap_helper module: No such file or directory** error message is displayed in the output of the command. Starting with Junos OS Release 12.2, PMAP information is correctly displayed in the output of this command for MX80 and ACX Series routers.

[*System Basics and Services Command Reference*]

- **New range for message-rate-limit** – The range for **message-rate-limit** under the **syslog** configuration for services has changed to 0 through 2147483647.
- **Configuration support to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure**—You can now configure the **prefer-status-control-active** statement with the **status-control standby** configuration at the **[edit interfaces aeX aggregated-ether-options mc-ae]** hierarchy level to prevent the LACP MC-LAG system ID from reverting to the default LACP system ID on ICCP failure. Use this configuration only if you can ensure that ICCP does not go down unless the router is down. You must also configure the **hold-time down** value (at the **[edit interfaces interface-name]** hierarchy level) for the interchassis link with the **status-control standby** configuration to be higher than the ICCP BFD timeout. This configuration prevents traffic loss by ensuring that when the router with the **status-control active** configuration goes down, the router with the **status-control standby** configuration does not go into standby mode.
- **Layer 2 port mirroring**—Starting in Junos OS Release 13.2, you can enable Layer 2 port mirroring of host-generated outbound packets only on MPCs on MX Series 3D Universal Edge Routers.
- **Changes to DDoS protection policers for PIM and PIMv6 (MX Series with MPCs, T4000 with FPC5)**—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.

- **Preventing the filtering of packets by ARP policers (MX Series routers)**—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.

[Network Interfaces, Protocol Family and Interface Address Properties]

IPv6

- **Change in automatically generated virtual-link-local-address for VRRP over IPv6—**The seventh byte in the automatically generated virtual-link-local-address for VRRP over IPv6 is 0x02. This change makes the VRRP over IPv6 feature in Junos OS 12.2R5, 12.3R3, 13.1R3, and later releases inoperable with Junos OS 12.2R1, 12.2 R2, 12.2R3, 12.2R4, 12.3R1, 12.3R2, 13.1R1, and 13.3R2 releases if an automatically generated virtual-link-local-address ID used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

J-Web

- On all M Series, MX Series, and T Series platforms, the username field does not accept HTML tags or the < and > characters. The following error message appears: **A username cannot include certain characters, including < and >.**

Junos OS XML API and Scripting

- **IPv6 address text representation is stored internally and displayed in command output using lowercase—**Starting with Junos OS Release 11.1R1, IPv6 addresses are stored internally and displayed in the command output using lowercase. Scripts that match on an uppercase text representation of IPv6 addresses should be adjusted to either match on lowercase or perform case-insensitive matches.
- **<get-configuration> RPC with inherit="inherit" attribute returns correct time attributes for committed configuration—**In Junos OS Release 12.3R1, when you configured some interfaces using the interface-range configuration statement, if you later requested the committed configuration using the <get-configuration> RPC with the inherit="inherit" and database="committed" attributes, the device returned **junos:changed-localtime** and **junos:changed-seconds** in the RPC reply instead of **junos:commit-localtime** and **junos:commit-seconds**. This issue is fixed in Junos OS Release 12.3R2 and later releases so that the device returns the expected attributes in the RPC reply.

MPLS

- Starting in Junos OS Release 9.3, when you run the **show route table mpls.0 protocol ccc** command, the next-hop information includes the outgoing interface and the name of the label-switched path. Previously, the next-hop information included the outgoing interface and the MPLS label value.

[MPLS]

- **Policers for MPLS LSPs (T Series Core Routers)**—You can now configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. To do so, you must configure the LSP policer as a logical interface policer. Include the **logical-interface-policer** statement at the **[edit firewall policer policer-name]** hierarchy level. Previously, you could not configure an MPLS LSP policer as a logical interface policer. When you configure an MPLS LSP policer as a logical interface policer, that single policer polices traffic for all protocol families for an LSP. An MPLS LSP policer not configured as a logical interface policer continues to police traffic for a specific protocol family only.

[Firewall Filters and Traffic Policers Configuration Guide, MPLS Applications Configuration Guide]

- Starting in Junos OS Release 12.2, at the end of each **adjust-interval**, LSP's **max_average** for the **auto-bandwidth** functionality does not reset to zero. The **max_average** retains the value from the last interval until the first sample of the current interval is received. When the first sample of the current interval is received, the **max_average** is updated to the first sample value.

In the **show mpls lsp** command output, the value for **Max AvgBW util** now displays the value of the maximum average bandwidth utilization from the previous interval until the first sample of the current interval is obtained.

[MPLS Operational Mode Commands]

Multicast

- In a bootstrap router (BSR)-enabled bidirectional PIM domain, mixing Junos OS Release pre-12.1R7 releases and later releases can cause unexpected outages. If you have a deployment with routers running Junos OS Release pre-12.1R7 and if you upgrade a subset of the routers to Junos OS Release 12.1R7 or later, the group-to-RP mapping across the domain breaks and an outage occurs.

Network Address Translation (NAT)

- **Protection of MX, M, and T Series routers from denial of service (DOS) attacks**—New CLI options provide improved protection against DOS attacks.
 - NAT mapping refresh behavior—Prior to Junos OS Release 12.3, a conversation was kept alive when either inbound or outbound flows were active. This remains the default behavior. As of this release, you can also specify mapping refresh for only inbound flows or only outbound flows. To configure mapping refresh behavior, include the **mapping-refresh (inbound | outbound | inbound-outbound)** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.
 - EIF inbound flow limit—Previously, the number of inbound connections on an EIF mapping was limited only by the maximum flows allowed on the system. You can now configure the number of inbound flows allowed for an EIF. To limit the number of inbound connections on an EIF mapping, include the **EIF-flow-limit number-of-flows** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.

[Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions]

- **Limitation on number of terms for NAT rules applied to inline services interfaces**—You are limited to a maximum of 200 for a NAT rule that is applied to an inline services (type si) interface. If you specify more than 200 terms, you will receive following error when you commit the configuration:

```
[edit]
'service-set service-set-name'
  NAT rule rule-name with more than 200 terms is disallowed for si-x/y/z.n
error: configuration check-out failed
```

- The method for computing the block size for deterministic port block allocation for network port translation (NAPT) when the configured block size is zero has changed, and is computed as follows:

$$\text{block-size} = \text{int}(64512 / \text{ceil}[(\text{Nr_Addr_PR_Prefix} / \text{Nr_Addr_PU_Prefix})])$$

where:

64512 is the maximum available port range per public IP address.

Nr_Addr_PR_Prefix is the number of usable pre-NAT IPv4 subscriber addresses in a **from** clause match condition

Nr_Addr_PU_Prefix is the number of usable post-NAT IPv4 addresses configured in the NAT pool

Network Management and Monitoring

- Each Routing Engine runs its own SNMP process (**snmpd**), allowing each Routing Engine to maintain its own engine boots. However, if both Routing Engines have the same engine ID and the Routing Engine with the lesser **snmpEngineBoots** value is selected as the master Routing Engine during the switchover process, the **snmpEngineBoots** value of the master Routing Engine is synchronized with the **snmpEngineBoots** value of the other Routing Engine.

[Network Management Configuration Guide]

- **New system log message indicating the difference in the Packet Forwarding Engine counter value (M Series, MX Series, and T Series)**—Effective in Junos OS Release 12.3R9, if the counter value of a Packet Forwarding Engine is reported lesser than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the **MIB2D_COUNTER_DECREASING** system log message for that specific counter.
- **SNMP MIB support for subscriber interface index**—The Juniper Networks enterprise-specific Subscriber MIB, whose object ID is **{jnxSubscriberMibRoot 1}**, supports a new MIB table, **jnxSubscriberInterfaceHardwareIndexTable**, to display the index of subscriber interfaces. The **jnxSubscriberInterfaceHardwareIndexTable**, whose object identifier is **{jnxSubscriberGeneral 4}**, contains **jnxSubscriberInterfaceHardwareIndexEntry** that maps to the specification of each subscriber. You must provide the session ID of the subscriber in the SNMP **Get** and **GetNext** queries. When you perform an SNMP walk operation, you need to provide only the name of the subscriber interface index table or the name of the object.

Each `jnxSubscriberInterfaceHardwareIndexEntry`, whose object identifier is `{jnxSubscriberInterfaceHardwareIndexTable 1}`, contains the objects listed in [Table 3 on page 170](#).

Table 3: jnxSubscriberInterfaceHardwareIndexTable

Object	Object ID	Description
<code>jnxSubscriberInterfaceHardwareIndexHandleHiWord</code>	<code>jnxSubscriberInterfaceHardwareIndexEntry 1</code>	Subscriber handle associated with each subscriber. Returns the most significant 32 bits of the 64-bit subscriber ID. The value of the subscriber handle is a monotonically increasing number.
<code>jnxSubscriberInterfaceHardwareIndexHandleLoWord</code>	<code>jnxSubscriberInterfaceHardwareIndexEntry 2</code>	Subscriber handle associated with each subscriber. Returns the least significant 32 bits of the 64-bit subscriber ID. The value of the subscriber handle is a monotonically increasing number.
<code>jnxSubscriberInterfaceHardwareIndex</code>	<code>jnxSubscriberInterfaceHardwareIndexEntry 3</code>	The hardware index of the subscriber interface.

[SNMP MIBs and Traps Reference]

Routing Policy and Firewall Filters

- **Firewall filter option to force premium treatment for traffic (MX Series routers)—**
By default, a hierarchical policer processes the traffic it receives according to the traffic's forwarding class. Premium, expedited-forwarding traffic has priority for bandwidth over aggregate, best-effort traffic. Now you can include the **force-premium** option at the `[edit firewall filter filter-name term term-name]` hierarchy level to ensure that traffic matching the term is treated as premium traffic by a subsequent hierarchical policer, regardless of its forwarding class. This traffic is given preference over any aggregate traffic received by that policer.

Consider a scenario where a firewall filter is applied to an interface that receives both expedited-forwarding voice traffic and best-effort video traffic. Traffic that matches the first term of the filter is passed to a hierarchical policer in the second term. The hierarchical policer also receives best-effort data traffic from another source. The filtered video traffic is treated the same as this data traffic, as aggregate traffic with a lower priority than the premium voice traffic. Consequently, some of the video traffic might be dropped and some of the data traffic passed on.

To avoid that situation, include the **force-premium** option in the firewall filter term that passes traffic to the hierarchical policer. This term forces the video traffic to be marked

as premium traffic. The hierarchical policer gives both the voice traffic and the video traffic priority over the aggregate data traffic.



NOTE: The **force-premium** filter option is supported only on MPCs.

- **Extends support for Layer 2 policers to MX Series routers with MPC3**—You can now configure Layer 2 policers for the ingress and egress interfaces on MX Series routers with MPC3. Policers include single-rate two-color, single-rate three-color (color-blind and color-aware), and two-rate three-color (color-blind and color-aware). To configure Layer 2 policing, include the **policer** statement at the **[edit firewall]** hierarchy level.

[Junos OS Firewall Filters and Traffic Policers]

Routing Protocols

- Bidirectional Forwarding Detection (BFD) is a protocol that verifies the liveness of data paths. One desirable application of BFD is to detect connectivity to routers that span multiple network hops and follow unpredictable paths. On M Series, MX Series, and T Series platforms only, starting in Junos OS Release 12.3, multihop BFD runs on the CPU in the FPC, DPC, or MPC. Previously, multihop BFD ran from the Routing Engine.
- Junos OS Release 12.3 supports a new **show firewall templates-in-use operational** command. This command enables you to display the names of filters configured using the filter statement at either the **[edit firewall]** or **[edit dynamic-profiles profile-name firewall]** hierarchy level and that are being used as templates for dynamic subscriber filtering. The command also displays the number of times the filter has been referenced by subscribers accessing the network. *[Routing Protocols and Policies Command Reference]*
- When configuring the **advertise-external** statement for an AS confederation, we recommend that EBGp peers belonging to different autonomous systems be configured in a separate EBGp peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.

[Routing Protocols Guide]

- If you configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**. To avoid a conflict in the two route distinguisher values, we recommend ensuring that the first half of the route distinguisher obtained by configuring the **route-distinguisher** statement be different from the first half of the route distinguisher obtained by configuring the **route-distinguisher-id** statement.

[Routing Protocols Guide]

- **BGP hides a route received with a label block size greater than 256 (M Series, MX Series, and T Series)**—When a BGP peer (running Junos OS) sends a route with a label block size greater than 256, the local speaker hides the route and does not re-advertise this route. The output of the **show route detail/extensive hidden/all** displays the hidden route and states the reason as **label block size exceeds max supported value**. In earlier

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

- **Configure and establish targeted sessions with third-party controllers using LDP targeted neighbor (M Series and MX Series)**—Starting with Junos OS Release 12.3R10, you can configure LDP targeted neighbor to third-party controllers for applications such as route recorder that wants to learn label-FEC bindings of an LSR. LDP targeted neighbor helps to establish a targeted session with controllers for a variety of applications.

Security

- **DDoS protection support for more protocol groups and packet types (MX Series 3D Universal Edge Routers)**—DDoS protection now supports the following additional protocol groups and packet types:
 - **amtv4**—IPv4 automatic multicast (AMT) traffic.
 - **amtv6**—IPv6 AMT traffic.
 - **frame-relay**—Frame relay traffic.
 - **inline-ka**—Inline service interfaces keepalive traffic.
 - **inline-svcs**—Inline services traffic.
 - **keepalive**—Keepalive traffic.
 - **l2pt**—Layer 2 protocol tunneling traffic.

Two packet types are available for the **frame-relay** protocol group:

- **frf15**—Multilink frame relay FRF.15 packets.
- **frf16**—Multilink frame relay FRF.16 packets.

The PPP protocol group has an additional packet type available, **mlppp-lcp** for MLPPP LCP packets.

[*System Basics and Services Command Reference*]

- In all supported Junos OS releases, regular expressions can no longer be configured if they require more than 64MB of memory or more than 256 recursions for parsing.

This change in the behavior of Junos OS is in line with the Free BSD 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 process (rpd).

Services Applications

- Starting in Junos OS Release 12.3R3, the **destination-address** statement in a firewall rule **from** statement might not have the address value of 0::00 with IPv6.

```
[edit services stateful-firewall rule rule-name term term-name from]
destination-address (address | any-unicast) <except>;
```

This issue is being tracked by [PR857106](#).

- In Junos OS Release 12.3R3 and earlier, when peers in a security association (SA) became unsynchronized, packets with invalid security parameter index (SPI) values could be sent out, and the receiving peer dropped those packets. The only way to recover was to manually clear the SAs or wait for them to time out. Starting in Junos OS release 12.3R4, you can enable automatic recovery by using the new **respond-bad-spi max-responses** configuration statement, which appears under the **[edit services ipsec-vpn ike policy]** hierarchy level. This command results in a resynchronization of the SAs when invalid SPIs are received.
- **New statement for resynchronization of SAs**—In Junos OS Release 12.3R3 and earlier, when peers in a security association (SA) became unsynchronized, packets with invalid security parameter index (SPI) values could be sent out, and the receiving peer dropped those packets. The only way to recover was to manually clear the SAs or wait for them to time out. Starting in Junos OS release 12.3R4, you can enable automatic recovery by using the **respond-bad-spi max-responses** configuration statement, which appears under the **[edit services ipsec-vpn ike policy]** hierarchy level. This statement results in a resynchronization of the SAs when invalid SPIs are received.

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

```
[edit services ipsec-vpn ike policy]
respond-bad-spi max-responses;
```

- **Protection of routers from denial of service (DOS) attacks**—New CLI options provide improved protection against DOS attacks.
 - NAT mapping refresh behavior—Prior to this release, a conversation was kept alive when either inbound or outbound flows were active. This remains the default behavior. As of 12.3R6, you can also specify mapping refresh for only inbound flows or only outbound flows. To configure mapping refresh behavior, include the **mapping-refresh (inbound | outbound | inbound-outbound)** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.
 - EIF inbound flow limit—Previously, the number of inbound connections on an EIF mapping was limited only by the maximum flows allowed on the system. Starting in Release 12.3R6, you can configure the number of inbound flows allowed for an EIF. To limit the number of inbound connections on an EIF mapping, include the **eif-flow-limit number-of-flows** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.

The **show services nat pool detail** command now shows the current number of EIF flows and the flow limit.

```
Current EIF Inbound flows count: 0
EIF flow limit exceeded drops: 0
```

- Maximum dropped flows—There is a default maximum of 2000 drop flows allowed per PIC at a given instance of time. You can now configure the maximum number of drop flows allowed per direction (ingress and egress) at any given instance of time. This enables you to limit the creation of drop flows during a denial-of-service (DOS)

attack. When the maximum number of drop flows exceeds the configured or default limit, drop flows are not created and packets are dropped silently, meaning that no syslog message is generated for the dropped packets. If maximum dropped flows is configured, the appropriate error counters are incremented for packets dropped due to exceeded limits.

To limit the number of drop flows, include the **max-drop-flows ingress *ingress-flows* egress *egress-flows*** statement at the **[edit services service-set *service-set-name*]** hierarchy level.

The **show services stateful-firewall statistics extensive** command now shows the maximum flow drop counters when **max-drop-flows** is configured.

```
Drop Flows:
Maximum Ingress Drop flows allowed: 20
Maximum Egress Drop flows allowed: 20
Current Ingress Drop flows: 0
Current Egress Drop flows: 0
Ingress Drop Flow limit drops count: 0
Egress Drop Flow limit drops count: 0
```

- TWAMP connection/session will come 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.

Subscriber Access Management

- **Effect of changing the forwarding class configuration with PPP fast keepalive (MX Series routers with MPC/MIC interfaces)**—To change the default queue assignment (forwarding class) for outbound traffic generated by the Routing Engine, you can include the **forwarding-class *class-name*** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level.

For PPP fast (inline) keepalive LCP Echo-Request and LCP Echo-Reply packets transmitted between an MX Series router with MPCs/MICs and a PPP client, changing the forwarding class configuration takes effect immediately for both new PPP-over-Ethernet (PPPoE), PPP-over-ATM (PPPoA), and L2TP network server (LNS) subscriber sessions created after the configuration change, and for existing PPPoE, PPPoA, and LNS subscriber sessions established before the configuration change.

In earlier Junos OS releases with PPP fast keepalive, forwarding class configuration changes applied only to new PPPoE, PPPoA, and LNS subscriber sessions created after the configuration change. The forwarding class setting was fixed for existing PPPoE, PPPoA, and LNS subscriber sessions, and could not be changed until the session was terminated and re-established.

[Junos OS Subscriber Access Configuration Guide, Junos OS Class of Service Configuration Guide]

- **Display of a warning message for enhanced policer statistics (MX Series routers)**—When you commit a configuration that contains the **enhanced-policer** statement at the **[edit chassis]** hierarchy level, a warning message is displayed stating that all the FPCs in the router need to be rebooted for the configuration changes to become effective. At this point, you must confirm that you want to proceed with the

reboot of the FPCs. If you do not reboot the FPCs, the FPCs return all 0s (zeros) when you perform a query for the retrieval of detailed statistics—for example, when you issue the **show firewall detail** command.

[*System Basics, Chassis-Level Features*]

- When an MX Series router configured as an L2TP network server (LNS) sends an Access-Request message to RADIUS for an LNS subscriber, the LNS now includes the Called-Station-ID-Attribute when it receives AVP 21 in the ICRQ message from the L2TP network concentrator (LAC).
- The **user *username*** option for the **clear services l2tp session** command is no longer available in the CLI for LNS on MX Series routers. Added to the option's previous unavailability for LAC on MX Series routers, this means that L2TP on MX Series routers does not support clearing L2TP sessions based on subscriber username. As an alternative, you can determine the session ID for the username by issuing the **show subscribers detail** command, and then remove the session with the **clear services l2tp session local-session-id *session-id*** command.

[*Subscriber Access*]

- The **user *username*** option for the **show services l2tp session** command is no longer available in the CLI for L2TP LAC or L2TP LNS on MX Series routers. To view L2TP session information organized by subscriber username, you can issue the **show subscribers detail** command or the **show network-access aaa subscribers username** command.

[*Subscriber Access*]

- **Enhanced filtering for tracing PPP and PPPoE operations (MX Series routers)**—Capturing relevant traces for particular PPP and PPPoE subscribers increases in complexity as the number of subscribers increases. New filter options have been added to simplify tracing PPP service operations and PPPoE subscriber operations in a scaled subscriber environment. You can include one or more of the following options at the [**edit protocols ppp-services traceoptions filter**] or [**edit protocols pppoe traceoptions filter**] hierarchy levels:
 - **aci *regular-expression***—Regular expression to match the agent circuit identifier provided by PPP or PPPoE client.
 - **ari *regular-expression***—Regular expression to match the agent remote identifier provided by PPP or PPPoE client.
 - **service *regular-expression***—Regular expression to match the name of PPP or PPPoE service.
 - **underlying-interface *interface-name***—Name of a PPP or PPPoE underlying interface. You cannot use a regular expression for this filter option.

When you apply more than one of these trace filters, events for a particular connection are traced only when it matches all of the filter conditions. For example, when you configure the following filter options, PPP (jpppd) events are traced only for PPP connections where the agent circuit identifier begins with the string west-metro-ge and the agent remote identifier includes the string CUST-0102:

```
user@host1> set protocol ppp-service traceoptions filter aci west-metro-ge*
user@host1> set protocol ppp-service traceoptions filter ari *CUST-0102*
```

Similarly, when you configure the following filter options, PPPoE events are traced only for PPPoE connections where the subscribers are on static interface pp0.50001 and receive the premium service:

```
user@host1> set protocol pppoe traceoptions filter interface pp0.50001
user@host1> set protocol pppoe traceoptions filter service premium
```

The amount of information logged when a connection matches the filters is considerably less than when no filters are applied. If the connection does not match the configured filters, some information is still logged, but only a minimal amount.

[Subscriber Access]

- **Increased visibility for PPP session state in trace logs (MX Series routers)**—Log files generated by tracing jpppd (ppp-service) operations now display the interface name for each line of the traced events. The new information might also include the module or the session state and event type for each event. This new information appears immediately after the timestamp and makes it easier to distinguish PPP packet exchange and session states in the logs.

[Subscriber Access]

- **Interface names logged for PPPoE messages (MX Series routers)**—Log output for PPPoE PADI, PADM, PADN, PADO, PADR, PADS, and PADT packets now explicitly includes the interface name rather than just the index.

[Subscriber Access]

- **Microsecond timestamps for certain tracing operations (MX Series routers)**—The logs generated when tracing authd, jpppd, and pppoe operations have been enhanced to provide more precise timestamps. The timestamps now record events at microsecond intervals.

[Subscriber Access]

- On MX80 routers, you can configure only four inline services physical interfaces as anchor interfaces for L2TP LNS sessions: si-1/0/0, si-1/1/0, si-1/2/0, si-1/3/0. You cannot configure si-0/0/0 for this purpose on MX80 routers.
- **Source Class Parameterized Match Condition (MX Series routers with MPCs/MICs)**—You can now reference **source-class** in the parameterized match condition of the dynamic profile filter. Source class usage allows you to limit traffic to specific subscribers from specific network zones. These limits are per subscriber and the profile name is communicated using RADIUS. The source-class parameterized match condition is supported for both IPv4 and IPv6.

[Subscriber Access Configuration Guide]

- **L2TP support for SNMP statistics (MX Series routers)**—By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in the table have a default value of zero.

Table 4: SNMP Counters for L2TP Statistics

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	Tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 30,000 sessions, none of these statistics is more than 15 minutes old.



BEST PRACTICE: The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

Subscriber Access

- **RADIUS accounting support of duplicate reporting for nondefault VRFs (MX Series routers)**—You can now configure duplicate RADIUS accounting records to be sent to a nondefault VRF; that is, to an LS:RI combination other than default:default. You can also specify up to five access profiles in the target VRF that list the RADIUS servers that receive the duplicate reports. Include the **vrf-name** statement at the new **[edit access profile profile-name accounting duplication-vrf]** hierarchy level to designate the single nondefault target VRF. Include the **access-profile-name** statement at the same hierarchy level to designate the access profiles listing the RADIUS servers.
- **DNS address assignment in DHCPv6 IA_NA and IA_PD environments (MX Series)**—Starting in Junos OS Release 12.3R3 and Release 13.3R1, DHCPv6 local server returns the DNS server address (DHCPv6 attribute 23) as a global DHCPv6 option, rather than as an IA_NA or IA_PD suboption. DHCPv6 returns the DNS server address that is specified in the IA_PD or IA_NA pools—if both address pools are requested, DHCPv6 returns the address specified in the IA_PD pool only, and ignores any DNS address in the IA_NA pool.

In releases prior to 12.3R3, and in releases 13.1 and 13.2, DHCPv6 returns the DNS server address as a suboption inside the respective DHCPv6 IA_NA or IA_PD header. You can use the **multi-address-embedded-option-response** statement at the **[edit system services dhcp-local-server dhcpv6 overrides]** hierarchy level to revert to the prior behavior. However, returning the DNS server address as a suboption can create interoperability issues for some CPE equipment that cannot recognize the suboption information.

[Subscriber Access]

- **Support for VLAN ID none configuration for MC-LAG bridge domains in active-active mode (MX Series)**—To facilitate forwarding and media access control (MAC) and Address Resolution Protocol (ARP) synchronization among multichassis link aggregation (MC-LAG) peers when the VLAN identifier is **none**, you must now configure a service identifier within bridge domains in active-active mode.

To configure a service identifier for a bridge domain, configure the **service-id** statement at the **[edit bridge domain bridge-domain-name]** hierarchy level. You must configure the same service identifier for MC-LAG peers.

- ANCP sessions may not persist across a graceful Routing Engine switchover and may need to be reestablished.
- **New vpi and vci options in show subscribers command (MX Series routers with MPCs and ATM MICs with SFP)**—Adds the following two new options to the **show subscribers** operational command to enable you to display information about active subscribers using PPPoE-over-ATM, PPP-over-ATM (PPPoA), IP-over-ATM (IPoA), or bridged IP-over-Ethernet-over-ATM to access the router over an ATM network:
 - **vpi**—ATM virtual path identifier (VPI) on the subscriber's physical interface, in the range 0 through 65535
 - **vci**—ATM virtual circuit identifier (VCI) for each VPI configured on the subscriber interface, in the range 0 through 255

In earlier Junos OS releases, the **vpi** and **vci** options were not available for the **show subscribers** command.

To display information about ATM subscriber interfaces based on their VPI and VCI values so you can better distinguish ATM-based subscribers from Ethernet-based subscribers, you can use the new **vpi** and **vci** options for the **show subscribers** command together or separately. For example, the following **show subscribers** command includes both the **vpi** and **vci** options to display extensive information about the active PPPoE-over-ATM subscriber using VPI 40 and VCI 50. The **ATM VPI** and **ATM VCI** fields are new in this output.

```
user@host> show subscribers vpi 40 vci 50 extensive
Type: PPPoE
User Name: testuser
IP Address: 100.0.0.2
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: pp0.0
Interface type: Static
MAC Address: 00:00:65:23:01:02
State: Active
```

```

Radius Accounting ID: 2
Session ID: 2
ATM VPI: 40
ATM VCI: 50
Login Time: 2012-12-03 07:49:26 PST
IP Address Pool: pool_1
IPv6 Framed Interface Id: 200:65ff:fe23:102

```

[*Subscriber Access, System Basics and Services Command Reference*]

- **Updated AAA Terminate Reason Mappings (MX Series routers)**—The AAA **idle-timeout** terminate reason is now mapped to the RADIUS accounting Idle Timeout (4) terminate cause, and the AAA **session-timeout** terminate reason is now mapped to the RADIUS Session Timeout (5) terminate cause. In earlier releases, both terminate reasons were mapped to the RADIUS accounting NAS Request (10) terminate cause.

To support backward compatibility, you can configure the router to support the previous behavior—use the **terminate-code aaa shutdown (idle-timeout | session-timeout) radius 10** statement at the **[edit access]** hierarchy level.

[*Subscriber Access*]

- **ATM subscriber enhancements for configuring RADIUS NAS-Port extended format (MX Series routers with MPCs/MICs)**—Enables you to use the same access profile to configure an extended format for the NAS-Port (5) RADIUS IETF attribute for both ATM subscribers and Ethernet subscribers. In earlier Junos OS releases, you used the access profile to configure the NAS-Port extended format only for Ethernet-based subscribers.

For ATM subscribers, the NAS-Port extended format configures the number of bits (bit width) in the **slot-width**, **adapter-width**, **port-width**, **vpi-width**, and **vci-width** fields in the NAS-Port attribute. Each field can be 1 through 32 bits wide; however, the combined total of the widths of all fields must not exceed 32 bits, or the configuration fails.

To configure the NAS-Port extended format for ATM subscribers in an access profile, include the new **atm** stanza and appropriate ATM bit width options in the **nas-port-extended-format** statement at the **[edit access profile profile-name radius options]** hierarchy level.

Instead of globally configuring an extended format for the NAS-Port attribute in an access profile, you can configure the NAS-Port extended format on a per-physical interface basis for both Ethernet interfaces and ATM interfaces. In earlier Junos OS releases, you configured the NAS-Port extended format only for Ethernet interfaces.

To configure the NAS-Port extended format for an ATM interface, include one or both of the **vpi-width** and **vci-width** options in the **nas-port-extended-format** statement at the **[edit interfaces interface-name radius-options nas-port-options nas-port-options-name]** hierarchy level.

- **DHCPv6 Relay Agent (MX Series)**—Starting in Junos OS Release 12.3R3, during the subscriber authentication or client authentication process, you can identify a subset of the DHCPv6 Relay Agent Remote-ID option (option 37) in the client PDU name to be concatenated with the username instead of concatenating the entire Remote-ID. You can use the **enterprise-id** and **remote-id** statements at the **[edit forwarding-options**

dhcp-relay dhcpv6 authentication username-include relay-agent-remote-id] and the **[edit system services dhcp-local-server dhcpv6 authentication username-include relay-agent-remote-id]** hierarchy levels.

- **DHCP client IP address (MX Series)**—Starting in Junos OS Release 12.2, you can configure the subnet to which the DHCP local server matches the requested IP address. The server accepts and uses an active client's requested IP address to address assignment only when the requested address and the IP address of the DHCP server interface are in the same subnet. The server accepts and uses a passive client's requested IP address only when the requested address and the IP address of the relay interface are in the same subnet.

User Interface and Configuration

- **Enhancement to set date ntp command**—You can now specify an authentication-key number for the NTP server used to synchronize the date and time on the router or switch. Include the new **key number** option with the **set date ntp** command. The key number you include must match the number you configure for the NTP server at the **[edit system ntp authentication-key number]** hierarchy level.
- **TFEB Slot**—On MX80 routers, the FPC Slot output field has been changed to TFEB Slot for the **show services accounting flow inline-jflow**, **show services accounting errors inline-jflow**, and **show services accounting status inline-jflow** commands.

VPNs

- On a Layer 3 VPN PE routing device, a direct subnet route on a LAN PE-CE interface is advertised with a matching next-hop label. Previously, when there were multiple matching next hops, one of the next-hop labels was selected for the direct subnet route. There was room for improvement because a packet with a destination address matching the subnet route might need to be sent to another next hop in the LAN. Starting in Release 12.3, Junos OS no longer advertises the direct subnet route on a LAN PE-CE interface when there are multiple matching next hops. The direct subnet route on LAN PE-CE interface is advertised only if there is a single matching next hop. *[VPNs]*
- Starting in Junos OS Release 11.4, vrf-import policies must reference a target community in the from clause. If the import policy does not reference a specific community target or if the referenced community is a wildcard, the commit operation fails. As an exception, the policy does not need to reference a community target in the from clause when the policy action in the then clause is "reject." Prior to Junos OS Release 11.4, when the vrf-import policy did not reference a specific community target in the from clause, the commit operation succeeded, but the import policy had a non-deterministic effect. *[VPNs]*

Changes Planned for Future Releases

The following are changes planned for future releases.

Routing Protocols

- **Change in the Junos OS support for the BGP Monitoring Protocol (BMP)**—In Junos OS Release 13.3 and later, the currently supported version of BMP, BMP version 1, as

defined in Internet draft draft-ietf-grow-bmp-01, is planned to be replaced with BMP version 3, as defined in Internet draft draft-ietf-grow-bmp-07.txt. Junos OS can support only one of these versions of BMP in a release. Therefore, Junos OS Release 13.2 and earlier will continue to support BMP version 1, as defined in Internet draft draft-ietf-grow-bmp-01. Junos OS Release 13.3 and later support only the updated BMP version 3 defined in Internet draft draft-ietf-grow-bmp-07.txt. This also means that beginning in Junos OS Release 13.3, BMP version 3 configurations are not backwards compatible with BMP version 1 configurations from earlier Junos OS releases.

[Routing Protocols]

- **Removal of support for provider backbone bridging (MX Series routers) from Release 14.1**—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. When you upgrade MX Series routers running Junos OS Releases 12.3, 13.2, or 13.3 to Junos OS Release 14.1 and if your deployment contains PBB settings in configuration files, the configuration files after the upgrade need to be modified to remove the PBB-specific attributes because PBB is not supported in Release 14.1 and later.

[Provider Backbone Bridging]

**Related
Documentation**

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 181](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 357](#)

Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 12.3R11 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 \(CoS\) on page 182](#)
- [MPLS on page 183](#)
- [Routing Policy and Firewall Filters on page 183](#)
- [Software Installation and Upgrade on page 183](#)

Class of Service (CoS)

- If you define more than one forwarding class for a given queue number, do not use the name of a default forwarding class for one of the new classes, because doing so causes the forwarding class with the default name to be deleted. For example, do not configure the following, because doing so deletes the **best-effort** class:

```
user@host# set class-of-service forwarding-classes class be queue-num 0
user@host# set class-of-service forwarding-classes class best-effort queue-num 0
user@host# commit
```

MPLS

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

Routing Policy and Firewall Filters

- On MX Series, subscriber management uses firewall filters to capture and report the volume-based service accounting counters that are used for subscriber billing. You must always consider the relationship between firewall filters and service accounting counters, especially when clearing firewall statistics. When you use the **clear firewall** command (to clear the statistics displayed by the **show firewall** command), the command also clears the service accounting counters that are reported to the RADIUS accounting server. For this reason, you must be cautious in specifying which firewall statistics you want to clear. When you reset firewall statistics to zero, you also zero the counters reported to RADIUS.

Software Installation and Upgrade

- **Downgrading to Junos OS 12.3 when the configuration includes the targeted-distribution statement**—In Junos OS Release 12.3, the **targetted-distribution** statement at the **[edit interfaces demux0 unit logical-unit-number]** hierarchy level is misspelled. Starting in Junos OS Release 13.3, the spelling for this statement is corrected to **targeted-distribution**. If you use the misspelled **targetted-distribution** statement in Junos OS Release 13.3 or higher, the CLI corrects the spelling to **targeted-distribution** in your configuration, so existing scripts still work. The correct spelling is not backward compatible; Junos OS Release 12.3 supports only the **targetted-distribution** spelling. If you downgrade from Release 13.3 or higher to Release 12.3, all correctly spelled **targeted-distribution** statements are removed from the configuration and configuration scripts with the correct spelling fail.

Related Documentation

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)

- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers](#) on page 357

Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

Errata

Hardware

- The *Protocols and Applications Supported by MX240, MX480, MX960, MX2010, and MX2020 MPCs* topic erroneously states that support was introduced in Junos OS Release 10.4 for IEEE 802.3ah OAM (discovery and link monitoring, fault signaling and detection, and remote loopback). In fact, this support was introduced in Junos OS Release 11.1.

Class of Service

- The *Example: Configuring Scheduling Modes on Aggregated Interfaces* topic fails to mention the following additional information regarding the parameters that are scaled for aggregated interface member links when the scheduler parameters are configured using scheduler maps:

Apart from transmit rate and buffer size that are scaled when the parameters are configured using scheduler maps, shaping rate is also scaled if you configure it in bits per second (bps). Shaping rate is not scaled if you configure it as a percentage of the available interface bandwidth.

[*Class of Service, Schedulers on Aggregated Ethernet and SONET/SDH Interfaces*]

- The following additional information regarding the processing of custom EXP rewrite rules on MPCs applies to the *Rewriting IEEE 802.1p Packet Headers with an MPLS EXP Value* topic:

For MPCs, default EXP rewrite rules do not exist for logical interfaces. The EXP CoS bits for MPLS labels are obtained from the IP precedence bits for IP traffic. The EXP bits for labels that are pushed or swapped are inherited from the current label of the MPLS packets. For non-IP and non-MPLS packets, the EXP bits are set to 0. If a custom EXP rewrite rule is configured on the core-facing interface, then it overrides the EXP bits.

[*Class of Service, CoS Re-Marking of Packets Entering or Exiting the Network*]

- The *Configuring Tunnel Interfaces on MX Series Routers* topic in the *Services Interfaces Configuration Guide* fails to state that Ingress queuing and tunnel services cannot be configured on the same MPC as it causes the Packet Forwarding Engine forwarding to stop. Each feature can, however, be configured and used separately.
- The *enhanced-policer* topic in the *Junos OS Subscriber Access Configuration Guide* fails to include a reference to the *Enhanced Policer Statistics Overview* topic. The overview topic explains how the enhanced policer enables you to analyze traffic statistics for debugging purposes.

The enhanced policer statistics are as follows:

- Offered packet statistics for traffic subjected to policing.
- OOS packet statistics for packets that are marked out-of-specification by the policer. Changes to all packets that have out-of-specification actions, such as discard, color marking, or forwarding-class, are included in this counter.
- Transmitted packet statistics for traffic that is not discarded by the policer. When the policer action is discard, the statistics are the same as the statistics that are within specification; when the policer action is non-discard (loss-priority or forwarding-class), the statistics are included in this counter.

To enable collection of enhanced statistics, include the **enhanced-policer** statement at the **[edit chassis]** hierarchy level. To view these statistics, include the **detail** option when you issue the **show firewall**, **show firewall filter *filter-name***, or **show policer** command.

High Availability (HA) and Resiliency

- The *MPC3E on MX Series Routers Overview* topic in the *Junos OS System Basics: Chassis-Level Features Configuration Guide* incorrectly states that configuration of an MX Series Virtual Chassis is not supported on an MPC3E module. In fact, the MPC3E module supports all MX Series Virtual Chassis features, including Layer 2 and IEEE 802.3ad link aggregation features.

An MX Series Virtual Chassis configuration on an MPC3E module or on MPC/MIC interfaces does not support the Spanning Tree Protocol.

[See [MPC3E on MX Series Routers Overview](#).]

- In Junos OS Release 11.4 and later releases, the *Example: Replacing a Routing Engine in a Virtual Chassis Configuration for MX Series 3D Universal Edge Routers* topic in the *MX Series Interchassis Redundancy Using Virtual Chassis* pathway page failed to mention that for a replacement Routing Engine shipped from the factory that you plan to install in an MX Series Virtual Chassis member router, you must modify the default factory configuration to enable proper operation of the Virtual Chassis. The documentation has been updated to include this information in Junos OS Release 13.2 and later releases, as follows:

A Routing Engine shipped from the factory is loaded with a default factory configuration that includes the following stanza at the [edit] hierarchy level:

```
[edit]
system {
  commit {
    factory-settings {
      reset-virtual-chassis-configuration;
    }
  }
}
```

When this configuration stanza is present, the Routing Engine can operate only in a standalone chassis and *not* in an MX Series Virtual Chassis member router. As a result, if you install this Routing Engine in the standby slot of a Virtual Chassis member router (**member1-re1** in this example), the Routing Engine does not automatically synchronize with the master Routing Engine and boot in Virtual Chassis mode.

To ensure that the standby factory Routing Engine successfully synchronizes with the master Routing Engine, you must remove this standalone chassis configuration stanza from the standby factory Routing Engine and verify that it reboots in Virtual Chassis mode before you install the Junos OS release.

To modify the Routing Engine factory configuration to ensure proper operation of the MX Series Virtual Chassis:

1. Log in to the console of the new Routing Engine as the user **root** with no password.
2. Configure a plain-text password for the **root** (superuser) login.

```
{local:member1-re1}[edit system]
root# set root-authentication plain-text-password
New password: type password here
Retype new password: retry password here
```

3. Delete the standalone chassis configuration.

```
{local:member1-re1}[edit]
root# delete system commit factory-settings reset-virtual-chassis-configuration
```

4. Commit the configuration.

The new Routing Engine synchronizes the Virtual Chassis member ID with the master Routing Engine and boots in Virtual Chassis mode.

5. Verify that the new Routing Engine is in Virtual Chassis mode.

During the boot process, the router displays the following output to indicate that it has synchronized the Virtual Chassis member ID (1) with the master Routing Engine and is in Virtual Chassis mode.

```
...
virtual chassis member-id = 1
virtual chassis mode      = 1
...
```

- For a two-member MX Series Virtual Chassis to function properly, you must enable enhanced IP network services on both member routers when you first set up the Virtual Chassis. If necessary, you can also enable enhanced IP network services for an existing Virtual Chassis.

Enhanced IP network services defines how the router recognizes and uses certain modules. When you set each member router's network services to **enhanced-ip**, only MPC/MIC modules and MS-DPC modules are powered on in the router. Non-service DPCs do not work with enhanced IP network services.

In Junos OS Release 11.4 and later releases prior to Release 13.2, the documentation for MX Series Virtual Chassis fails to mention the required procedures for enabling enhanced IP network services.

Use the following procedure to enable enhanced IP network services as part of the initial Virtual Chassis configuration. Perform these steps immediately after you create the preprovisioned member configuration on the master router, and before you enable graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) on both member routers.

To enable enhanced IP network services when you first set up an MX Series Virtual Chassis:

1. Configure enhanced IP network services on member 0.

- a. Log in to the console on member 0.
- b. Access the chassis hierarchy.

```
[edit]  
user@hostA# edit chassis
```

- c. Configure enhanced IP network services for member 0.

```
[edit chassis]  
user@hostA# set network-services enhanced-ip
```

- d. Commit the configuration on member 0 by using the **commit synchronize** command.



NOTE: Immediately after you commit the configuration, the software prompts you to reboot the router. You can proceed without rebooting the router at this point because a reboot occurs when you configure the member IDs to enable Virtual Chassis mode.

2. Configure enhanced IP network services on member 1.

- a. Log in to the console on member 1.
- b. Access the chassis hierarchy.

```
[edit]  
user@hostB# edit chassis
```

- c. Configure enhanced IP network services for member 1.

```
[edit chassis]
user@hostB# set network-services enhanced-ip
```

- d. Commit the configuration on member 1 by using the **commit synchronize** command.



NOTE: Immediately after you commit the configuration, the software prompts you to reboot the router. You can proceed without rebooting the router at this point because a reboot occurs when you configure the member IDs to enable Virtual Chassis mode.

3. (Optional) After the Virtual Chassis forms, verify that enhanced IP network services has been properly configured.

- a. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis master router (member0-re0).

```
{master:member0-re0}
user@hostA> show chassis network services
```

Network Services Mode: Enhanced-IP

- b. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis backup router (member1-re0).

```
{backup:member1-re0}
user@hostB> show chassis network services
```

Network Services Mode: Enhanced-IP

Use the following procedure to enable enhanced IP network services for an existing Virtual Chassis configuration.

To configure enhanced IP network services for an existing Virtual Chassis:

1. Log in to the console for the master Routing Engine in the Virtual Chassis master router (member0-re0).

2. Access the chassis hierarchy.

```
{master:member0-re0}[edit]
user@hostA# edit chassis
```

3. Configure enhanced IP network services on member 0.

```
{master:member0-re0}[edit chassis]
user@hostA# set network-services enhanced-ip
```

4. Commit the configuration by using the **commit synchronize** command.
5. When prompted to do so, reboot both Routing Engines in each member router forming the Virtual Chassis.

- For Junos OS Releases 11.4, 12.1, 12.2, 12.3R1, and 12.3R2:

```
{master:member0-re0}
user@hostA> request system reboot member 0 other-routing-engine
```

```
user@hostA> request system reboot member 1 other-routing-engine
user@hostA> request system reboot
```

- For Junos OS Release 12.3R3 and later releases:

```
{master:member0-re0}
user@hostA> request system reboot
```

Rebooting all Routing Engines in the Virtual Chassis propagates the enhanced IP network services configuration to both member routers.

6. (Optional) Verify that enhanced IP network services has been properly configured for the Virtual Chassis.

- a. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis master router (member0-re0).

```
{master:member0-re0}
user@hostA> show chassis network services
```

Network Services Mode: Enhanced-IP

- b. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis backup router (member1-re0).

```
{backup:member1-re0}
user@hostB> show chassis network services
```

Network Services Mode: Enhanced-IP

- 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.

.....

Infrastructure

- The following additional information regarding the configuration of peer IP addresses for ICCP peers and multichassis protection for MC-LAG applies to the *Configuring ICCP for MC-LAG* topic:

For Inter-Chassis Control (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer ip-address** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer ip-address** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0.0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

[*Network Interfaces, Ethernet Interfaces*]

- The following additional information regarding the behavior of the **accept-data** statement for MC-LAG in an active-active bridge domain applies to the *Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview* topic:

For a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain and with VRRP configured over an integrated routing and bridging (IRB) interface, you must include the **accept-data** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other Layer 3 protocols and applications to work properly over multichassis aggregated Ethernet (mc-aeX) interfaces.

[*Network Interfaces, Ethernet Interfaces*]

- The following additional information regarding the support of **vlan-id none** statement for MC-LAG applies to the *Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview* topic:

In an IPv6 network, you cannot configure a multichassis link aggregation group (MC-LAG) in an active-active bridge domain if you specified the **vlan-id none** statement at **[edit bridge-domain bd-name]** hierarchy level. The **vlan-id none** statement that enables the removal of the incoming VLAN tags identifying a Layer 2 logical interface when packets are sent over VPLS pseudowires is not supported for IPv6 packets in an MC-LAG.

[*Network Interfaces, Ethernet Interfaces*]

Interfaces and Chassis

- The *Redundancy Fabric Mode on Active Control Boards* subsection in the *Corrective Actions for Fabric Failures on MX Series Routers* topic and the *Configuring Redundancy Fabric Mode for Active Control Boards on MX Series Routers* topic incorrectly contain the following information about the default mode of redundant operation of active control boards on MX Series routers:

Until Junos OS Release 12.1, the MX Series routers that contain the enhanced Switch Control Board (SCB) with MX Series router with MPCs or MICs and the MPC3E, the control boards operate in redundancy fabric mode (all the FPCs use 4 fabric planes as active planes). Starting with Junos OS Release 12.2, on MX Series routers that contain the enhanced SCB with Trio chips and the MPC3E, the control boards operate in increased fabric bandwidth mode by default (all the available fabric planes are used).

The preceding description is incorrect because the enhanced SCB operates by default in redundancy fabric mode and not increased fabric mode in Junos OS Release 12.2 and later. The correct default operation of enhanced SCBs with Trio chips and the MPC3E is as follows:

The MX Series routers that contain the enhanced Switch Control Board (SCB) with Trio chips and the MPC3E, the control boards operate in redundancy fabric mode (all the FPCs use 4 fabric planes as active planes) by default.

[*System Basics, Chassis-Level Features*]

- The **redundancy-mode** configuration statement topic fails to state the following additional information regarding the default behavior for enhanced Switch Control Board (SCB) with Trio chips and the MPC3E on MX Series routers:

The MX Series routers that contain the enhanced Switch Control Board (SCB) with Trio chips and the MPC3E, the control boards operate in redundancy fabric mode (all the FPCs use 4 fabric planes as active planes) by default.

[*System Basics, Chassis-Level Features*]

- The following additional information regarding the binding of multiple port-mirror instances at the FPC level of M320 routers applies to the *Filter-Based Forwarding with Multiple Monitoring Interfaces* section in the *Configuring Port Mirroring* topic:

Because M320 routers do not support multiple bindings of port-mirror instances per FPC, the **port-mirror-instance** action is not supported in firewall filters for these routers.

[*Services Interfaces*]

- The **forwarding-mode (100-Gigabit Ethernet)** configuration statement topic fails to mention that this statement is supported on MX Series routers from Junos OS Release 12.1. The Supported Platforms section of this topic fails to list MX Series routers on which this command is supported.

[*Network Interfaces, Ethernet Interfaces*]

- The **show chassis fabric unreachable-destinations** command is incorrectly mentioned as supported on MX240, MX480, and MX960 routers from Junos OS Release 11.4R2 and Junos OS Release 12.1. The Supported Platforms section of this topic also incorrectly states MX240, MX480, and MX960 routers as supported routers for this command. This command is not available on the MX240, MX480, and MX960 routers. Instead, the correct command is the **show chassis fabric destinations** command, which you can use to view the state of fabric destinations for all FPCs.

[*System Basics and Services Command Reference*]

- The *Limiting traffic black-hole time by detecting Packet Forwarding Engine destinations that are unreachable over the fabric (MX240, MX480, and MX960 routers)* subsection under the *New Features in Junos OS Release for M Series, MX Series, and T Series Routers* main section of the Junos OS 11.4 Release Notes and Junos OS 12.1 Release Notes erroneously describes that the **show chassis fabric unreachable-destinations** command has been introduced. The correct command that has been introduced is the **show chassis fabric destinations** command, which is available in Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later.

[*Release Notes*]

- The following additional information regarding the working of unnumbered interfaces applies to the *Example: Configuring an Unnumbered Ethernet Interface* section in the *Configuring an Unnumbered Interface* topic:

The sample configuration that is described works correctly on M Series and T Series routers. For unnumbered interfaces on MX Series routers, you must additionally configure static routes on an unnumbered Ethernet interface by including the **qualified-next-hop** statement at the **[edit routing-options static route destination-prefix]** hierarchy level to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route.

[*Services Interfaces, Flow-Tap*]

- The following enhancements and additions apply to the *Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers* topic:

- The *Topology Diagram* section fails to mention that interface **ge-1/0/2** functions as the ICCP link between the two PE devices, interface **ge-1/1/1** is the ICL-PL link, and interface **ge-1/1/4** is the link that connects to the server or the MC-LAG client device.
- As a best practice, we recommend that you configure the ICCP and ICL interfaces over aggregated Ethernet interfaces instead of other interfaces such as Gigabit Ethernet interfaces, depending on your topology requirements and framework.
- You must disable RSTP on the ICL-PL interfaces for an MC-LAG in an active-active bridging domain.
- The *Step-by-Step Procedure* section for Router PE2 that is illustrated in the example is missing, although the quick configuration statements are presented.

To configure Router PE2:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE2# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae1
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE2# set ge-1/0/3 flexible-vlan-tagging
user@PE2# set ge-1/0/3 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/3 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/3 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/4 flexible-vlan-tagging
user@PE2# set ge-1/0/4 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/4 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/4 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae0
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae1
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```

```
[edit interfaces ae1]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1
```

```
[edit interfaces ae1 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 5
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active
```

```
[edit interfaces ae1 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 10
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE2 uses **chassis-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE1 uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE2# set domain-type bridge
user@PE2# set vlan-id all
user@PE2# set service-id 20
```

```
user@PE2# set interface ae0.0
user@PE2# set interface ae1.0
user@PE2# set interface ge-1/0/3.0
user@PE2# set interface ge-1/1/1.0
user@PE2# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE2# set local-ip-addr 100.100.100.2
user@PE2# set peer 100.100.100.1 redundancy-group-id-list 10
user@PE2# set peer 100.100.100.1 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE2# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

[*Network Interfaces, Ethernet Interfaces*]

- The following additional information regarding the compatibility of modules for the interoperation of RPM clients and RPM servers applies to the *Configuring RPM Probes* section in the *Configuring Real-Time Performance Monitoring* topic:

Keep the following points in mind when you configure RPM clients and RPM servers:

- You cannot configure an RPM client that is PIC-based and an RPM server that is based on either the Packet Forwarding Engine or Routing Engine to receive the RPM probes.
- You cannot configure an RPM client that is Packet Forwarding Engine-based and an RPM server that receives the RPM probes to be on the PIC or Routing Engine.
- The RPM client and RPM server must be located on the same type of module. For example, if the RPM client is PIC-based, the RPM server must also be PIC-based, and if the RPM server is Packet Forwarding Engine-based, the RPM client must also be Packet Forwarding Engine-based.

[*System Basics, Chassis-Level Features*]

- The following corrections apply to the *Example: Configuring One VPLS Instance for Several VLANs* topic:

The following sentence is erroneously presented:

If VLANs 1 through 1000 for customer C1 span the same sites, then the **vlan-id all** and **vlan-id-list-range** statements provide a way to switch all of these VLANs with a minimum configuration effort and fewer switch resources.

The correct description is as follows:

If VLANs 1 through 1000 for customer C1 span the same sites, then the **vlan-id all** and **vlan-id-list** statements provide a way to switch all of these VLANs with a minimum configuration effort and fewer switch resources.

The following example replaces the existing example that illustrates the use of the **vlan-id all** statement:

```
[edit]
interfaces ge-1/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-1000; # Note the use of the VLAN id list statement.
    }
  }
}
unit 11 {
  encapsulation vlan-vpls;
  family bridge {
    interface-mode trunk;
    vlan-id-list 1500;
  }
}
interfaces ge-2/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  unit 1 {
    encapsulation vlan-vpls;
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-1000; # Note the use of the VLAN id list statement.
    }
  }
}
interfaces ge-3/0/0 {
  encapsulation flexible-ethernet-services;
  flexible-vlan-tagging;
  family bridge {
    unit 1 {
      encapsulation vlan-vpls;
      interface-mode trunk;
      vlan-id-list 1-1000; # Note the use of the VLAN id list statement.
    }
  }
}
interfaces ge-6/0/0 {
  encapsulation flexible-ethernet-services;
```

```

flexible-vlan-tagging;
family bridge {
  unit 11 {
    encapsulation vlan-vpls;
    interface-mode trunk;
    vlan-id-list 1500;
  }
}
routing-instances {
  customer-c1-virtual-switch {
    instance-type virtual-switch;
    interface ge-1/0/0.1;
    interface ge-2/0/0.1;
    interface ge-3/0/0.1;
    bridge-domains {
      c1-vlan-v1-to-v1000 {
        vlan-id all; # Note the use of the VLAN id all statement
      }
    }
  } # End of customer-c1-v1-to-v1000
  customer-c2-virtual-switch {
    instance-type virtual-switch;
    interface ge-1/0/0.11;
    interface ge-6/0/0.11;
    bridge-domains {
      c1-vlan-v1500 {
        vlan-id all; # Note the use of the VLAN id all statement
      }
    }
  } # End of customer-c1-v1500
} # End of routing-instances

```

Note the use of the **vlan-id all** statement in the virtual-switch instance called **customer-c1-v1-to-v1000**.

[*MX Series Ethernet Services Routers Solutions Guide*]

- The *Junos OS 12.3 Release Notes* did not report that the default bandwidth and burst values for DDoS protection changed for the **aggregate**, **host**, **pfe**, **syslog**, and **tap** packet types in the **sample** protocol group. In releases before Junos OS Release 12.3, the default bandwidth value was 10,000 pps and the default burst size value was 10,000 packets. Beginning in Junos OS Release 12.3, the default values are 1000 pps for bandwidth and 1000 packets for burst size.
- The *open-timeout* configuration statement topic and the *Configuring Default Timeout Settings for Services Interfaces* topic incorrectly state that the default value of the timeout period for TCP session establishment is 30 seconds. The correct default value is 5 seconds.

[*System Basics, Chassis-Level Features*]

- The Supported Platforms section of the *set chassis display message* command topic erroneously states that this command is supported on MX Series routers. This command is not available on MX Series routers.

[*System Basics, Chassis-Level Features*]

- The *IP Demux Interfaces over Static or Dynamic VLAN Demux Interfaces* topic incorrectly states that both DPCs and MPCs support VLAN demux subscriber interfaces. In fact, only MPCs support these interfaces.

Junos OS XML API and Scripting

- The *NETCONF XML Management Protocol Guide* incorrectly states that when performing a confirmed commit operation using the `<commit>` element, the `<confirm-timeout>` value specifies the number of minutes for the rollback deadline. The value of the `<confirm-timeout>` element actually specifies the number of seconds for the rollback deadline.

[*NETCONF XML Management Protocol Guide*]

Layer 2 Ethernet Services

- The following information regarding the differences in the default limit on MAC addresses that can be learned on an access port and a trunk port is inadvertently omitted from the *Limiting MAC Addresses Learned from an Interface in a Bridge Domain* topic:
 - For an access port, the default limit on the maximum number of MAC addresses that can be learned on an access port is 1024. Because an access port can be configured in only one bridge domain in a network topology, the default limit is 1024 addresses, which is same as the limit for MAC addresses learned on a logical interface in a bridge domain (configured by including the `interface-mac-limit limit` statement at the `[edit bridge-domains bridge-domain-name bridge-options interface interface-name]` or `[edit bridge-domains bridge-domain-name bridge-options]` hierarchy level.
 - For a trunk port, the default limit on the maximum number of MAC addresses that can be learned on a trunk port is 8192. Because a trunk port can be associated with multiple bridge domains, the default limit is the same as the limit for MAC addresses learned on a logical interface in a virtual switch instance (configured by including the `interface-mac-limit limit` statement at the `[edit routing-instances routing-instance-name switch-options interface interface-name]` for a virtual switch instance).

MPLS

- The "Configuring Miscellaneous LDP Properties," "Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols," "authentication-key-chain (LDP)," and "authentication-key-chain (BGP and BMP)" topics should include the following information: You must also configure the authentication algorithm using the `authentication-algorithm algorithm` statement. This statement must be included at the `[edit protocols (bgp | ldp)]` hierarchy level when you configure the `authentication-key-chain key-chain` statement at the `[edit protocols (bgp | ldp)]` hierarchy level.

[*MPLS*]

- The "Path Computation for LSPs on an Overloaded Router" topic should state that when you set the overload bit on a router running IS-IS, only new LSPs are prevented from transiting through the router. Any existing Constrained Path Shortest First (CPSF) LSPs remain active and continue to transit through the router. The documentation incorrectly states that any existing LSPs transiting through the router are also rerouted when you configure the overload bit on an IS-IS router.

Network Management and Monitoring

- The documentation fails to clearly describe the characters that can be used for SNMPv3 authentication passwords. Besides numbers, uppercase letters, and lowercase letters, the following special characters are supported:

,./\<>;:'[]{}~!@#\$%^*_+=-`

In addition, the following special characters are also supported, but you must enclose them within quotation marks ("") if you enter them on the CLI; if you use a Network Management System to enter the password, the quotation marks are not required:

| & () ?

The documentation also fails to clearly state that characters entered by simultaneously pressing the Ctrl key and additional keys are not supported. [PR/883083: This issue has been resolved]

- The syntax of the **filter-interfaces** statement in the *SNMP Configuration Statement* section is incorrect. The correct syntax is as follows:

```
filter-interfaces {
  all-internal-interfaces;
  interfaces interface-names{
    interface 1;
    interface 2;
  }
}
```

Routing Policy and Firewall Filters

- The *OSPF Configuration Guide* incorrectly includes the **transmit-interval** statement at the **[edit protocols ospf area area interface interface-name]** hierarchy level. The **transmit-interval** statement at this hierarchy level is deprecated in the Junos OS command-line interface.

[*OSPF Configuration Guide*]

Routing Protocols

- In routing instances, when a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away. [*Routing Protocols*]

Services Applications

- The **show services stateful-firewall subscriber-analysis** command should be included in the *System Basics and Services Command Reference Guide*. This command displays information about the number of active subscribers on the service physical interface card (PIC).
- The **show services stateful-firewall flow-analysis** command should be included in the *System Basics and Services Command Reference Guide*. This command displays stateful firewall flow statistics.”
- In the *Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions Guide*, the section “Configuring Address Pools for Network Address Port Translation” should be revised as follows: The following variables should be added
Nr_Addr_PR_Prefix – Number of usable pre-NAT IPv4 subscriber addresses in a “from” clause match condition
Nr_Addr_PU_Prefix – Number of usable post-NAT IPv4 addresses configured in the NAT pool
Rounded_Port_Range_Per_IP = $\text{ceil}[(\text{Nr_Addr_PR_Prefix}/\text{Nr_Addr_PU_Prefix})]$
* Block_Size The Forward Translation formulas should be: 1. $\text{Pr_Offset} = \text{Pr_Prefix} - \text{Base_Pr_Prefix}$ 2. $\text{Pr_Port_Offset} = \text{Pr_Offset} * \text{Block_Size}$ 3. $\text{Rounded_Port_Range_Per_IP} = \text{ceil}[(\text{Nr_Addr_PR_Prefix}/\text{Nr_Addr_PU_Prefix})] * \text{Block_Size}$ 4. $\text{Pu_Prefix} = \text{Base_Public_Prefix} + \text{floor}(\text{Pr_Port_Offset}/\text{Rounded_Port_Range_Per_IP})$ 5. $\text{Pu_Start_Port} = \text{Pu_Port_Range_Start} + (\text{Pr_Port_Offset} \% \text{Rounded_Port_Range_Per_IP})$
The Reverse Translation formulas should be: 1. $\text{Pu_Offset} = \text{Pu_Prefix} - \text{Base_Pu_Prefix}$ 2. $\text{Pu_Port_Offset} = (\text{Pu_Offset} * \text{Rounded_Port_Range_Per_IP}) + (\text{Pu_Actual_Port} - \text{Pu_Port_Range_Start})$ 3. $\text{Subscriber_IP} = \text{Base_Pr_Prefix} + \text{floor}(\text{Pu_Port_Offset} / \text{Block_Size})$
- The following information should be added to the syntax of the “service-set (Services)” configuration statement topic in the *Services Interfaces Configuration Guide*. This information should appear under the **service-set service-set-name** level:

```
service-set-options {  
  bypass-traffic-on-exceeding-flow-limits;  
  bypass-traffic-on-pic-failure;  
  enable-asymmetric-traffic-processing;  
  support-uni-directional-traffic;  
}
```


This issue was being tracked by PR888803.

- The following information should replace Table 1 and the section “Sample Output” in the “show services stateful-firewall statistics” topic in the *System Basics and Services Command Reference*:

Table 5: show services stateful-firewall statistics Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
New flows	Rule match counters for new flows: <ul style="list-style-type: none"> • Rule Accepts—New flows accepted. • Rule Discards—New flows discarded. • Rule Rejects—New flows rejected.
Existing flow types packet counters	Rule match counters for existing flows: <ul style="list-style-type: none"> • Accepts—Match existing forward or watch flow. • Drop—Match existing discard flow. • Rejects—Match existing reject flow.
Hairpinning Counters	Hairpinning counters: <ul style="list-style-type: none"> • Slow Path Hairpinned Packets—Slow path packets that were hairpinned back to the internal network. • Fast Path Hairpinned Packets—Fast path packets that were hairpinned back to the internal network.
Drops	Drop counters: <ul style="list-style-type: none"> • IP option—Packets dropped in IP options processing. • TCP SYN defense—Packets dropped by SYN defender. • NAT ports exhausted—Hide mode. The router has no available Network Address Translation (NAT) ports for a given address or pool. • Sessions dropped due to subscriber flow limit—Sessions dropped because the subscriber's flow limit was exceeded.
Errors	Total errors, categorized by protocol: <ul style="list-style-type: none"> • IP—Total IP version 4 errors. • TCP—Total Transmission Control Protocol (TCP) errors. • UDP—Total User Datagram Protocol (UDP) errors. • ICMP—Total Internet Control Message Protocol (ICMP) errors. • Non-IP packets—Total non-IPv4 errors. • ALG—Total application-level gateway (ALG) errors.

Table 5: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
IP Errors	<p>IPv4 errors:</p> <ul style="list-style-type: none"> • IP packet length inconsistencies—IP packet length does not match the Layer 2 reported length. • Minimum IP header length check failures—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes. • Reassembled packet exceeds maximum IP length—After fragment reassembly, the reassembled IP packet length exceeds 65,535. • Illegal source address 0—Source address is not a valid address. Invalid addresses are, loopback, broadcast, multicast, and reserved addresses. Source address 0, however, is allowed to support BOOTP and the destination address 0xffffffff. • Illegal destination address 0—Destination address is not a valid address. The address is reserved. • TTL zero errors—Received packet had a time-to-live (TTL) value of 0. • Illegal IP protocol number (0 or 255)—IP protocol is 0 or 255. • Land attack—IP source address is the same as the destination address. • Non-IPv4 packets—Packet was not IPv4. (Only IPv4 is supported.) • Bad checksum—Packet had an invalid IP checksum. • Illegal IP fragment length—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • IP fragment overlap—Fragments have overlapping fragment offsets. • IP fragment reassembly timeout—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. • IP fragment limit exceeded: 0—Fragments that exceeded the limit. • Unknown: 0—Unknown fragments.

Table 5: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
TCP Errors	

Table 5: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	TCP protocol errors:
	<ul style="list-style-type: none"> • TCP header length inconsistencies—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes. • Source or destination port number is zero—TCP source or destination port is zero. • Illegal sequence number and flags combinations — Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set. • SYN attack (multiple SYN messages seen for the same flow)—Multiple SYN packets received for the same flow are treated as a SYN attack. The packets might be retransmitted SYN packets and therefore valid, but a large number is cause for concern. • First packet not a SYN message—First packets for a connection are not SYN packets. These packets might originate from previous connections or from someone performing an ACK/FIN scan. • TCP port scan (TCP handshake, RST seen from server for SYN)—In the case of a SYN defender, if an RST (reset) packet is received instead of a SYN/ACK message, someone is probably trying to scan the server. This behavior can result in false alarms if the RST packet is not combined with an intrusion detection service (IDS). • Bad SYN cookie response—SYN cookie generates a SYN/ACK message for all incoming SYN packets. If the ACK received for the SYN/ACK message does not match, this counter is incremented. • TCP reconstructor sequence number error—This counter is incremented in the following cases: The TCP seqno is 0 and all the TCP flags are also 0. The TCP seqno is 0 and FIN/PSH/URG TCP flags are set. • TCP reconstructor retransmissions—This counter is incremented for the retransmitted packets during connection 3-way handshake. • TCP partially opened connection timeout (SYN)—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder. • TCP partially opened connection timeout (SYN-ACK)—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder. • TCP partially closed connection reuse—Not supported. • TCP 3-way error - client sent SYN+ACK—A SYN/ACK should be sent by the server on receiving a SYN. This counter is incremented when the first message received from the initiator is SYN+ACK. • TCP 3-way error - server sent ACK—ACK should be sent by the client on receiving a SYN/ACK from the server. This counter is incremented when the ACK is received from the Server instead of from the Client. • TCP 3-way error - SYN seq number retransmission mismatch—This counter is incremented when the SYN is received again with a different sequence number from the first SYN sequence number. • TCP 3-way error - RST seq number mismatch—A reset could be received from either side. The server could send a RST on receiving a SYN or the client could send a RST on receiving SYN/ACK. This counter is incremented when the RST is received either from the client or server with a non-matching sequence

Table 5: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	<p>number.</p> <ul style="list-style-type: none"> • TCP 3-way error - FIN received—This counter is incremented when the FIN is received during the 3-way handshake. • TCP 3-way error - invalid flags (PSH, URG, ECE, CWR)—This counter is incremented when any of the PSH, URG, ECE, or CWR flags were received during the 3-way handshake. • TCP 3-way error - SYN recvd but no client flows—This counter is incremented when SYN is received but not from the connection initiator. The counter is not incremented in the case of simultaneous open, when the SYN is received in both the directions. • TCP 3-way error - first packet SYN+ACK—The first packet received was SYN+ACK instead of SYN. • TCP 3-way error - first packet FIN+ACK—The first packet received was FIN+ACK instead of SYN. • TCP 3-way error - first packet FIN—The first packet received was FIN instead of SYN. • TCP 3-way error - first packet RST—The first packet received was RST instead of SYN. • TCP 3-way error - first packet ACK—The first packet received was ACK instead of SYN. • TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR)—The first packet received had invalid flags. • TCP Close error - no final ACK—This counter is incremented when ACK is not received after the FINs are received from both directions. • TCP Resumed Flow—Plain ACKs create flows if rule match permits, and these are classified as TCP Resumed Flows. This counter is incremented in the case of a TCP Resumed Flow.
UDP Errors	<p>UDP protocol errors:</p> <ul style="list-style-type: none"> • IP data length less than minimum UDP header length (8 bytes)—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes. • Source or destination port is zero—UDP source or destination port is 0. • UDP port scan (ICMP error seen for UDP flow)—ICMP error is received for a UDP flow. This could be a genuine UDP flow, but it is counted as an error.
ICMP Errors	<p>ICMP protocol errors:</p> <ul style="list-style-type: none"> • IP data length less than minimum ICMP header length (8 bytes)—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes. • ICMP error length inconsistencies—Minimum length of an ICMP error packet is 48 bytes, and the maximum length is 576 bytes. This counter is incremented when the received ICMP error falls outside this range. • Duplicate ping sequence number—Received ping packet has a duplicate sequence number. • Mismatched ping sequence number—Received ping packet has a mismatched sequence number. • No matching flow—No matching existing flow was found for the ICMP error.

Table 5: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
ALG errors	<p>Accumulation of all the application-level gateway protocol (ALG) drops counted separately in the ALG context:</p> <ul style="list-style-type: none"> • BOOTP—Bootstrap protocol errors • DCE-RPC—Distributed Computing Environment-Remote Procedure Call protocols errors • DCE-RPC portmap—Distributed Computing Environment-Remote Procedure Call protocols portmap service errors • DNS—Domain Name System protocol errors • Exec—Exec errors • FTP—File Transfer Protocol errors • H323—H.323 standards errors • ICMP—Internet Control Message Protocol errors • IIOB—Internet Inter-ORB Protocol errors • Login—Login errors • NetBIOS—NetBIOS errors • Netshow—NetShow errors • Real Audio—RealAudio errors • RPC—Remote Procedure Call protocol errors • RPC portmap—Remote Procedure Call protocol portmap service errors • RTSP—Real-Time Streaming Protocol errors • Shell—Shell errors • SIP—Session Initiation Protocol errors • SNMP—Simple Network Management Protocol errors • SQLNet—SQLNet errors • TFTP—Trivial File Transfer Protocol errors • Traceroute—Traceroute errors
Drop Flows	<ul style="list-style-type: none"> • Maximum Ingress Drop flows allowed—Maximum number of ingress flow drops allowed. • Maximum Egress Drop flows allowed—Maximum number of egress flow drops allowed. • Current Ingress Drop flows—Current number of ingress flow drops. • Current Egress Drop flows—Current number of egress flow drops. • Ingress Drop Flow limit drops count—Number of ingress flow drops due to maximum number of ingress flow drops being exceeded. • Egress Drop Flow limit drops count—Number of egress flow drops due to maximum number of egress flow drops being exceeded.

```
user@host> show services stateful-firewall statistics extensive
Interface: ms-1/3/0
Service set: interface-svc-set
New flows:
  Rule Accepts: 907, Rule Discards: 0, Rule Rejects: 0
Existing flow types packet counters:
  Accepts: 3535, Drop: 0, Rejects: 0
Hairpinning counters:
  Slow Path Hairpinned Packets: 0, Fast Path Hairpinned Packets: 0
Drops:
  IP option: 0, TCP SYN defense: 0
  NAT ports exhausted: 0, Sessions dropped due to subscriber flow limit: 0

Errors:
  IP: 0, TCP: 0
  UDP: 0, ICMP: 0
  Non-IP packets: 0, ALG: 0
IP errors:
  IP packet length inconsistencies: 0
  Minimum IP header length check failures: 0
  Reassembled packet exceeds maximum IP length: 0
  Illegal source address: 0
  Illegal destination address: 0
  TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
  Land attack: 0
  Non-IPv4 packets: 0, Bad checksum: 0
  Illegal IP fragment length: 0
  IP fragment overlap: 0
  IP fragment reassembly timeout: 0
  IP fragment limit exceeded: 0
  Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combination: 0
  SYN attack (multiple SYN messages seen for the same flow): 0
  First packet not a SYN message: 0
  TCP port scan (TCP handshake, RST seen from server for SYN): 0
  Bad SYN cookie response: 0
  TCP reconstructor sequence number error: 0
  TCP reconstructor retransmissions: 0
  TCP partially opened connection timeout (SYN): 0
  TCP partially opened connection timeout (SYN-ACK): 0
  TCP partially closed connection reuse: 0
  TCP 3-way error - client sent SYN+ACK: 0
  TCP 3-way error - server sent ACK: 0
  TCP 3-way error - SYN seq number retransmission mismatch: 0
  TCP 3-way error - RST seq number mismatch: 0
  TCP 3-way error - FIN received: 0
  TCP 3-way error - invalid flags (PSH, URG, ECE, CWR): 0
  TCP 3-way error - SYN recvd but no client flows: 0
  TCP 3-way error - first packet SYN+ACK: 0
  TCP 3-way error - first packet FIN+ACK: 0
  TCP 3-way error - first packet FIN: 0
  TCP 3-way error - first packet RST: 0
  TCP 3-way error - first packet ACK: 0
  TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR): 0
  TCP Close error - no final ACK: 0
  TCP Resumed Flow: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
```

```
Source or destination port is zero: 0
UDP port scan (ICMP error seen for UDP flow): 0
ICMP errors:
  IP data length less than minimum ICMP header length (8 bytes): 0
  ICMP error length inconsistencies: 0
  Duplicate ping sequence number: 0
  Mismatched ping sequence number: 0
  No matching flow: 0
ALG errors:
  BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
  DNS: 0, Exec: 0, FTP: 0
  H323: 0, ICMP: 0, IIOP: 0
  Login: 0, NetBIOS: 0, Netshow: 0
  Real Audio: 0, RPC: 0, RPC portmap: 0
  RTSP: 0, Shell: 0, SIP: 0
  SNMP: 0, SQLNet: 0, TFTP: 0
  Traceroute: 0
Drop Flows:
  Maximum Ingress Drop flows allowed: 20
  Maximum Egress Drop flows allowed: 20
  Current Ingress Drop flows: 0
  Current Egress Drop flows: 0
  Ingress Drop Flow limit drops count: 0
  Egress Drop Flow limit drops count: 0

**If max-drop-flows is not configured, the following is shown**
Drop Flows:
  Maximum Ingress Drop flows allowed: Default
  Maximum Egress Drop flows allowed: Default
```

- The following information should be added after the second paragraph of the “Configuring Inline Sampling” topic in the *Services Interfaces Configuration Guide*:

The following limitations exist for inline sampling:

- Flow records and templates cannot be exported if the flow collector is reachable through any management interface.
- The flow collector should be reachable through the default routing table (inet.0 or inet6.0). If the flow collector is reachable via a non-default VPN routing and forwarding table (VRF), flow records and templates cannot be exported.
- If the destination of the sampled flow is reachable through multiple paths, the IP_NEXT_HOP (Element ID 15) and OUTPUT_SNMP (Element ID 14) in the IPv4 flow record would be set to the Gateway Address and SNMP Index of the first path seen in the forwarding table.
- If the destination of the sampled flow is reachable through multiple paths, the IP_NEXT_HOP (Element ID 15) and OUTPUT_SNMP (Element ID 14) in the IPv6 flow records would be set to 0.
- The user-defined sampling instance gets precedence over the global instance. When a user-defined sampling instance is attached to the FPC, the global instance is removed from the FPC and the user-defined sampling instance is applied to the FPC.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If OIF is in a different VRF, DST_MASK (Element ID 13), DST_AS (Element ID

17), IP_NEXT_HOP (Element ID 15), and OUTPUT_SNMP (Element ID 14) would be set to 0 in the flow records.

- Each Lookup Chip (LU) maintains and exports flows independent of other LUs. Traffic received on a media interface is distributed across all LUs in a multi-LU platform. It is likely that a single flow will be processed by multiple LUs. Therefore, each LU creates a unique flow and exports it to the flow collector. This can cause duplicate flows records to be seen on the flow collector. The flow collector should aggregate PKTS_COUNT and BYTES_COUNT for duplicate flow records to derive a single flow record.

This issue is being tracked by PR907991

- The *System Basics and Services Command Reference* should include the following commands in the chapter “Dynamic Application Awareness Operational Mode Commands”:

request services application-identification application: Copy, disable, or enable a predefined application signature.

request services application-identification group: Copy, disable, or enable a predefined application signature group.

show services application-identification application: Display detailed information about a specified application signature, all application signatures, or a summary of the existing application signatures and nested application signatures. Both custom and predefined application signatures and nested application signatures can be displayed.

show services application-identification group: Display detailed or summary information about a specified application signature group or all application signature groups. Both custom and predefined application signature groups can be displayed.

show services application-identification version: Display the Junos OS application package version.

- The following command should appear in the network address operational mode commands:

```
clear services nat statistics
<interface interface-name>
<service-set service-set-name>
```

The <**interface *interface-name***> option clears NAT statistics for the specified interface only.

The <**service-set *service-set-name***> option clears NAT statistics for the specified service set only.

The **clear services inline nat statistics** command should include the following option:

```
<interface interface-name>
```

The <**interface *interface-name***> option clears inline NAT statistics for the specified interface only.

- The following additional information regarding the interoperation of sample actions in firewall filters and traffic sampling applies to the *Minimum Configuration for Traffic Sampling* section in the *Configuring Traffic Sampling* topic:

The following prerequisites apply to M Series, MX Series, and T Series routers when you configure traffic sampling on interfaces and in firewall filters:

- If you configure a sample action in a firewall filter for an inet or inet6 family on an interface without configuring the forwarding-options settings, operational problems might occur if you also configure port mirroring or flow-tap functionalities. In such a scenario, all the packets that match the firewall filter are incorrectly sent to the service PIC.
- If you include the **then sample** statement at the **[edit firewall family inet filter *filter-name* term *term-name*]** hierarchy level to specify a sample action in a firewall filter for IPv4 packets, you must also include the **family inet** statement at the **[edit forwarding-options sampling]** hierarchy level or the **instance *instance-name* family inet** statement at the **[edit forwarding-options sampling]** hierarchy level. Similarly, if you include the **then sample** statement at the **[edit firewall family inet6 filter *filter-name* term *term-name*]** hierarchy level to specify a sample action in a firewall filter for IPv6 packets, you must also include the **family inet6** statement at the **[edit forwarding-options sampling]** hierarchy level or the **instance *instance-name* family inet6** statement at the **[edit forwarding-options sampling]** hierarchy level. Otherwise, a commit error occurs when you attempt to commit the configuration.
- If you configure traffic sampling on a logical interface by including the sampling input or sampling output statements at the **[edit interface *interface-name* unit *logical-unit-number*]** hierarchy level, you must also include the **family inet | inet6** statement at the **[edit forwarding-options sampling]** hierarchy level, or the **instance *instance-name* family inet | inet6** statement at the **[edit forwarding-options sampling]** hierarchy level.
- The *Configuring Port Mirroring* topic erroneously states that the **input** statement can be included under the **[edit forwarding-options port-mirroring family (inet | inet6) output]** hierarchy level. Only the **output** statement is available at the **[edit forwarding-options port-mirroring family (inet | inet6)]** hierarchy level. To configure the input packet properties for port mirroring, you must include the **input** statement at the **[edit forwarding-options port-mirroring]** hierarchy level.

To configure port mirroring on a logical interface, configure the following statements at the **[edit forwarding-options port-mirroring]** hierarchy level:

```
[edit forwarding-options port-mirroring]
input {
  maximum-packet-length bytes
  rate rate;
  run-length number;
}
family (inet|inet6) {
  output {
    interface interface-name {
      next-hop address;
    }
  }
}
```

```

no-filter-check;
}
}

```

Also, the note incorrectly states that the **input** statement can also be configured at the **[edit forwarding-options port-mirroring]** hierarchy level and that it is only maintained for backward compatibility. The note also mentions that the configuration of the **output** statement is deprecated at the **[edit forwarding-options port-mirroring]** hierarchy level.

The correct behavior regarding the port-mirroring configuration for the packets to be mirrored and for the destination at which the packets are to be received is as follows:



NOTE: The **input** statement is deprecated at the **[edit forwarding-options port-mirroring family (inet | inet6)]** hierarchy level and is maintained only for backward compatibility. You must include the **input** statement at the **[edit forwarding-options port-mirroring]** hierarchy level.

[Services Interfaces, Port Mirroring]

- In the *Output Fields* section of the **show services ipsec-vpn ipsec security-associations** command topic, the descriptions of the **Local Identity** and **Remote Identity** fields are not clear and complete. The following are the revised descriptions of these fields:
 - **Local Identity**—Protocol, address or prefix, and port number of the local entity of the IPsec association. The format is **id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)**. The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the **id-data-len** parameter can be one of the following, depending on the address configured in the IPsec rule:
 - For an IPv4 address, the length is 4 and the value displayed is 3.
 - For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7.
 - For a range of IPv4 addresses, the length is 8 and the value displayed is 7.
 - For an IPv6 address prefix, the length is 16 and the value displayed is 15.
 - For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31.
 - For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31.

The value of the **id-data-presentation** field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the local peer of the IPsec association, it is displayed instead of the address details.

- **Remote Identity**—Protocol, address or prefix, and port number of the remote entity of the IPsec association. The format is **id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)**. The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not

user-configurable in the IPsec rule. The value of the **id-data-len** parameter can be one of the following, depending on the address configured in the IPsec rule:

- For an IPv4 address, the length is 4 and the value displayed is 3.
- For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7.
- For a range of IPv4 addresses, the length is 8 and the value displayed is 7.
- For an IPv6 address prefix, the length is 16 and the value displayed is 15.
- For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31.
- For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31.

The value of the **id-data-presentation** field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the remote peer of the IPsec association, it is displayed instead of the address details.

[*Services Interfaces, IPsec Properties, Junos VPN Site Secure*]

- The *Supported Flow Monitoring and Discard Accounting Standards* topic fails to mention the following additional information:

On MX Series routers, Junos OS partially supports the following RFCs:

- RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
- RFC 5102, *Information Model for IP Flow Information Export*

[*Junos OS Supported Standards*]

- 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.

[*Services Interfaces, Flow Monitoring*]

- 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 change to the NAT pool name.
- 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 “Configuring Unicast Tunnels” topic incorrectly shows the backup-destination statement. This statement does not apply to unicast tunnels and should be removed.

Subscriber Access Management

- In the *Subscriber Access Configuration Guide*, there is an error in the *Example: Configuring RADIUS-Based Subscriber Authentication and Accounting* topic. In the example, the **profile** stanza incorrectly includes the **authentication** statement. The correct statement is **authentication-order**, as shown in the following sample:

```
profile isp-bos-metro-fiber-basic {
  authentication-order radius;
}
```

[Subscriber Access]

- The *L2TP for Subscriber Access Overview* topic in the *Junos OS Subscriber Access Configuration Guide* incorrectly states that L2TP is supported only on MX240, MX480, and MX960 routers. In fact, support for MX80 routers was added in Junos OS Release 12.3. In that release and later releases, the MX80 supports all L2TP features that were supported on the MX240, MX480, and MX960 routers as of Junos OS Release 11.4.

[Subscriber Access]

- The *MX Series 3D Universal Edge Router Interface Module Reference* does not state that VLAN demux configurations are not supported on MX Series routers that have any of the following line cards installed:

- Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q)
- Enhanced Queuing IP Services DPCs (DPCE-R-Q)

The nonsupport includes any configuration stacked on top of a VLAN demux. For example, although PPPoE is supported, PPPoE over aggregated Ethernet interfaces is not supported when one of these cards is installed, because this configuration requires PPPoE to be stacked on a VLAN demux.

- In the *AAA Service Framework Feature Guide for Subscriber Management*, the **parse-direction** (Domain Map) statement and the *Specifying the Parsing Direction for Domain Names* topic show an incorrect default setting for the **parse-direction** statement. The correct default is the **left-to-right** direction.
- In the *Junos OS Subscriber Access Feature Guide*, the **fail-over-within-preference** statement at the **[edit services l2tp]** hierarchy level is incorrectly spelled. The correct spelling for this statement is **failover-within-preference**.
- The *Example: HTTP Service Within a Service Set* topic in the *Subscriber Access Configuration Guide* erroneously describes how to configure captive portal content delivery rules in service sets.

Use the following procedure to configure captive portal content delivery rules in service sets:

1. Define one or more rules with the **rule** *rule-name* statement at the **[edit services captive-portal-content-delivery]** hierarchy level. In each rule you specify one or more terms to match on an application, destination address, or destination prefix list; where the match takes place; and actions to be taken when the match occurs,
 2. (Optional) Define one or more rule sets by listing the rules to be included in the set with the **rule-set** *rule-set-name* statement at the **[edit services captive-portal-content-delivery]** hierarchy level.
 3. Configure a captive portal content delivery profile with the **profile** *profile-name* statement at the **[edit services captive-portal-content-delivery]** hierarchy level.
 4. In the profile, specify a list of rules with the **cpcd-rules** *[rule-name]* statement or a list of rule sets with the **cpcd-rule-sets** *[rule-set-name]* statement. Both statements are at the **[edit services captive-portal-content-delivery profile profile-name]** hierarchy level.
 5. Associate the profile with a service set with the **captive-portal-content-delivery-profile** *profile-name* statement at the **[edit services service-set service-set-name]** hierarchy level.
- The *LAC Tunnel Selection Overview* topic in the *Subscriber Access Configuration Guide* incorrectly describes the current behavior for failover between preference levels. The topic states that when the tunnels at every preference level have a destination in the lockout state, the LAC cycles back to the highest preference level and waits for the lockout time for a destination at that level to expire before attempting to connect and starting the process over.

In fact, the current behavior in this situation is that from the tunnels present at the lowest level of preference (highest preference number), the LAC selects the tunnel that has the destination with the shortest remaining lockout time. The LAC ignores the lockout and attempts to connect to the destination.

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

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

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

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

User Access and Authentication

- The "Example: DHCP Complete Configuration" and "dhcp" topics should not include support for the MX Series Universal Edge 3D Routers. This feature is supported only on the M Series and the T Series.

User Access and Authentication

User Interface and Configuration

- For the **set system login format** command, the **des** option has been deprecated.
- The output for **show configuration | display inheritance | display set** now displays the set commands needed to duplicate the fully inherited configuration.

VPNs

- The following guideline regarding the support of LSI traffic statistics on M Series routers is missing from the *General Limitations on IP-Based Filtering* section in the *Filtering Packets in Layer 3 VPNs Based on IP Headers* topic:

Label-switched interface (LSI) traffic statistics are not supported for Intelligent Queuing 2 (IQ2), Enhanced IQ (IQE), and Enhanced IQ2 (IQ2E) PICs on M Series routers.

[VPNs, Layer 3 VPNs]

- The following limitation regarding firewall filters configured in conjunction with the **vrf-table-label** statement is missing from the *General Limitations on IP-Based Filtering* in the *Filtering Packets in Layer 3 VPNs Based on IP Headers* topic:

Firewall filters cannot be applied to interfaces included in a routing instance on which you have configured the **vrf-table-label** statement.

This documentation is applicable to all M Series, T Series, and SRX Series routers.

[VPNs, Layer 3 VPNs]

- The descriptions of the **pw-label-ttl-1** and **router-alert-label** options in the *control-channel (Protocols OAM)* configuration statement topic are incorrectly and interchangeably stated. The correct descriptions of these options are as follows:

- **pw-label-ttl-1**—For BGP-based pseudowires that send OAM packets with the MPLS pseudowire label and time-to-live (TTL) set to 1.
- **router-alert-label**—For BGP-based pseudowires that send OAM packets with router alert label.

[VPNs, Layer 2 VPNs]

Changes to the Junos OS Documentation Set

The following are the changes made to the Junos OS documentation set:

- A new topic, *CGN Implementation: Best Practices*, which provides experience-based recommendations for configuring carrier-grade NAT, has been added to the documentation set. The new topic is available at http://www.juniper.net/techpubs/en_US/junos12.3/topics/concept/nat-best-practices.html
- ALG documentation for MX Series platforms has been updated. The topic has been reorganized and expanded, with particular emphasis on SIP and SIP-NAT interaction. An updated version of the documentation is available at the following PR link location: [PR817816](#)

Related Documentation

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)
- [Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 181](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 357](#)

Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

The current software release is Release 12.3. For information about obtaining the software packages, see “[Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers](#)” on page 357.

- [Authentication and Access Control](#)
- [Class of Service \(CoS\)](#)
- [Forwarding and Sampling](#)
- [General Routing](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)

- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [MPLS](#)
- [Multicast](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Services Applications](#)
- [Software Installation and Upgrade](#)
- [Subscriber Access Management](#)
- [User Interface and Configuration](#)
- [VPNs](#)

Authentication and Access Control

- The syslog message "UI_OPEN_TIMEOUT: Timeout connecting to peer" might appear if "show version detail" command is executed. This log is a cosmetic log and could be ignored. [PR895320](#)

Class of Service (CoS)

- When FPC/PIC restarts or Routing Engine reboots, the physical interfaces are created and the Class of Service daemon (cosd) sends chassis scheduler ADD for all interfaces. If a group of physical interfaces share the same Packet Forwarding Engine stream (such as oversubscribed PIC PD-5-10XGE-SFPP) and user configured chassis-scheduler is applied on some (NOT all) of the interfaces, the user configured chassis-scheduler can get over-written by default scheduler when chassis scheduler ADD comes for other non configured interface in same stream group with default scheduler-map. The issue can happen on any queuing PIC where multiple physical interfaces on PIC share same Packet Forwarding Engine/Chassis stream on FPC, in this bug, the fix is ONLY for PD-5-10XGE-SFPP. [PR809528](#)
- COSD errors are seen while Routing Engine switchover without GRES enabled. [PR827534](#)
- With GRES enabled and LSQ interface configured, after Routing Engine switchover, the following error message might be seen: COSD_GENCFG_WRITE_FAILED: GENCFG write failed (op, minor_type) = (add, policy inline) for tbl 4 if 301 /0/0 Reason: File exists [PR827538](#)
- In PPPoE/DHCP subscriber management environment, with "burst-size \$junos-cos-shaping-rate-burst" configured in subscriber dynamic-profiles, while logging in/out subscribers, the class-of-service daemon (cosd) memory leaks due to cosd process doesn't free up memory used for parsing burst attributes of a traffic-control-profiles (tcp) guaranteed rate. The memory usage of cosd process can be monitored by the following CLI command: user@router> show system processes extensive | match "PID | cosd" (Note: The "RES" field means "Current amount of

resident memory, in kilobytes") PID USERNAME THR PRI NICE SIZE RES STATE TIME
WCPU COMMAND 1326 root 1 96 0 14732K 4764K select 0:01 0.00% cosd. [PR846615](#)

- Commit throws an error "Invalid rewrite rule rule-name for ifl <ifl name>. lfd <lfd name> is not capable to rewrite inner vlan tag 802.1p bits" even though there is no rewrite configuration related to inner-vlan tag. [PR849710](#)
- The output of the "show subscribers extensive" command displays the Effective shaping-rate field only if you have enabled the effective shaping rate at the [edit chassis] hierarchy level. [PR936253](#)
- CoS relevant misconfiguration (for example, configure classifier exp for LT interfaces implicitly using "interface all" way) might cause cosd to crash. If cosd experiences multiple crashes within a short time, it might not be able to restart. [PR969900](#)
- This issue is specific to rate-limit on trunk port in DPC due to a software issue that installing rate-limit variables to egress Packet Forwarding Engine does not work normally. [PR1022966](#)

Forwarding and Sampling

- With more than four archive-sites configured under [system archival configuration archive-sites] hierarchy, after committing the configuration changes, pfd process crashes and generates core file due to memory corruption or double free. The core files could be seen by executing CLI command "show system core-dumps". [PR849465](#)
- This is a cosmetic issue. Deleting interface policer with "commit synchronize" later will see the following message on backup. This behavior cannot be seen without "commit synchronize". < Conditions > 1. Use 64bit Junos. 2. Configure "graceful-switchover" and "policer". 3. Delete interface policer configuration and then enter "commit synchronize". < backup RE messages > dfw_update_local_shared_policer: new filter program should be NULL for op 3 If using fixed version, "system syslog" can be reconfigured. [PR873084](#)
- When we configure unsupported firewall filter on channelized interfaces, commit error message show without this fix was misleading. With this fix, commit error will have the following message: mgd: error: layer2-policer is not supported for interface so-3/2/0 [PR897975](#)
- On MX Series-based platform, when deleting firewall filter and the routing instance it is attached to, in some race conditions, the filter might not be deleted and remains in resolved state indefinitely. [PR937258](#)

General Routing

- The configuration statement route-memory-enhanced(hierarchy: set chassis) is hidden in platforms M320 and MX series. There is no functionality break but this configuration statement should not be hidden. [PR690100](#)
- For an IPv4 pool, only the all-0 host and the all-1 host addresses are precluded from allocation, both for gateway-assigned and external address assignment. [PR729144](#)
- The next-hop-group configuration statement is not supported under the routing-instance hierarchy, but this configuration statement is present under this

hierarchy. This PR is opened to remove next-hop-group configuration statement from routing-instance hierarchy. [PR731264](#)

- ICMP redirects are not disabled even after configuring no-redirects on irb interface. [PR819722](#)
- When we execute the CLI command "show app-engine virtual-machine instance detail", if the virtual-machine (VM) is not ACTIVE, there should be a message displayed if it is waiting for secondary disk space to be available or for a particular interface to come up. In the fix we add the message. [PR824665](#)
- Changing static route with qualified-next-hop and order option to next-hop option results in static route missing from the routing table. Need restart routing to see the routes again. [PR830634](#)
- The issue will happen under following conditions: * The "forwarding-options sampling input maximum-packet-length" configuration statement is configured to a non-zero value; * Packets are sent to be sampled from a Type 5 FPC to an ES-Type FPC housing Multiservices PIC. Trigger: * Receive packets that should be sampled and the IP payload size is 53 bytes or above. Then the respective packets will be dropped and some SONN error logs: SRCHIP(0): 1 Bad packets on p1 SRCHIP(0): 1 SONN errors on p1. [PR839696](#)
- Restarting the FPC can terminate the DFWD process and create a core file. This will require a restart of the OpenFlow daemon for the OpenFlow functionality to work properly again. [PR842923](#)
- When a prefix next-hop address resolution requires a recursive lookup, the next-hop might not be updated correctly after an egress interface is disabled. [PR862989](#)
- In subscriber management environment, with scaling subscribers login (110K DHCP and 20K PPPoE), after restarting one of the line cards which has subscribers, autoconf process might crash and dump core due to memory corruption or memory double free. Only 11.4X27.45 is affected by this issue. [PR870661](#)
- The lldpd process might crash if there are multiple unknown type, length, and value (TLV) elements included in received LLDP PDUs. [PR882778](#)
- There are two issues about "flow term-order" Issue 1: Config command "set routing-options flow term-order standard" has no effect on the router without "restart routing". Issue 2: When the flow term-order is set to "standard", if the command "delete routing-options flow" is executed, the flow term-order should revert to the default setting of "legacy", but in this case, this does not happen. [PR885091](#)
- The rpd might crash when deactivate rib-groups (inet and inet6) under protocols ISIS, also these rib-groups applied under interface-routes. The core files could be seen by executing CLI command "show system core-dumps". [PR885679](#)
- When the NSR switchover happened immediately after a lot of vrf routing-instances were being deleted, garbage lsi interfaces will remain in kernel, while they are removed from RPD. Those garbage interfaces will result in KRT queue stack issue upon later lsi re-configuration. [PR912861](#)
- Periodic "show subscribers" CLI requests during the GRES recovery (on a scaled system) might lead to spawning of too many subinfo processes. As a side effect, CoA requests might not be serviced while system is kept busy by subinfo processes as authd might

take long time to be recovered (it was observed that authd is not recovered after 1+ hours). [PR915677](#)

- PIC removal without PIC offline can cause FPC crash in case of PD-5-10XGE-SFPP PIC. [PR922655](#)
- Traceroute or SSD crash seen when as-number-lookup option is used when executing traceroute. [PR928769](#)
- The message [RPD_KRT_KERNEL_BAD_ROUTE] may start to display on backup Routing Engine logs. This message does not point to a real network issue. This message is displayed because all nexthops are read in the backup RPD including the L2 nexthops, which are not created by RPD process. Therefore, RPD reads these L2 nexthops and discards them, then it displays this message. Hence these messages are harmless. To address this, and to stop swamping the syslog with these messages, a fix aims to print these messages only for L3 nexthops i.e. rpd-created nexthops has been implemented. [PR931667](#)
- When P2MP LSP is protected by link protection, it could have active and multiple standby next-hop. If one of the next-hop, regardless whether it is active or standby one, is removed due to FPC power-off or failure, multicast diagnostics daemon (mcdiagd) falls into infinite loop during collecting next-hop information. [PR931380](#)
- While a duplicate interface address (IFA) is configured for an interface, software will accept that with a warning message. But the kernel side cannot accept duplicate IFA, and needs to delete the next-hop created for this operation. However, the cleanup might not remove the duplicated IFA under heavy kernel workload. As a result, the kernel might crash while trying to update this duplicated IFA to Packet Forwarding Engine side. [PR936807](#)
- FPC would crash when access corrupted unicast next hop data structure that chain composite nexthop's depends on. [PR946276](#)
- When a router is booted with AE having per-unit-scheduler configuration and hosted on an EQ DPC, AE as well as its children get default traffic control profile on its control logical interface. However, if a non-AE GE interface is created on the DPC with per-unit-scheduler configuration, it will get default scheduler map on its control logical interface. [PR946927](#)
- On systems running Junos OS 13.3R1 and Nonstop routing (NSR) is enabled, when "switchover-on-routing-crash" under [set system] hierarchy is set, Routing Engine switchover should happen only when routing protocol daemon (rpd) crashes. But unexpected Routing Engine switchover can be seen when we perform CLI command "request system core-dump routing running" to manually generate a rpd live core. [PR954067](#)
- If an Aggregated Ethernet (AE) interface has the "scaled" member-link scheduling mode(which is the default mode) and multiple forwarding-classes map to a same queue, then the actual transmit-percent might unable to reach the configured in scheduler. [PR954789](#)
- "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](#)

- MPLS traceroute causes "rttable-mismatch" syslog messages. [PR960493](#)
- Default threshold for ES-FPC errors is 1 for major errors and 10 for minor errors. When the threshold is reached, some actions (for example, alarm|offline-pic|log|get-state|offline|reset) will be taken by FPC as configured. This feature is designed for permanent/real errors. The issue here is that even some transient errors (for example, link flaps) will also trigger the default action. In some cases, it might cause panic for the FPC. [PR961165](#)
- On MX Series DPC line cards with redundancy System Control Boards (SCBs), when active SCB goes down ungracefully by unexpected event (such as turn off Power Entry Modules (PEMs)), traffic loss is observed and cannot recovered on standby SCB as expected. [PR961241](#)
- On T Series or M320 routers with OSPF configuration statement. If have large-scale routes (for example, 180K Composite Nexthop), when do costing-out and costing-in operations along with changing gigether-options of core router facing interface multiple times continuously, the Flexible PIC Concentrator (FPC) CPU utilization might increase to 100 percent, and then FPC might crash. [PR961473](#)
- For MXVC platform, the Packet Forwarding Engine freconnect timer extends from the default 15 seconds to 60 seconds temporarily. This will be reversed once Packet Forwarding Engine connection issues resolved. [PR963576](#)
- Display issue only. "show route cumulative vpn-family" command is using "inet.6" for vpnv6 routes instead of inet6.0 [PR966828](#)
- On T1600/T640/T320 with FPC installed or M320 with FPC type 1/2/3(non E3 FPC). In processing for fpc-resync and fab-liveness packets (a kind of periodic probe packet) if error occurs while sending packet, the buffer of packet does not be freed. This causes packets buffers to leak and eventually the packet heap runs out of memory. [PR973892](#)
- In scenario of next-generation-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, and besides, rpd process might crash while trying to generate a live core dump. [PR982959](#)
- In point-to-point (P2P) SONET/SDH interface environment, there is a destination route with this interface as next-hop. When this interface is disabled, the destination route is still kept in the forwarding table and might cause ping fails with "Can't assign requested address" error. [PR984623](#)
- In the VPLS environment with control-word configuration, when the "control-word" is changed to "no-control-word", there is a 5 minute service outage. [PR987216](#)
- If a user configures an MX-VC member with member ID 2, the Virtual Chassis master Routing Engine may eventually experience a kernel panic. [PR989291](#)
- There was a bug in the code path for 'show route resolution' which was causing an extra decrement of the refcount in the show handling. This was causing an early free of some shared object and hence causing a crash. [PR995170](#)
- 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](#)

- If encapsulation type is "ppp" for the SONET interface on IQE PIC, sometimes the MTU change might not work. [PR1001880](#)
- If with accounting/sampling enabled, an unnecessary update from the routing protocol process (rpd) to the route record database might be triggered by certain configuration change. This process causes jump in CPU utilization of all Packet Forwarding Engines. [PR1002107](#)
- When a static discard route is configured with no-install option but actual forwarding using different next hop, if egress sampling is enabled on the forwarding outgoing interface (OIF), traffic leaving that interface would have incorrect OIF on the flow records, resulting in unreliable flow records and incorrect billing. There is no traffic impact though. [PR1002287](#)
- A raw IP packet with invalid Memory Buffer(mbuf) length may trigger a kernel crash. The invalid mbuf length might be set by other daemons incorrectly. [PR1006320](#)
- With NSR enabled, when activating a BGP session in a routing instance, and the interface route is imported into the main routing instance, the TCP receive window might decrement until it hits 0 after receiving incoming BGP traffic arrives from the main routing instance. [PR1003576](#)
- This PR is implementing traceoptions debug enhancements to detect route-record corruption events. The route-record traceoptions debug will be enabled as follows:
----- user@router> edit Entering configuration mode [edit]
user@router# set routing-options traceoptions flag route-record [edit] user@router# commit ----- [PR1015820](#)
- When destinations are pointing to protocol next-hops as unilist type or IP forwarding next-hops as unilist, which in scenarios like using Loop-Free Alternate Routes for OSPF (LFA-OSPF) with link protection, link-protection for P2MP LSP, or MPLS FRR is enabled. If flapping the active interface very fast, especially an interface comes back up before Kernel gets a chance to delete all the unilist next-hops, those unilist next-hops which have not been deleted yet would be re-used. As a result, the corresponding destinations are pointing to discard next-hop(s) or replaced next-hop(s) in Packet Forwarding Engine Jtree. The "discard" next-hop(s) causes traffic blackhole while the "replaced" next-hop(s) diverts traffic to other active next-hop(s) in the unilist. Those unilist next-hops which have been already deleted are safe and get updated accordingly. This is a day one timing issue. [PR1016649](#)
- On M Series, MX Series, T Series platforms with DHCP relay configured, the router might keep filling a specific partition "/var/mfs/sdb" with files named log.XXXX and this would eventually cause DHCP relay fail. [PR1017642](#)
- In a rare case, rdd core is reported under /usr/sbin/rdd as soon as applying the group and commit is performed. [PR1029810](#)
- This issue only affects OC-48 MICs. If an SFP is inserted into an OC-48 MIC port that has been disabled the SFP will not show up in a >show chassis hardware command. The issue is fixed with a patch. [PR1031851](#)
- With VPLS BGP control word configured, intermittent packet loss might be seen in one direction on VPLS circuit due to the control-word not being programmed at Packet Forwarding Engine after member DPC reboot. The problem can happen on below

conditions: 1. LSI interface exists across two or more physical interfaces. 2. Those physical interfaces located in different FPCs. 3. Those physical interfaces consist of equal-cost paths. So, LSI will not be flapped with one member FPC down. 4. Flap the member DPC where one of physical interfaces situated. [PR1031863](#)

- On MX Series with MPCs and MICs, in Connection-less Network Service (CLNS) L3VPN over Ethernet scenario, because the IS-IS route may be incorrectly treated as ARP-type route (32 bits routes with the MAC and outgoing interface specified in the nexthop), the routing protocol process (rpd) of Provider Edge (PE) device may reject to add the route. [PR1041251](#)
- Incorrect SNMP interface index (index 0) is given to VCP interface. Due to wrong index value, VCP interface description does not appear in SNMP interface table. [PR1044331](#)
- On a fresh reboot, SCG is able to establish subscribers with 10 rules at a rate of 2k calls per second (cps). However if we have more than 10 rules per subscriber then we see some calls getting rejected due to call-admission-control. Also these errors are seen only if the rate is consistently maintained at 2k cps continuously. However in the subsequent iterations (that is after deleting the sessions created in previous iteration) we do not see any failures even if we have more than 10 rules for subscriber creation of 2k cps. [PR1047686](#)
- MPC with Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC (MIC-3D-4COC3-1COC12-CE) might crash. This problem is very difficult to replicate and a preventive fix is planned to be implemented to avoid the crash. [PR1050007](#)
- As a precautionary measure, a periodic sanity check is added to FPC situated on M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX Series with DPC. It checks FPC error conditions and performs the appropriate actions in case of an error. [PR1056161](#)
- In LDP tunneling over single hop RSVP based LSP environment, after enabling "chained-composite-next-hop", the router may fail to create the chained composite next hops if the label value of VPN is equal with the label value of LDP. [PR1058146](#)
- In subscriber management environment, the Change of Authorization (CoA) request is not responded when a failure (like service session creation failure, Structure of Management Information posting failure, etc.) occurs. The authd process would crash when client profile update is sent in COA request. [PR1062442](#)
- For MX Series Virtual Chassis (MX-VC) with scaled subscribers, for example, 100K DHCP/20K PPPoE subscribers. If the Virtual Chassis port (VCP) FPCs also house the uplink ports and the "indirect-next-hop-change-acknowledgements" and "krt-nexthop-ack-timeout" configuration statements are configured along with the protection mechanism, after the master Routing Engine in the Virtual Chassis master router (VC-Mm) is powered down, the traffic loss and subscriber loss might be observed due to the indirect next-hop change acknowledgement timeout. With this fix, the upper limit for "krt-nexthop-ack-timeout" is changed from 100 seconds to 250 seconds. [PR1062662](#)
- When "satop-options" is configured on an E1 with Structure-Agnostic TDM over Packet (SAToP) encapsulation, after Automatic Protection Switching (APS) switchover, some

SAToP EIs on the previously protect interface (now working) start showing drops.

[PR1066100](#)

- On MX Series routers with MPC based line cards in a setup involving Packet Forwarding Engine fast reroute (FRR) applications, when BFD session flaps the next-hop program in the Packet Forwarding Engine may get corrupted. It may lead to incorrect selection of next-hop or traffic blackhole. [PR1071028](#)
- After rebooting the BNG with scaled subscribers, a dynamic-profile add request might fail, causing bbe-smgd (subscriber management daemon) to crash, then some subscribers might fail to login. [PR1071850](#)
- If the hidden configuration statement "layer-4 validity-check" is configured, the Layer4 hashing will be disabled for fragmented IP traffic. Due to a defect, the Multicast MAC rewrite is skipped in this case, the fragmented multicast packets will be sent with incorrect destination MAC. [PR1079219](#)
- On MX Series platform with MS-MPC/MS-MIC, when Network Address Translation (NAT), Stateful Firewall (SFW), Traffic Detection Function (TDF), or IPsec service is configured and traffic flows, an ordered packet might miss the descriptor due to the software defect. It results in prolonged flow-control, all data and control path are blocked, the service PIC goes down and not come up. [PR1079745](#)
- In some rare conditions, depending on the order in which configuration steps were performed or the order in which hardware modules were inserted or activated, if PTP master and PTP slave are configured on different MPCs on MX Series router acting as BC, it might happen that clock is not properly propagated between MPCs. This PR fixes this issue. [PR1085994](#)
- Memory leak likely given VC-Bm Routing Engine memory utilization sustained at 97% usage with high active and inactive memory utilization compared with VC-Mm Routing Engine. [PR1088132](#)
- For Broadband Subscriber Management L2 wholesale solution to deliver large scale services, in this case 8K sessions, not all dynamic L2-wholesale sessions are established when ANCP Port Up messages are received for each access line at first attempt. [PR1088965](#)
- In a fib-localization scenario, IPv4 addresses configured on service PICs (SP) will not appear on FIB-remote FPCs although all local (/32) addresses should, regardless of FIB localization role, install on all Packet Forwarding Engines. There is no workaround for this and it implies that traffic destined to this address will need to transit through FIB-local FPC. [PR1092627](#)
- The mspmand process might crash due to prolonged flow-control with TCP ALGs under the following possible scenario, mostly when the following conditions happen together: 1. When the system is overloaded with TCP ALG Traffic 2. There are lots of retransmissions and reordered packets [PR1092655](#)
- If a service-PIC is configured to simultaneously function as both an MS interface and as a member of an AMS interface, then some settings under services-options may not apply correctly. These settings are A) syslog_rate_limit, B) fragment-limit, C) reassembly-timeout and D) jflow_log_rate_limit. [PR1096368](#)

- When the control path is busy/stuck for service PIC, the AMS member interface hoisted by it might be down, but when the busy/stuck condition is cleared, the member interface might not recover, and AMS bundle still shows the PIC as inactive. [PR1093460](#)
- On TCP ALG, if there are a lot of retransmissions and reordered TCP packets, and the system is overloaded due to the TCP traffic, the mspmand (which manages the service PIC) process might crash. [PR1093788](#)
- Extensive Header integrity checks will be done for packets which match a service set which has NAT/SFW configured. 1. Enable Header integrity checks by default when SFW or NAT is configured in same service set. This is inline with ukernel behavior 2. Retain the configuration statement for use by other plugins such as IPsec which may want to enforce header integrity if needed 3. Ensure that the cmd "show services service-sets statistics integrity-drops" works if sfw/nat is configured. [PR1095290](#)
- If fpc offline configuration statement is configured after the presence of Non-recoverable faults, then offline action will not be performed. [PR1103185](#)
- On MX Series platform, the output of CLI command "show system subscriber-management route" may be shown as empty. [PR1104808](#)
- An IPv4 filter configured to use the filter block with term that has both "from precedence" and another non 5-tuple (i.e. not port, protocol, address) will cause an XL/EA based board to reboot. Example: set firewall family inet filter FILTER fast-filter-lookup set firewall family inet filter FILTER term TERM from precedence PRECEDENCE set firewall family inet filter FILTER term TERM from tcp-established [PR1112047](#)

High Availability (HA) and Resiliency

- Configuring the maximum segment size (MSS) for the TCP connection for BGP neighbors, if "mtu-discovery" and "path-mtu-discovery" configuration statements are removed, the default MSS value of 512 will be used instead, this is not an expected behavior. [PR835220](#)
- PR 855661 will affect IQ2 PICs during unified ISSU on TX platform. When upgrading to Release 12.3R2 from releases prior to 12.3R2 through unified ISSU, IQ2 PICs will report an error. This error is because IQ2 PICs are not able to download the image during unified ISSU. [PR855661](#)
- During a router hardware upgrade procedure, in a dual Routing Engines system, the newly installed Routing Engine may overwrite the other Routing Engine configuration with the factory default configuration. As a result, both Routing Engines may bootup in "Amnesiac" mode. This situation can occur under following conditions: - RE0 has default factory configuration and, - RE1 has "commit synchronize" enabled - Both RE0 and RE1 boot-up simultaneously, or - RE0 is UP and running and RE1 is restarted. [PR909692](#)
- If NSR Routing Engine switchover has done right after committing the configuration change which deletes routing-instance(s), some of those instances might not be deleted from forwarding table. [PR914878](#)

Infrastructure

- Executing the 'show multicast route' CLI command with a crafted argument may cause the CLI process to crash, creating a cli.core.O.gz. The impact of the crash is limited to the CLI process executing the command, so malicious exposure is minimal. [PR939262](#)
- SSD failure does not trigger Routing Engine switchover. [PR1102978](#)

Interfaces and Chassis

- On logical tunnel (lt) interfaces, you might not be able to use the 'family vpls' option at the [edit interfaces lt-fpc/pic/port unit logical-unit-number] hierarchy level. [PR44358](#)
- When you have the following configuration on a logical interface, unit 2000 { encapsulation vlan-bridge; vlan-tags outer 40 inner-list [20 3000]; family bridge; } And you execute "show interface intf-name extensive" you will see the below: Under " Flags: SNMP-Traps Redundancy-Device 0x20004000 VLAN-Tag [0x8100.40 0x8100.2000 20,3000] ", you will see the unit number 2000 between outer and inner tags configured. This is just a display issue and no functionality is affected. [PR723188](#)
- To troubleshoot a particular subscriber, one can use 'monitor traffic interface <ifd> write-file xy.pcap'. Using this command on aggregated or demux interfaces can lead to corrupted ingress packets in the PCAP file. Customer traffic is not affected though. [PR771447](#)
- Collecting subscriber management control traffic via 'monitor traffic interface demux0 write-file xy.pcap', the logical unit number is incorrect when multiple demux IFL's are present. This problem is fixed and the correct interface logical unit number is reported in the juniper header of the captured PCAP file. [PR771453](#)
- On MX series, if a composite next hop entry which is deleted in Master NW-INE exists in Backup NW-INE, the kernel may crash with a vmcore file generated during NW-INE switchover. [PR793098](#)
- Issue is seen only when the following steps are followed: 1. Enable IRB MAC Sync feature. 2. Deactivate BD/MAC Sync/Service ID on the higher MAC node. 3. Activate virtual switch that configures MCAE under the virtual switch. The result: IRB MAC Sync happens even though the feature is not enabled. [PR793889](#)
- Master LED of craft interface keeps Green during Halt the system or Power off. [PR805213](#)
- Warning message added is syslog when external sync is not supported. [PR817049](#)
- Prior to this PR, the speed of a GE interface capable of working at FE speeds was set to 'auto' in the Packet Forwarding Engine level. This causes a problem when manually setting the speed on the Routing Engine. Now the behavior is to set the speed to '1 g' in the Packet Forwarding Engine. For automatic speed detection the interface should be set to 'speed auto' in the configuration. [PR821512](#)
- With Junos OS Release 11.4 or later and Enhanced SCB installed on a mix of MX Series routers and DPC cards, REG_ERR messages might be reported under certain traffic flow conditions from MX Series routers to the DPC card. On the receiving DPC card fabric cell received out of order will be re-ordered and merged to build the packet. If this out-of-order delivery is too high a reorder event will be triggered and all cells

belonging to the packets are dropped. The frequency is low rate. The following syslog entry will be reported Sep 29 20:43:10 node fpc8 ICHIP(3)_REG_ERR:first cell drops in ichip fi rord : 4122 Sep 29 20:43:10 node fpc8 ICHIP(3)_REG_ERR:Non first cell drops in ichip fi rord: 7910 [PR821742](#)

- In some circumstances, such as deleting a routing-instance containing active dynamic subscribers, a flag db_on_miss_queue is wrongly set for db_entry and leads the system trying to access the Null, hence causes the crashing of the DCD process. [PR826899](#)
- A request (like snmp query) for collecting input ipv6 stats of ae logical interface on abc chipset is not working properly. [PR831811](#)
- In PPPoE subscriber management environment, while subscribers login/logout, each subscriber will use an Event Rate Analyzer (ERA) until the outcome of the subscriber connection (whether it succeeds or fails). During a logout of a high number of subscribers (e.g. 16k), all the ERA events are quickly exhausted (there are 1250 in total), so that new logins are blocked until ERA events start to be freed. [PR842935](#)
- Whenever tunnel interface -pe/-pd gets created using the MS-DPC instead of the MPC, it will not be able to process register messages. Because of MPC and MS-DPC have different multicast architectures and they are incompatible, if chassis is configured in "enhanced-ip" mode this issue will be seen. Necessary changes have been made to code so that these interfaces will not be created on MS-DPC. [PR853995](#)
- On M Series, MX Series, and T Series platforms with Services PIC and dual Routing Engines, configure MPLS and set Services PIC in layer 2 mode. Apply class-of-service (CoS) configuration to sp-x/y/z control interface. After that perform graceful Routing Engine switchover (GRES), and the Services PIC might restart. [PR859036](#)
- In scaled MX-VC environment, AE interfaces may get removed from the Kernel after the GRES switchover. [PR860316](#)
- To configure FEC thresholds via CLI, use string format with mantissa and exponent:
Example: set interfaces et-1/0/0 otn-options signal-degrade
ber-threshold-signal-degrade 1.23E-4 set interfaces et-1/0/0 otn-options
signal-degrade ber-threshold-clear 2.34E-5 [PR886572](#)
- On MX Series router, the physical or logical interfaces (ifd/ift) might be created and marked UP before a resetting FPCs' fabric planes are brought up and ready to forward traffic, as a result, traffic might be black-holed during the time window. This window of traffic black-hole is particular long if the chassis is heavily populated with line-cards, for example, the router has large scale of configuration (routes or subscribers), and coupled with a lot of FPC reset, such as upon a node power up/reset. [PR918324](#)
- Non-Existent leg in AE bundle prevents DHCP subscribers from coming up. [PR918745](#)
- During MX-VC bootups, when initial role is determined, chassisd restarts. Until the role changes, we have two LCC contexts and no SCC yet. Chassisd restart at bootup time may result in additional reboots of slower cards along with the churn in the system. The fix is to separate the SCC and LCC chassisd contexts in order to make the (local) Routing Engine mastership logic the same for both standalone MX chassis and MX-VC. The code change has also fixed the following issues: PR868418: MXVC-Commit causes "Loss of communication with Backup RE" minor alarm PR805436: backup Routing Engine (VC-Bm) going offline momentarily observed during configuration commit

PR969650: MX-VC, commit-sync fails on member0 local backup Routing Engine. The fix is to separate the SCC and LCC chassisd contexts in order to make the (local) Routing Engine Mastership logic the same for both stand-alone MX chassis and MX-VC. The code change has also fixed the following issues: PR868418: MXVC-Commit causes "Loss of communication with Backup RE" minor alarm PR805436: backup Routing Engine (VC-Bm) going offline momentarily observed during configuration commit

PR969650: MX-VC, commit-sync fails on member0 local backup Routing Engine

[PR919894](#)

- PPPoA session would not come up on removal/addition of cable to the tester port. [PR939404](#)
- Strange FRU Insertion trap [RE PCMCIA card 0] is generated when Routing Engine master-switching is done on box with RE-1800. [PR943767](#)
- Following logs will be seen with certain config changes. Dec 24 06:22:16 dcd[4177]: ael: per-unit-scheduler is valid only on FR and VLAN encapsulation Dec 24 06:22:17 dcd[1660]: ael: per-unit-scheduler is valid only on FR and VLAN encapsulation Issue was, the check for allowing per-unit-scheduler on ae interface was done before encapsulation attribute of ae interface was read, and hence dcd was throwing log error message for any of the modification on ae interface. Issue is cosmetic in nature. [PR951434](#)
- When transit traffic of Ethernet frames of size less than 64 bytes are received by 1x 10GE(LAN/WAN) IQ2E PIC, the router forwards the frames instead of dropping them. [PR954996](#)
- When an ifl containing some VRRP group configuration is deleted, SNMP walk on VRRP MIB may loop continuously. [PR957975](#)
- In very uncommon situation, we will see LCCs chassisd state is inconsistent with SFC chassisd state, this is very misleading in troubleshooting stage. This PR fixed this issue. [PR963342](#)
- After changing the speed of fxp0 interface (the management Ethernet interface) to 1G (the maximum speed), the interface process (dcd) configures the interface but reads the speed even before the change takes effect. Although the hardware speed is updated to 1G, from dcd perspective, the speed is still not changed. Then if you change back to the original speed, the change is ignored by dcd. [PR976825](#)
- On MX Series platform, when an aggregated Ethernet bundle participating as L2 interface within bridge-domain goes down, the following syslog messages could be observed. The messages would be associated with FPC0 even if there are no link(s) from this FPC0 participating in the affected aggregate-ethernet bundle. mib2d[2782]: SNMP_TRAP_LINK_DOWN: ifIndex 636, ifAdminStatus up(1), ifOperStatus down(2), ifName xe-3/3/2 mib2d[2782]: SNMP_TRAP_LINK_DOWN: ifIndex 637, ifAdminStatus up(1), ifOperStatus down(2), ifName xe-3/3/3 mib2d[2782]: SNMP_TRAP_LINK_DOWN: ifIndex 740, ifAdminStatus up(1), ifOperStatus down(2), ifName ae102 fpc0 LUCHIP(0) Congestion Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f:f:f:f:f fpc0 LUCHIP(0) Congestion Detected, Active Zones 2:0:0:0:0:8:a:0:0:0:0:0:8:4:0:a alarmd[1600]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Major Errors craftd[1601]: Major alarm set, FPC 0 Major Errors fpc0 LUCHIP(0) Congestion Detected, Active Zones 2:0:0:0:0:8:a:0:0:0:0:0:8:4:0:a alarmd[1600]: Alarm cleared: FPC color=RED,

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

- In the PPPoE environment, when the subscriber login successfully but profile activate fails, due to code processing error, the address entry is not deleted in the authd's DAP pool. So when the subscriber tries to login again, it connects fails. [PR995543](#)
- In Ethernet OAM connectivity-fault-management, Junos OS default encodes MAID(MD name and MA name) in character format. Currently only 43 octets are supported in Junos OS for the MD + MA name. Junos OS needs to support maximum length of 44 octets for MAID per the standards. [PR997834](#)
- On MX Series-based line cards, in virtual private LAN service (VPLS) environment, the next hop in the kernel allocated by connectivity-fault management process (cfmd) may not be freed even after the CFM session has been removed (for example, deactivating the routing-instance). In this situation, after re-activating the routing-instance, the interface within the routing instance would fail to come up because the nexthop is not freed by the cfmd application and hence the VPLS connection is down. [PR1000060](#)
- When IEEE 802.3ah OAM link-fault management action profile is configured to define event and the resulting action, the link might flap after it is brought down by an event but brought up by other events erroneously. [PR1000607](#)
- With vrf-table-label configured on the routing-instances, when an FPC with Enhanced IQ (IQE) PIC is sharing the same Forwarding Engine Board (FEB) with another FPC, and the FEB has two core-facing interfaces configured with the family mpls on aforementioned FPCs separately, the label-switched interface (LSI) might be removed incorrectly on the working FPC when the other FPC with IQE PIC is set to offline. [PR1027034](#)
- During the test of login/logout of couple of thousands of PPPoE subscribers, pppoe may crash by generating a core dump file. [PR1041367](#)
- On MX Series platform, when ACI VLAN interface sets are configured for PPPoE subscribers, PPPoE process (pppoed) crash may occur during PPPoE control packet processing when the ACI VLAN interface set needs to be created. If this pppoe crash

is seen, then the ACI VLAN interface set will not be created and PPPoE subscriber login will not make progress. There is no workaround for this issue and an upgrade to a release that includes the fix for this PR is recommended. [PR1057343](#)

- During subscriber login/logout the below error log might occur on the device configured with GRES/NSR. /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV (pp0.1073751222) [PR1058958](#)
- For multichassis link aggregation groups (MC-LAGs) running in active-active mode with back-to-back square topology, when the Inter-chassis Control Protocol (ICCP) is broken between any MC-LAG devices, the non preferred device reverts to its own local system ID. But its Link Aggregation Control Protocol (LACP) partner on the remote side does not remove the flap link from AE bundle and it remains UP. This might cause a network wide loop resulting in traffic outage until manual intervention. [PR1061460](#)
- After chassis restart or non graceful Routing Engine switchover, few vrrp sessions may start sending vrrp packets while still in vrrp backup state. [PR1062929](#)
- Deactivating/activating logical interfaces may cause BGP session flapping when BGP is using VRRP VIP as source address. This is caused by a timing issue between dcd and VRRP overlay file. When dcd reads the overlay file, it is not the updated one or yet to be updated. This results in error and dcd stops parsing VRRP overlay file. [PR1089576](#)
- During scaling login/logout different types of subscribers (e.g. 17K) on LAC router, there might be some L2TP LAC subscribers stuck in terminating state and never get cleared, blocking new sessions from establishing on the same interface. [PR1094470](#)
- On PB-2OC12-ATM2-SMIR PIC, port 0 and port 1 are configured with clock source as external. If Loss of signal (LOS) is inserted on port 0, the port 0 will down. The expected behavior is clock being used from port 1. But in this case, port 0 down will result in port 1 flapping and reporting SONET phase lock loop (PLL) errors. [PR1098540](#)
- Due to the fact that the error injection rate configured by user on Routing Engine via CLI command "bert-error-rate" may not be programmed in the hardware register, the PE-4CHOC3-CE-SFP, PB-4CHOC3-CE-SFP, MIC-3D-4COC3-1COC12-CE, and MIC-4COC3-1COC12-CE-H may fail to inject bit errors during a Bit Error Ratio Test (BERT). [PR1102630](#)
- The jpppd process (which is used to authenticate subscribers) might crash after restarting MPC in live network, and then some subscribers might be found stuck in INIT state. [PR1114851](#)

Layer 2 Features

- When directly apply sampling on VPLS interface (i.e interface ge-4/0/1 unit 0 family vpls sampling input), if customer configures logical interface and sampling input/output together first time, then deactivating sampling input/output through CLI, kernel will then not disable the sampling. Also note that, the action of sampling is a hidden command for VPLS interfaces and would not be listed in "possible completion" list when combined with "?". [PR772270](#)
- It can happen that when changing an interface framing from lan-phy (default) to wan-phy and back a few times, the interface doesn't show up any more in "show interfaces terse". [PR836382](#)

- DHCPv6 fails for clients using DUID type 2 (Vendor-assigned unique ID), the software was using the DUID to extract MAC address information. This behavior is fixed and tested. [PR838404](#)
- "show bridge mac-table interface X vlan-id Y" is empty on trunk port. This is just a display issue. This MAC is present on the forwarding table that can be confirmed using command "show route forwarding-table family bridge". [PR873053](#)
- In DHCP relay scenario, some DHCP relay bindings might get stuck in "RELEASE(RELAY_STATE_WAIT_AUTH_REQ_RELEASE" state due to the LOGOUT Request is not processed correctly by authentication manager process (authd) and this is causing clients not to be able to get a lease. [PR850187](#)
- 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](#)
- After change the way of getting site ID of VPLS from fixed site-id to automatic-site-id on one site while other sites are still using the fixed site-id in the network, the rpd process might crash due to the site ID get by "automatic-site-id" may conflict to site ID which was configured as fixed site ID on other sites. [PR1054985](#)
- With Dynamic Host Configuration Protocol (DHCP) maintain subscriber feature enabled, when the subscriber's incoming interface index is changed, for example, the interfaces go away and come back after changing the MTU configuration of interface, the existing subscribers may get dropped and new subscribers fail in connection. [PR1059999](#)
- During interface flaps a high amount of TCN (Topology Change Notification) might get propagated causing other switches to get behind due to high amount of TCN flooding. This problem is visible after the changed done from 11.4R8 onwards which propagates TCN BPDU immediate and not in the pace of the 2 second BPDU Hello interval to speed up topology change propagation. The root cause is the TCNWHILE timer of 4 seconds is always reset upon receiving TCN notifications causing the high churn TCN propagation. [PR1089580](#)

MPLS

- RSVP graceful restart does not function for LSPs that have a forwarding adjacency (FA) label-switched path (LSP) as a next hop. [PR60256](#)
- Statistics for a P2MP LSP used for CCC connection will not be displayed on a Ingress PE. [PR444336](#)
- Customer upgrading network using features involving Non-Penultimate Hop Popping Behavior and Out-of-Band Mapping should upgrade routers involved together to Release 13.1 or later. [PR852808](#)
- In current Junos OS, lsping/lsptrace utilities have compatibility issue with other vendor routers. Millisecond field will show huge value which results in incorrect RTD calculated.

```
user@router> ping mpls ldp 192.168.228.7/32 source 192.168.199.193/32 exp 5 count 5 size 100 detail Request for seq 1, to interface 510, label 1102, packet size 100 Reply for seq 1, return code: Egress-ok, time: 3993729.963 ms <---- Local transmit time: 2013-04-29 12:05:06 IST 873.491 ms Remote receive time: 2013-04-29 12:05:06 IST 3994603.454 <---- This is cosmetic issue and current software limitation. PR891734
```


- The RSVP bandwidth of the aggregate ethernet (AE) bundle does not adjust properly when a member link is added to AE interface, and at the same time an IP address is removed from this AE bundle. [PR948690](#)
- During SNMP walk on table MPLS cross-connect table (mplsXCTable) in case of flood nexthop, the rpd might crash. [PR964600](#)
- If RSVP preemption is set to be aggressive mode, when MPLS LSP Make-Before-Break (MBB) event happens due to auto-bandwidth or manual optimization, traffic loss might be seen. [PR984741](#)
- LSP metric modification leads to Constrained Shortest Path First (CSPF) computation and resignaling. It should update RSVP routes directly. [PR985099](#)
- In next-generation MVPN extranet scenario, if there is a mix of VT interface and LSI (vrf-table-label is used) interface on next-generation MVPN egress node, after changing some vrf policies, the routing protocol process (rpd) might crash and reset. [PR1045523](#)
- In Resource Reservation Protocol (RSVP) environment, if CoS-Based Forwarding (CBF) for per LSP (that filter out traffic not related to that LSP) is configured, and either the feature fast-reroute or link-protection is used on the device, when the primary link is down (for example, turning off the laser of the link), due to some next hops of the traffic may be deleted or reassigned to different class of traffic, and the RSVP local repair may fail to process more than 200 LSPs at one time, the traffic may get dropped by the filter on the device before the new next hop is installed. In this situation, the feature (fast reroute or link protection) may take longer time (for example, 1.5 seconds) to function and the traffic loss might be seen at the meantime. In addition, the issue may not be seen if the CBF for per LSP is not configured on the device. [PR1048109](#)
- With BGP prefix-independent convergence (PIC) edge feature enabled, more than one BGP next-hop association will be installed in the Packet Forwarding Engine for MPLS VPN and Internet transit traffic. Deactivating/activating the IGP protocol (IS-IS or OSPF) might cause the backup session to stay down on Packet Forwarding Engine. [PR1058190](#)
- When fast-reroute, node-link-protection or link-protection is configured, if a Shared Risk Link Group (SRLG) is associated with a link used by a LSP ingressing at a router, then on deleting the SRLG configuration from the router, the SRLG entry still stays in the SRLG table even after the re-optimization of this LSP. [PR1061988](#)
- The point-to-multipoint (P2MP) label-switched path (LSP) is unable to re-establish after certain links are down. This issue might be encountered when the links are those that contain the primary and backup LSPs for the P2MP LSP. The P2MP LSP can be restored after the links are up. [PR1064710](#)
- When a primary LSP gets re-routed due to better metric, Link/Node protection for this LSP is expected to come up within 7 seconds provided the bypass-lsp protecting the next-hop link/node is already available. However in some corner cases, the Link/Node protection for re-routed primary LSP will not come up within 7 seconds even with bypass-lsp availability. The PR fixes this issue and reduces the delay of associating bypass-lsp with primary-lsp from 7 seconds to 2 seconds. [PR1072781](#)

Multicast

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

Network Management and Monitoring

- When OID dot3adAggPortTable is polled on the router, snmpd tries to fetch interface/interface-unit data from kernel in ASYNC mode, it is possible that by the time kernel replies with stats the interface/interface-unit state is already changed or deleted. This problem is most commonly seen when kernel replies for interface/interface-unit are late, more than expected window. Accessing interface/interface-unit stats for which the state has already been changed can cause the mib2d core-dump. At such times mib2d must first check if the interface/interface-unit entries are still valid (i.e. not changed or deleted) before accessing the associated data-structure. A check has been added in Junos OS Release 13.2R2 and later via this PR which marks the interface/interface-unit as stale if state has been changed. This prevents the code from accessing any associated data-structure if the entry is marked as stale. [PR852282](#)
- Customer might see the following messages. mib2d[xxxx]:
MIB2D_RTSLIB_SEQ_MISMATCH: mdb_async_mismatch: sequence mismatch (15,14),
resyncing mib2d[xxxx]: SNMP_RTSLIB_FAILURE: mdb_conv_update_seq_err: mdb
conv: socket seq error [PR889652](#)
- While some set operation is in progress, there is a huge pile-up of pending requests in netsnmp_agent_queued_list Queue, which is running into several thousands of requests and causing the memory consumption increase in snmpd, and finally, running out of 256 MB of memory and crashing. [PR920471](#)
- 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](#)
- Mib2d cores while trying to re-add a lag child into the internal DB. Since the entry is already present in the internal DB. Before adding the child link mib2d does a lookup on the tree, to know if the entry is not already there. However, this lookup returns no results, since the child link is part of snmp filter-interface configuration. [PR1039508](#)
- In rare cases, when the mib2d process attempts connection with the snmpd process and there are pending requests waiting to be finished, the mib2d process might crash and the CPU utilization is high around the same time as the crash happens. [PR1076643](#)

- Due to a bug in jnxIfcInline mib, a high order interface churn such as the one done by submitter in this case, can lead to a mib2d core. The situation is recovered after the core and no other impact is seen. [PR1105438](#)
- While the router is rebooting and SNMP polling is not stopped, SNMP requests might land on mib2d process before Routing Engine protocol mastership is resolved, causing the mib2d process crash. [PR1114001](#)

Platform and Infrastructure

- Commit time warning is changed to trace message. [PR480082](#)
- On the process details page (Monitor > System View > Process Details) of the J-Web interface, there are multiple entries listed for a few processes that do not impact any functionality. [PR661704](#)
- On the JCS-1200 RE-JCS-1X2400-48G-S Routing Engine configuration of the MAC address on the external interfaces em0 and em1 is not allowed. You cannot configure the MAC address on fxp0 on the other routing engines supported on the JCS-1200 as well. Therefore, the Junos OS CLI to configure the MAC address on em0 and em1 interfaces has been disabled. [PR770899](#)
- XML tags for get-software-information output missing some elements of new Junos OS service release naming convention. [PR783653](#)
- There is a problem going from 12.2 to 12.3 using unified ISSU. The blobs being created in 12.2 are using the newer format which is not compatible with 12.3 code. [PR818947](#)
- Adaptive load-balance functionality is only supported for unicast traffic. If the aggregate bundle contains logical interfaces for bridge or vpls domains, flooded traffic might get dropped. [PR821237](#)
- CLI command 'show route forwarding-table' would only display <= 16 ecmp paths when CBF is used. [PR832999](#)
- Deny-commands is not working for "show route community-name" [PR836624](#)
- cscript core generated during pressure test of ServiceNow [PR843062](#)
- Distributed protocol adjacencies (LFM/BFD/etc) may experience a delay in keepalives transmission and/or processing due to a prolonged CPU usage on the FPC microkernel on T4000 Type 5-3D FPCs. The delay in keepalive transmission/processing may result in a mis-diagnosis of a link fault by the peer devices. The issue is seen several seconds after a Routing Engine mastership switch with NSR enabled and the fault condition will clear after a couple of minutes. [PR849148](#)
- With inline jflow enabled, if the low 12 bits of the packet counter are zero (0x000) while copying packets count from hash record into flow export packet, the packetDeltaCount counter might be incorrect in inline jflow records. There is no traffic impact but may impact billing. [PR886222](#)
- In DHCP relay agent scenario, DHCP offer message with option82 (relay-agent-option) is discarded by UDP Forwarding process (fud) after receiving the reply back from DHCP server. This issue happens when the length of interface name (including underlying

and parent interface) greater than 23. For example: `irb.1011/0/0.1011` - 22 characters works `irb.1011/0/0.10011` - 23 characters fails [PR886463](#)

- On all high-end MX Series devices, when a router is acting as an NTP broadcast server, broadcast addresses must be in the default routing instance. NTP messages are not broadcasted when the address is configured in a VPN virtual routing and forwarding (VRF) instance. [PR887646](#)
- The `jcs:dampen()` function will not perform correctly if the system clock is moved to an earlier time. [PR930482](#)
- With MX Series based line cards, change MTU on one interface might cause L2 traffic interruption on other interfaces in the same FPC. [PR935090](#)
- On a router which does a MPLS label POP operation (penultimate hop router for example) if the resulting packet (IPv4 or IPv6) is corrupted then it will be dropped. [PR943382](#)
- If an unnumbered interface is configured as the next-hop interface for a static route, when issue "`clear arp hostname <hostname>`" command twice, the ARP entry for static route would get deleted and could not get updated any more. The ARP entry will appear again only after deactivating and then activating static route.
- Backing up the configuration with transfer-on-commit does not work in a MX-VC environment. [PR947444](#)
- Bad udp checksum for incoming DHCPv6 packets as shown in monitor traffic interface output. The UDP packet processing is normal, this is a monitor traffic issue as system decodes checksum=0000. [PR948058](#)
- With FPC3-E3 type FPC, the internal pc- interface statistics on the IQ/IQ2 PIC will be the same as the ingress interface statistics of the physical interface if family mpls is configured. It is a cosmetic display issue. [PR953183](#)
- When both FPC CPU are already very busy and a very large firewall filters (thousands of terms long) are to be used, a port being used for port mirroring goes down due to an external factor (such as a fiber cut or the remote side rebooting), the FPC CPU may rise to 100% for 4 minutes and then followed by a FPC reboot with a reason of "pfeman watchdog expired". This issue will only be observed occasionally, and if any of these three factors is not present the issue will not occur. So, disabling the port being used for port mirroring prior to bringing down that link is sufficient to avoid this issue. [PR968393](#)
- 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](#)
- In the dual Routing Engines scenario with NSR configuration, the configuration statement "`groups re0 interfaces fxp0 unit 0`" is configured. If disable interface `fxp0`, backup Routing Engine is unable to proceed with commit processing due to SIGHUP not received, the `rpd` process on backup Routing Engine might crash. [PR974430](#)

- 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](#)
- no-propagate-ttl doesn't work for L3VPN when PE is configured with l3vpn-composite-nexthop and its core interfaces are hosted on MPC based FPC. [PR985688](#)
- XML traceroute does not display as-numbers. [PR988727](#)
- When receiving traffic coming on MPC and going out on DPC, the MAC entry on a Packet Forwarding Engine might not be up-to-date and the frames targeted to a known MAC address will be flooded across the bridge domain. [PR1003525](#)
- In PPPoE over ATM subscriber management environment with active subscribers present, when you issue the "show arp" command, an ARP core file is generated. [PR1006306](#)
- For MX Series platform with inline Network Address Translation (NAT) service, when using "source-prefix" or "destination-prefix" in a NAT translation rule, a pool is implicitly created, appending "_jinpool_" with the rule name and term name with a form : _jinpool_{rule_name}_{term_name}. The name might be cropped due to the maximum length limitation (64 characters). If that happens, both pools might get the same name and result in the indeterminate behavior (statistic issue, drop or incorrect translation). [PR1020033](#)
- LSI logical interface input packet and byte stats are also added to core logical interface stats, but when the LSI logical interface goes down and the core logical interface stats are polled, there is a dip in stats. The fix is to restore LSI logical interface stats to core logical interface before deleting the LSI logical interface. [PR1020175](#)
- CPQ RLDRAM ECC single and double bit errors will generate CM alarm. "show chassis alarms" command can be used to view CM alarm. Details ===== 1> CPQ RLDRAM ECC single bit error in last 10 secs will raise minor CM alarm. 2> No CPQ RLDRAM ECC single bit error in last 10 secs will clear minor CM alarm. 3> CPQ RLDRAM ECC double bit error will raise Major CM alarm (this alarm will not be cleared until the FPC is restarted) [PR1023146](#)
- On all high-end MX Series devices, the packets per second (pps) and bits per second (bps) counters are not reporting accurate values while checking the "monitor traffic interface interface-name" command or the "show interface interface-name extensive" command. [PR1033222](#)
- Recurring local memory (LMEM) data errors may cause lookup chip on MX Series based FPC wedge and eventually FPC crash. [PR1033660](#)
- The Priority code point (PCP) and Drop eligible indicator (DEI) bit in 802.1Q header are preserved while packet gets routed within the same Packet Forwarding Engine . The expected behavior is resetting the PCP and DEI bit when the packet is routed. [PR1036756](#)
- On MX Series routers with MPCs/MICs, when using inline Two-Way Active Measurement Protocol (TWAMP) server (the server address is the inline service interface address),

because the TWAMP server may incorrectly calculate the packet checksum, the packet may get dropped on the TWAMP client. [PR1042132](#)

- ISC BIND software included with Junos for MX Series devices is affected by CVE-2014-8500. This may allow a network based attacker to cause a denial of service condition on SRX devices. This issue only affects SRX devices where "set system services dns dns-proxy" has been configured. This is not enabled by default on SRX devices. This issue does not affect other Junos OS based devices as they do not have BIND DNS server feature. This issue has been assigned CVE-2014-8500. Please see <https://kb.juniper.net/JSA10676> [PR1048628](#)
- For a Routing Matrix, if different Routing Engine models are used on switch-card chassis (SCC)/switch-fabric chassis (SFC) and line-card chassis (LCC) (for example, RE-1600 on SCC/SFC and RE-DUO-C1800 on LCC), where the out-of-band (OoB) management interfaces are named differently (for example, fxp0 on SCC/SFC Routing Engine and em0 on LCC Routing Engine), then the OoB management interface configuration for LCC Routing Engine will not be propagated from SCC/SFC Routing Engine during commit. [PR1050743](#)
- Under very rare situations, Packet Forwarding Engines on the following linecards, as well as the compact MX80/40/10/5 series, may stop forwarding transit traffic: - 16x10GE MPC - MPC1, MPC2 This occurs due to a software defect that slowly leaks the resources necessary for packet forwarding. Interfaces handled by the Packet Forwarding Engine under duress may exhibit incrementing 'Resource errors' in consecutive output of 'show interfaces extensive' output. A Packet Forwarding Engine reboot via the associated linecard or chassis reload is required to correct the condition. [PR1058197](#)
- When MX Series platform acts as Virtual Extensible Local Area Network (VXLAN) gateway, if there are multiple Packet Forwarding Engines, VXLAN packets will be distributed to available Packet Forwarding Engines in the chassis to perform VXLAN encapsulation/decapsulation. This is not expected (Expect behavior: VXLAN packet processing will be done on the same Packet Forwarding Engine on which it is received). This might result in unexpected packet drop and also overlay ping/traceroute not working. [PR1063456](#)
- For inline-jflow service on MPC3 and MPC4 linecards of MX Series, the packets get sprayed across all four lookup units (LUs). Normally the records from different LUs should have different observation domain ID (obs-id), but in this case they all have the same obs-id, causing four different sets of sequence numbers for the flow records. [PR1066319](#)
- When under configure batch mode, the commit error is noticed when the system is missing the 'commit-queue' folder under /var/db for some unknown reason. [PR1069440](#)
- Aggregate Ethernet (AE) interfaces in combination with shared-bandwidth-policer might lead to Packet Forwarding Engine policer corruption if there are child member links configured on the same Packet Forwarding Engine and the AE interface is being reconfigured (add/delete of logical units). This corruption could alter the policer rate programmed in hardware and lead to unexpected policer behavior. A different trigger with physical interface flap illustrating the same symptoms are tracked in [PR1035845](#). [PR1084912](#)

- On MX Series-based line cards, when the prefix-length is modified from higher value to lower value for an existing prefix-action, heap gets corrupted. Due to this corruption, the FPC might crash anytime when further configurations are added/deleted. The following operations might be considered as a workaround: Step 1. Delete the existing prefix-action and commit Step 2. Then re-create the prefix-action with newer prefix-length [PR1098870](#)
- DHCP End options (option 255) is missing by DHCP-relay agent (where 20 bytes DHCP options 82 inserted) for client DHCP discover message with 19bytes padding. [PR110939](#)
- The monitor component in sshd in OpenSSH before 7.0 on non-OpenBSD platforms accepts extraneous username data in MONITOR_REQ_PAM_INIT_CTX requests, which allows local users to conduct impersonation attacks by leveraging any SSH login access in conjunction with control of the sshd uid to send a crafted MONITOR_REQ_PWNAM request, related to monitor.c and monitor_wrap.c (CVE-2015-6563). [PR116227](#)

Routing Policy and Firewall Filters

- The auto-complete feature is not working for the "show policy" command. [PR471332](#)

Routing Protocols

- This issue reported captures a change in behavior observed from previous releases. The adjacency hold down is taking longer than expected on passive interfaces and subsequently the issue disappears. This will not cause any functionality break since the functionality is restored eventually and seen only on passive interfaces immediately after unified ISSU. [PR780684](#)
- Continuous soft core-dump may be observed due to bgp-path-selection code. RPD forks a child and the child asserts to produce a core-dump. The problem is with route-ordering. And it is auto-corrected after collecting this soft-assert-coredump, without any impact to traffic/service. [PR815146](#)
- This PR fixes a bug by which the receiving EBGp speaker mistakenly accepts a session establishment attempt from an EBGp peer address that is not directly connected because it did not check to see if the address to which peer wants to establish a session belongs to the receiving interface or not. [PR816531](#)
- OSPF route will not be deleted from routing/forwarding table if configuration satisfies below simultaneously. 1. Router ID is not specified and it can be changed due to interface down. 2. There is an interface where OSPF is not running. Suppose OSPF is running on interface A and it is not running on interface B. IP address of interface A is selected as router ID. When interface A goes down and router ID is changed to the IP address of interface B, OSPF on interface A will lose adjacency to the remote OSPF router but router will keep routes learned via OSPF. [PR820909](#)
- In a rare condition, the periodic packet management process (ppmd) may crash during freeing connections. [PR825522](#)
- When a Bidirectional Protocol Independent Multicast (PIM) rendezvous point (RP) is configured on a physical interface, such as fe-0/0/0 not the loopback interface, after restarting the routing, the Reverse Path Forwarding (RPF) interface might not be added

to the accepting interface list for the affected groups, then some traffic can not be forwarded normally. [PR842623](#)

- When pim traceoptions "flag all" and "flag hello disabled" are configured, traces about hello from pppmd are still seen. The work-around is to configure "flag hello detail disabled" as well [PR842627](#)
- When bfd-holddown is configured for BGP, and if interface goes down, bfd session is not destroyed currently, and if BGP-session is brought up with a new interface, then the stale BFD's flapping causes the BGP session to flap. [PR846981](#)
- On Inter-AS PIM SM, where RP and Source are located in another BGP AS domain, if there are duplicate upstream and link flapped happened to primary upstream, multicast will get stuck with secondary upstream path and will not revert back to primary upstream, which will rpf mismatch and traffic outage. [PR847370](#)
- Whenever a configuration change is made and a commit is issued, the Routing Engine's CPU utilization might go up due to BGP reprocessing all the routes. This could happen for any commits unrelated to policy, BGP configuration and most common with scaled BGP environment. [PR853670](#)
- When an import-policy change rejects a BGP-route previously contributing to BGP-Multipath formation, the Peer Active-route-counters in "show bgp neighbor" may not get updated correctly. [PR855857](#)
- When you perform a commit or a commit check of an invalid subnet configuration on a multicast group, the routing protocol process (rpd) crashes, and core files are generated. [PR856925](#)
- Prefixes that are marked with 2 or more route target communities (matching multiple configured targets configured in policies) will be using more CPU resources. The time it takes to process this kind of prefixes depends on the number of VRFs and the number of routes that are sharing this particularity. This can lead to prolonged CPU utilization in rpd. [PR895194](#)
- If Node-link protection is required in case of multiple ECMP primary paths, Node-link protection command: ("set protocols ospf area <area_id> interface <interface_name> node-link-protection") needs to be configured on all the outgoing-interfaces of PLR(Point of Local Repair)node that falls on the ECMP path to the primary. For example in the following diagram: PLR: RTA Destination: RTC Primary paths:
RTA-->lt-1/2/10.102-->RTB-->lt-1/2/10.203-->RTC;
RTA-->lt-1/2/10.122-->RTB-->lt-1/2/10.203-->RTC; Outgoing interfaces on PLR:
lt-1/2/10.102 lt-1/2/10.122 Node-link protection needs to be enabled on both lt-1/2/10.102 and lt-1/2/10.122 if backup route avoiding RTB needs to be computed. (cost 1)
|----|-----lt-1/2/10.102(81.1.2.2)-----|----| || (cost 1) || RTA
|-----lt-1/2/10.122(82.11.22.2)-----| RTB | | | | |lt-1/2/10.203
| 81.3.3.3 || (cost 1000) |----| | |lt-1/2/10.103(81.1.3.1) ----| RTC |-----|
|----| The behavior is corrected from release 14.1 and Node-link protection can be configured on any one of the interfaces on the ECMP path. [PR924290](#)
- When a Junos OS router with multicast enabled receives IGMP packets with protocol DVMRP (IGMP_PROTO_DVMRP) to the IGMP port is 0x5 (DVMRP_ASK_NEIGHBORS2), IGMP builds a neighbor list and responds back to the source IP address of the sender.

This source IP address can be a unicast address or a multicast address. There is no throttling of responses. The requests are answered at the highest rate possible. Secondary impacts are that the routing protocol daemon (rpd) IGMP utilization goes very high and the host path and interface network control queues can get congested. Refer to KB29553 for more information and mitigation. [PR945215](#)

- In rare cases, rpd may write a core file with signature "rt_notbest_sanity: Path selection failure on ..." The core is 'soft', which means there should be no impact to traffic or routing protocols. [PR946415](#)
- RPD crashes during stress test. [PR962610](#)
- In a scaling setup a restart routing or NSR switchover can result in duplicate MSDP entries. [PR977841](#)
- In multicast environment, if one interface is changed from dense mode to sparse mode, the multicast traffic might experience loss on all interfaces for few seconds. [PR999728](#)
- There are two scenarios that the rpd might crash. The first scenario is when all BGP peers flap with bgp route target proxy configured. The second scenario is when BGP session is configured in a way that one side is configured with family l2vpn auto-discovery-only, while on the other side is configured with both family l2vpn signaling and keep all configuration statements. [PR1002190](#)
- With multicast discard route present, if a RP router has no pd- interface, it might not generate (S,G) join to upstream when receiving MSDP source active (SA) message. [PR1014145](#)
- With BGP multipath configured, if a BGP route's multiple protocol nexthops are resolved to different types of IGP routes with a same metric, high rpd process utilization might be observed due to the BGP multipath task. [PR1017372](#)
- On the provider PE in Carrier-of-Carriers VPN scenario, a route in the vrf.inet.3 table is copied to vrf.inet.0 automatically. It is because the provider carrier's iBGP session has family inet-vpn and will only advertise routes from vrf.inet.0. Then the route in vrf.inet.0 is further auto-exported to bgp.l3vpn.0 table. The rpd process might crash when BGP is trying to advertise the route in bgp.l3vpn.0 table while its original route in vrf.inet.0 table is in the middle of deletion. This is a timing issue and not easy to be reproduced. [PR1024470](#)
- The multicast traffic might be pruned with a static IGMP join configuration upon receiving an IGMP leave group message when the interface is not a querier on the corresponding interface. [PR1034270](#)
- IBGP learned route has an indirect next hop and due to bug in code, indirect next hops were not being accepted as valid next hops for MSDP SA functionality. [PR1036415](#)
- Either "rib inet.3" or "resolve-vpn" feature is available to be configured in the lower hierarchy for BGP labeled-unicast family routes. These two features are mutually exclusive and only one of them could be used at a single BGP group. If the administrator swaps the two features (for example, using the "resolve-vpn" first, then deactivate it and using "rib inet.3" instead, then use "resolve-vpn" back), the secondary routes (routes in inet.3 which including the ones from this BGP group and from other BGP

groups) may get accidentally removed every time on "commit" operation take place. [PR1052884](#)

- In multi-topologies IS-IS scenario, there is huge difference between estimated free bytes and actual free bytes when generating LSP with IPv6 Prefix. It might cause LSP fragment exhaustion. [PR1074891](#)
- There are two issues in the PR: (1) In multicast environment, Incoming interface list (IIF) list has only RPF interface, designated forwarder (DF) winners are not added in the list in backup Routing Engine. (2) "Number of downstream interfaces" in show pim join extensive is not accounting Pseudo-VXLAN interface. [PR1082362](#)
- Due to software bug Junos OS cannot purge so-called doppelganger LSP, if such LSP is received over newly formed adjacency shortly after receiving CSNP from the same neighbor. [PR1100756](#)

Services Applications

- When sending traffic through IPsec tunnels for above 2.5Gbps on an MS-400 PIC, the Service-PIC might bounce due to prolonged flow control. [PR705201](#)
- In the Adaptive Service PIC (Service PIC II) scenario, configure the command "root@user# set services service-set <service set name> stateful-firewall-rules", because the command is not supported by 12.1R4, so Adaptive Service PIC goes offline. [PR819833](#)
- In L2TP subscriber management environment, on L2TP Access Concentrator (LAC), L2TP tunnel idle timer is started when the last session on the tunnel is deleted. If the tunnel idle timer expires, then L2TP keeps the tunnels/session/destinations in dying state for the duration of destruct timer (which by default is 5 minutes (300 secs)) before they get destructed. During this phase, jl2tpd process tries to resurrect the tunnel in dying state, causing jl2tpd process to crash and generate a core file. When issue happens, the following logs could be observed: init: l2tp-universal-edge (PID 50230) terminated by signal number 6. Core dumped! /kernel: pid 50230 (jl2tpd), uid 0: exited on signal 6 (core dumped) The impact of l2tpd process crash is, for short period of time, tunneled subscribers cannot connect while processes restarts. Existing connections are not expected to drop. The unexpected result of continued crashes (which have been found in production and been replicated in the lab) is that some subscribers are left in stale states. Subscriber disconnects and reconnects but original session gets stuck on LAC in stale state. This will cause memory jump of many different processes (for example, authd, jpppd, dcd, dfwd, rpd, cosd). [PR824760](#)
- When rollback from v9 to v5 is done, Sampling logic was not rolling back, as sampling registers are not getting released from Packet Forwarding Engine and because in v5 the sampling is Routing Engine based it was not working. [PR824769](#)
- When an MS-DPC PIC reboots due to a crash or manual intervention, it might get stuck in a booting loop if the MS-DPC up-time is more than 49 days and 17 hours. After 5 consecutive boot failures, the MS-DPC PIC will go offline automatically and gives the following error message: [15:21:22.344 LOG: Err] ICHIP(0): SPI4 Training failed while waiting for PLL to get locked, ichip_sra_spi4_rx_snk_init_status_clk [15:21:22.344 LOG: Err] CMSPC: I-Chip(0) SPI4 Rx Sink init status clock failed, cmsdpc_spi4_init [15:21:22.344 LOG: Err] CMX: I(0) ASIC SPI4 init failed [15:21:22.379 LOG: Err] Node for

service control ifl 68, is already present [15:21:23.207 LOG: Err] ASERO SPI-4 XLR source core OOF did not go low in 20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 stop src train failed! [15:21:23.208 LOG: Err] ASERO XLR SPI-4 sink core DPA incomplete in 20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 sink core init failed! [15:21:24.465 LOG: Err] ICHIP(0): SPI4 Stats Unexpected 2'b 11 Error, isra_spi4_parse_panic_errors [15:21:24.465 LOG: Err] ICHIP(0): SPI4 Tx Lost Sync Error, isra_spi4_parse_panic_errors In order to recover from this state the whole MS-DPC needs to be rebooted. [PR828649](#)

- The jnxNatSrcNumPortInuse counter is not refreshing when polling the jnxNatSrcNumPortInuse OID via SNMP after RSP switchover. [PR829778](#)
- In the case of a stateful proxy, SIP hairpinning does not function, and two SIP users behind the NAT device might be unable to connect through a phone call. [PR832364](#)
- With RTSP ALG enabled, RTSP keep-alive packets might be dropped if they are already Ack'ed by the receiver. [PR834198](#)
- On an MS-DPC/service PIC with any of the affected releases and with a NAT hide mode (napt44) configuration, system may generate SIP core in SipFreeXInfo caused by portbmap incorrectly set on us xlate information. SIP transactions with the following characteristics: * the other end is swapping from and to when replying to requests (it should not but some bad implementations do it with little consequences, except for us) * all requests include SDP payload (invites for instance) and are made using the same callID (because they belong to the same dialog or the client was just implemented that way) * from and to tags can change, register callID can also have been different will cause the crash with the reported stack trace. [PR834309](#)
- When DHCP subscribers log in and radius hands down flow-tap variables, the following errors are seen in the log: "/kernel: GENCFG: op 24 (Lawful Intercept) failed; err 5 (Invalid)." [PR837877](#)
- If flow-tap or radius-flow-tap is configured and logging, dynamic flow control daemon (dfcd) may be leaking file descriptors. Over time these leaked file descriptors reach the limit and following error message will be seen. /kernel: kern.maxfiles limit exceeded by uid 0, please see tuning(7). Then routing protocol daemon (rpd) may crash and generate a core file. [PR842124](#)
- On M Series, MX Series, and T Series routers (platforms) with Services PIC, the Session Initiation Protocol (SIP) Application Layer Gateway (ALG) is deployed. When the SIP invite message is received, in rare condition, the Services PIC might crash. [PR843047](#)
- Service PIC might crash in corner cases when receiving specific SIP REGISTER. [PR843479](#)
- When Session Initiation Protocol (SIP) Application Level Gateway (ALG) is enabled, if there are SIP conversations such that a leaf or branch has a master flow (for example, if a SIP message does not match any existing flow), after issuing CLI command "clear services stateful-firewall flows", service PIC might crash and generate a core file. A practical example is a SIP INVITE initiated hierarchy with no reply and retransmissions using a different source UDP port. Clearing the flows in such condition will cause service PIC to crash and generate a core file. [PR845746](#)

- In SIP Application Layer Gateway (ALG) with NAT scenario, "VIA headers" in "183 Session Progress" messages from outside NAT to inside NAT are not getting changed properly, which causes ringback tone failed. [PR845934](#)
- In CGNAT environment, SIP call broken after putting in hold button of a phone which is in inside NAT. For example, as shown here: [Phone A] ----- [router] (NAT) ----
[Phone B] Call from Phone B to Phone A and then push hold button on Phone A. After then 200 OK with Session Description Protocol (SDP) from Phone B but IP address is not getting changed on router which cause SIP call broken. [PR846838](#)
- Service PIC might crash in corner cases when EIM is enabled for SIP ALG. [PR847124](#)
- When allocate the memory from shared memory for bitmaps used in port blocks. Junos OS requests as many bytes as the size of the block. If customers assign for example, 10K block size for deterministic NAT or PBA, then Junos OS allocates 10K bytes for that bitmap. However, it only needs 10K/8 bytes as one byte can represent 8 ports. These huge allocations are leading to memory depletion when many source addresses are behind the NAT, and port blocks are big. [PR851724](#)
- spd core file generated during switchover when CGAT configuration is there. Issue is well understood now and has been fixed in later releases. [PR854206](#)
- In a CGNAT scenario with Port Block/Bucket Allocation (PBA) configured, when the port is exhausted due to receive ICMP or ICMPv6 echo requests fast with changing ID, the services PIC will have no more ports to allocate but create state objects for these newly packets. The state objects then cannot be released anymore, and memory leak will occur. If the service PIC used memory that reaches 2GB, then it will no longer allocate new port blocks and some logs will be seen "port block memory allocation errors". The memory usage of service PIC can be seen by using following command:
user@router> show services nat pool detail Jan 10 11:52:37 Interface: sp-11/0/0, Service set: MOBILE-1 NAT pool: POOL1-MOBILE, Translation type: dynamic Address range: 151.71.180.0-151.71.181.255 Port range: 512-65535, Ports in use: 48, Out of port errors: 196197999, Max ports used: 344898 AP-P out of port errors: 75964912 Max number of port blocks used: 55371, Current number of port blocks in use: 15, Port block allocation errors: 4098769297, Port block memory allocation errors: 196197999 Port blocks limit exceeded errors: 75979500 [PR854428](#)
- MS-DPC may crash in certain scenarios when using CGNAT PBA and junos-rsh, junos-rlogin, junos-rpc-services-udp and junos-rpc-services-tcp ALGs (either one) in combination with EIM. [PR862756](#)
- In IPsec environment, after performing the Routing Engine switchover (for example, performing Graceful Routing Engine Switchover) or chassis reboot (that is, whole device is powered down and powered UP again), due to the key management daemon (kmd) may be launched before the Routing Engine mastership is finalized, it may stop running on the new master Routing Engine. [PR863413](#)
- When DHCP subscribers log in and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: rts_gencfg_dependency_ifstate(): dependency type (2) is not supported." [PR864444](#)

- "replicate-services" configuration command-line interface (CLI) under "set services service-set ..." is hidden, but it can be seen according to "root@user# run show configuration services | display set" [PR930521](#)
- FSAD can sometimes crash on systems with high number of intelligent PICs. These crashes have no impact but creating corefiles for fsad. [PR940522](#)
- 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 below example, the Network Address Translation (NAT) rule and term names will be used to generate an internal jpool with a form : _jpool_{rule_name}_{term_name}. If the generated jpool name exceeds 64 characters in length, it will get truncated. If the truncated jpool name get overlapped with other generated jpool name it will lead to an inconsistent pool usage. user@router# show services nat rule A_RULE_NAME_WHICH_IS_LONG_12345 { ... term A_TERM_ALSO_WITH_LONG_NAME_1 { from { source-address { 10.20.20.1/32; } } then { translated { source-prefix 10.10.10.1/32; <--- translation-type { source static; } } } } term A_TERM_ALSO_WITH_LONG_NAME_2 { from { source-address { 10.20.20.22/32; } } then { translated { source-prefix 10.10.10.2/32; <--- translation-type { source static; } } } } } First jpool = _jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_LONG_NAME_1 > 64 characters. Second jpool = _jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_LONG_NAME_2 > 64 characters. The resulted jpool "_jpool_A_RULE_NAME_WHICH_IS_LONG_1234_A_TERM_ALSO_WITH_" will be used wrongly in both terms. [PR973465](#)
- On M Series, MX Series, T Series routers (platforms) with Services PIC, the incoming interface is a services interface. If the services interface receives "ICMP MTU Exceeded" message, the message might be dropped. [PR977627](#)
- 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](#)
- On MX240/480/960 Series router with MS-DPC with "deterministic-port-block-allocation block-size" configuration. In rare condition, when the "block-size" is set to a larger value (in this case, block-size=16128), the Services PIC might crash. [PR994107](#)
- In the NAT environment, the jnxNatSrcPoolName OID is not implemented in jnxSrcNatStatsTable. [PR1039112](#)

Software Installation and Upgrade

- Filesystem corruption might lead to Routing Engine boot up failure. This problem is observed when directory structure on hard disk (or SSD) is inconsistent. Such a failure should not result in boot up problem normally, but due to the software bug, the affected Junos OS releases mount /var filesystem incorrectly. The affected platforms are M/T/MX/TX/TXP. [PR905214](#)

Subscriber Access Management

- When an MX Series router is acting as the Dynamic Host Configuration Protocol (DHCP) local server and interacting with Session and Resource Control (SRC) for subscriber authorization and provisioning, SRC passes back "framed-ip-address" during subscriber login to the local address pool. In this scenario, the OFFER and ACK messages sent by the MX Series router do not include dhcp-option 1, subnet-mask. [PR851589](#)
- There was a software bug related to shmlog locking that is exposed when too many ifinfo process invocations are triggered by CLI show interface commands. [PR855677](#)
- When the MX Series router acting as the Policy and Charging Enforcement Function (PCEF) uses Gx-Plus to request service provisioning from the Policy Control and Charging Rules Function (PCRF), the authentication service process (authd) might crash during the subscribers logout. [PR1034287](#)
- In subscriber management environment, when performing sustained concurrent login/logout for scaling subscribers, because the rpd may incorrectly destroy the route, the Broadband Edge subscriber management daemon (bbe-smgd) may crash and restart. [PR1064647](#)
- In subscriber management environment, when dual-stack service is activated by the Change of Authorization (CoA) request from the Radius Server, both families will be activated in the same profile response. Due to a software defect, the service accounting session id is not generated properly and the Service Accounting Messages and Interim-updates failed to be sent out. [PR1071093](#)
- On MX Series routers when adding the LI-Action attribute for mirroring the traffic of dual stack subscribers, due to the loop of the service requests lookup and adding, the authentication process (authd) CPU utilization may stay high indefinitely and traffic mirroring is not happening. [PR1077940](#)

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](#)
- 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](#)
- Protected sections of the group hierarchy do not have their protection status displayed correctly and are not prevented from adding new elements into existing groups. [PR717527](#)
- "annotate" is not valid under firewall filter then hierarchy level and displayed "No valid completions", and lead to the configuration could not be committed under "edit private" mode. [edit] user@router# show | compare [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] + /* Don't process the packet here; it's IPv6, not IPv4. + * Accept it and have it be processed by the IPv6 ACL. */ accept; syntax error. user@router# commit full [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] 'accept' outgoing comment does not match patch: [PR812111](#)
- On the J-Web interface, Configure > Routing > OSPF > Add > Interface Tab is showing only the following three interfaces by default: - pfh-0/0/0.16383 - lo0.0 - lo0.16385 To overcome this issue and to configure the desired interfaces to associated ospf area-range, perform the following operation on the CLI: - set protocols ospf area 10.1.2.5 area-range 12.25.0.0/16 - set protocols ospf area 10.1.2.5 interface fe-0/3/1. [PR814171](#)
- When PIM is enabled via apply-groups to one routing-instance whose instance-type is not defined (no-forwarding type is set), incorrect constraint check of PIM will cause routing protocol daemon (rpd) to crash upon any configuration change later. [PR915603](#)

VPNs

- BGP community 0xFF04 (65284) is a well known community (NOPEER), but it is incorrectly displayed as "mvpn-mcast-rpt" in the cli command "show route". This is a show command issue only. No operational mis-behavior will be observed on the router/network. [PR479156](#)
- VPLS traffic gets flooded back over the ingress interface on the local PE as the split-horizon gets disabled upon interface flap. [PR818926](#)
- When a receiver already receiving multicast traffic for a group leaves the group, router connected to the receiver sends a Prune upstream and starts its upstream Prune timer. When the egress PE receives the Prune it will withdraws Type-4 route. During this time, if we 'clear pim join instance vrf' or (set routing-instances vrf protocols pim disable/enable) is done on egress PE and when the Receiver joins the group again, egress PE receives PIM Graft message but, drops it because it does not have matching SG state. This resulting in egress PE not able to get trigger to send Type-4 and thereby is not able to pull traffic from ingress. [PR888901](#)

- For NG-MVPN with RPT-SPT mode, configuring the leaf tunnel limit to block the TYPE 4 routes sending from egress PE after it reaches the configured limit. In certain corner scenarios, the leaf route count is not updated correctly and then the TYPE 4 routes are blocked on receiving the TYPE 3 routes. This timing issue could happen only when a new C-multicast route is bound to the selective provider multicast service interface (S-PMSI) which is marked for deletion however not yet deleted. [PR953449](#)
- In NG-MVPN route-group scenario, configuring leaf tunnel limit on egress PE, flapping the source route on preferred serving site, during the transition from preferred to non-preferred serving site, if the selective provider multicast service interface (S-PMSI) limit count is exceeded temporarily on member site, the member site might fail to originate Type 4 route even though it has Type 5 and Type 3 routes from non-preferred serving site. This is a timing issue and difficult to be reproduced. [PR994687](#)
- In the Rosen MVPN environment, the RP-PE is an assert loser, another PE is sending traffic over the data-mdt. If a new receiver PE with higher rate comes up, because internal workflow processes wrong, the receiver PE might reset data-mdt. This leads to traffic loss. [PR999760](#)
- In NG-MVPN rpt-spt scenario with C-multicast limitation configured, there are some corner circumstances to make the receiver site not receiving all the traffic from the serving site: 1. When the C-multicast limitation is relaxed or some of the routes deleted making way for the new routes to be installed, if MVPN installed a route first, and PIM tries to install the same route prefix, since the limit was reached it was not installing the route. Similarly when MVPN tries to install a route which is installed by PIM already. This was leading to miscounting and hence the egress PE advertising lesser number of routes than limit. 2. If RP PE is also connected to source, the limit should not be applied as it is ingress PE. However when the source route changes to remote, then the limit should be applied as it is egress PE now. The miscounting on RP PE also causing lesser routes to be advertised than configured limit. [PR1001861](#)
- In NG-MVPN scenario with multiple source PEs for a same group, if an inactive source PE has local receivers, the routing protocol process (rpd) on this PE might cause multicast traffic loss and continuous IFF-MISMATCH error. [PR1009215](#)
- In the 12.3 release after issuing a "request pim multicast-tunnel rebalance" command the software may place the default encapsulation and decapsulation devices for a Rosen MVPN on different tunnel devices. [PR1011074](#)
- Problem Description The problem is that MSDP is periodically polling PIM for S,G's to determine if the S,G is still active. This check helps MSDP determine if the source is active and therefore the SA still be sent. There is a possibility that PIM will return that the S,G is no longer active which causes MSDP to remove the MSDP state and notify MVPN to remove the Type 5. One of the checks PIM makes is to determine if it is the local RP for the S,G. During a re-configuration period where any commit is done, PIM re-evaluates whether it is a local RP. It waits until all the configuration is read and all the interfaces have come up before making this determination. The local rp state is cleared out early in this RP re-evaluation process, however, which allows for a window of time where the local RP state was cleared out but it has not yet been re-evaluated. During this window PIM may believe it is not the local rp and return FALSE to MSDP for the given source. If MSDP makes the call into PIM during this window after a configuration change(commit), then it is possible that the Source Active(Type 5) state

will be removed. Fix The fix will be to clear out the local rp state right before it is re-evaluated ie after it reads configuration for all interfaces; to not allow any time gap where it could be inconsistent. [PR1015155](#)

- In PIM Draft-Rosen Multicast VPN (MVPN) environment, in a setup where active C-PR, standby C-RP, C-receivers, and C-source are located in different VPN sites of MVPN instance, once the link to active C-RP is flapped, the PE router which connects to C-receivers would send (*g) join and (s,g,rpt) prune towards standby C-RP. When the PE that connects to standby C-RP receives the (*g) join and (s,g,rpt) prune over mt-, it ends up updating the (s,g) forwarding entry with mt- as downstream, which is already the incoming interface (IIF). This creates a forwarding loop due to missing check if IIF is same as OIF when PIM make-before-break (MBB) join load-balancing feature is enabled. As a result traffic gets looped back into the network. Loop once formed will remain at least for 210 seconds till the delayed prune timer expires. After this, IIF is updated to the interface towards standby C-RP. [PR1085777](#)
- In NG-MVPN spt-only mode with a PE router acts as the rendezvous point (RP), if there are only local receivers, the unnecessary multicast traffic continuously goes to this RP and gets dropped though it is not in the shortest-path tree (SPT) path from source to receiver. [PR1087948](#)

Related Documentation

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)
- [Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 181](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 357](#)

Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

The current software release is Release 12.3. For information about obtaining the software packages, see “[Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers](#)” on page 357.

- [Resolved Issues on page 249](#)
- [Previous Releases on page 251](#)

Resolved Issues

Class of Service (CoS)

- In SNMP environment, when performing multiple walks or parallel snmpget for same interface at the same time (for example, SNMP bulk get/walk, or SNMP polling from multiple devices) on CoS related MIBs (jnxCos table), if the interface state changes or the request gets timeout when FPC is responding the request, memory leak of Class-of-Service process (cosd) about 160 bytes (up to 1500 bytes) may occur, which may cause cosd to crash eventually when limit is exceeded. [PR1058915](#): This issue has been resolved.

Forwarding and Sampling

- In rare cases, MX Series routers might crash while committing inline sampling related configuration for INET6 Family only. [PR1091435](#): This issue has been resolved.

General Routing

- Upon a `write core` on an FPC in T-series (Gimlet/Stoli) since Junos 12.3 (PR/843389), we will also dump the contents of the R/SR ASIC memory (Jtree SRAM). If there is a memory parity error present in the SRAM then the coredump will abort and the FPC will crash unless `set chassis pfe-debug flag disable-asic-sram-dump` is configured. [PR1105721](#): This issue has been resolved.

Interfaces and Chassis

- When the Ethernet Link Fault Management (LFM) action profile is configured, if there are some errors (refer to the configuration, for example, frame errors or symbol errors) happening in the past (even a long past), due to the improper handling of error stats fetching from kernel, the LFM process (lfmd) may generate false event PDUs and send the false alarm to the peer device. [PR1077778](#): This issue has been resolved.
- After removing a child link from AE bundle, in the output of "show interface <AE> detail", the packets count on the remaining child link spikes, then if you add back the previous child link, the count recovers to normal. [PR1091425](#): This issue has been resolved.

MPLS

- Junk characters are being displayed in output of "show connections extensive" command. [PR1081678](#): This issue has been resolved.

Platform and Infrastructure

- When the system parses IPv6 packet, it does not copy the GRE key id from the GRE header to the packet context. Due to this, the IPv6 GRE header hash value is not computed correctly, hence the session lookup based on the hash value fails and duplicate sessions are created. [PR916028](#): This issue has been resolved.
- VRRP advertisements might be dropped after enabling delegate-processing on the logical tunnel (lt) interface. It would result in VRRP master state observed on both routers. [PR1073090](#): This issue has been resolved.

- On MX Series router, if ifl (logical interface) is configured with VID of 0 and parent ifd (physical interface) with native-vlan-id of 0, when sending L2 traffic received on the ifl to Routing Engine, the VID 0 will not be imposed, causing the frames to get dropped at Routing Engine. [PR1090718](#): This issue has been resolved.

Routing Protocols

- RIP is applying the RIB import-policy for the primary RIB table. As per the policy configured, evaluation fails and routes are removed from primary RIB. But import-policy is applied only for secondary tables. RIP should apply only to the protocol import policy and add routes to primary RIB. Routes are leaked to secondary routing table according to import-policy. As suggested by rpd infrastructure team, removed the import policy filter application to primary routing table by protocol rip. Now import policy application is handled by policy module within RPD. [PR1024946](#): This issue has been resolved.

Services Applications

- In Network address translation (NAT) environment, if the translation type is "dynamic-nat44", when processing bursts of packets (for example, packets coming in one after the other at a delay close to the interpacket gap, based on the replication, issue was seen when processing line rate traffic on a 1GE port), because the device may fail to free up dynamic NAT addresses, the address pool may get exhausted quite fast. In this situation, since no new addresses could be allocated, the incoming traffic drop might be seen on the device, also memory leak (not the main issue here, may become a problem way after the addresses leak) may occur on the device. [PR1098583](#): This issue has been resolved.

Previous Releases

Resolved Issues in Release 12.3R10

Class of Service (CoS)

- Add chassis scheduler map support on gr interface on MS-PIC, which means there will be no commit error if scheduler-map-chassis is applied on gr interface. [PR1066735](#): This issue has been resolved.

Forwarding and Sampling

- If the template of the policer is changed (for example, change the bandwidth-limit value of policer), shared-bandwidth-policer knob may not function properly anymore. [PR1056098](#): This issue has been resolved.

General Routing

- On dual Routing Engine platforms, after performing unified graceful Routing Engine switchover (GRES) with 8K subscribers, the ksyncd process may crash due to the replication error on a next hop change operation. The issue is hit when there is memory pressure condition on the Routing Engine and in that case, it may lead to null pointer de-reference and ksyncd crash. Or in some case, the kernel on the new master Routing Engine might crash after Routing Engine switchover if Routing Engine is under memory pressure due to missing null check when trying to add a next hop and the next hop is not found at the time. [PR942524](#): This issue has been resolved.
- 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 backup

Routing Engine might bounce the BGP sessions on the master Routing Engine. [PR942981](#): This issue has been resolved.

- In a scaled network (for example, more than 100K unicast next-hop entries) with BGP ECMP configured, when the master Routing Engine crashes and switchover happens, during when the neighbor session of BGP over aggregate Ethernet (AE), interface might get broken. This is because the unicast next-hops of the AE is stuck at standby state and therefore no traffic can be transmitted through. [PR953365](#): This issue has been resolved.
- The configuration statement 'gratuitous-arp-on-ifup' should send a gratuitous arp on each unit of a physical interface, but in Junos OS Release 12.3 and higher versions, only the first unit is seeing the configured behavior. [PR986262](#): This issue has been resolved.
- Whenever the logical tunnel (lt-) interface with IPv6 family configured goes down and comes up upon hardware initialization (MPC/FPC replacement/reboot or chassis reboot), due to Duplicate Address Detection (DAD), functionality is not being performed for the logical interface up/down event. The "lt-" interface may get stuck in "tentative" state and thus IPv6 traffic cannot pass over it. [PR1006203](#): This issue has been resolved.
- On MX Series router with IPv6 subscribers, after performing GRES or reloading one line card that has underlying interfaces for demux, some demux interfaces might get stuck in Tentative state, and some other demux interfaces which have the same link local addresses might be unable to send any IPv6 RA message. [PR1026724](#): This issue has been resolved.
- Configuring a routing policy with the "no-route-localize" option to ensure that the routes matching a specified filter are installed on the FIB-remote Packet Forwarding Engines, after removing the routing policy and changing the next-hop for the routes, the previously installed routes using "no-route-localize" policy will not get removed from PFE 1 but will from PFE 0 on the same FPC. Then traffic received on PFE 1 will not forward received packets to the FIB-local PFEs to perform full IP table lookup but using the staled routes instead. This situation does also apply if the interface is getting disabled. If traffic destined to the local-address is still received on PFE 1, those stale route lookup entries might have incorrect entries and might lead to one of the following possible symptoms. fpc1 RCHIP(1): 8 Multicast list discard route entries fpc1 PFE: Detected error nexthop: fpc1 RCHIP(1): RKME int_status 0x10000000 RKME and Detected error nexthop will per default will trigger a FPC restart. [PR1027106](#): This issue has been resolved.
- On the Type 5 PIC, when the "hold-time down" of the interface is configured less than 2 seconds and the loss of signal (LOS) is set and cleared repeatedly in a short period (for example, performing ring path switchover within 50ms), the "hold-time down" may fail to keep the interface in "up" state within the configured time period. [PR1032272](#): This issue has been resolved.
- This issue is applicable to a case which inline NAT configured on an interface belongs to either an MPCE or an MP3E/MPC4E/T4000-FPC5. Ingress and egress traffic traversing between an MPCE and these cards may cause the router to drop packets. [PR1042742](#): This issue has been resolved.

- When querying specific entries of the JUNIPER-SUBSCRIBER-MIB, memory leak may occur on the smihelperd process which provides the necessary information over SNMP. [PR1048469](#): This issue has been resolved.
- This problem is because of a race condition, where other FPCs are not able to drain "which is 1 second" Fabric Streams connecting to FPC which is getting offline. With this situation - even when FPC comes online, other FPCs which have observed message "xmchip_dstat_stream_wait_to_drain" will not be able to send traffic to that particular FPC over fabric. There is no workaround. Rebooting FPCs which observed error message "xmchip_dstat_stream_wait_to_drain" is a recovery. [PR1052472](#): This issue has been resolved.
- In subscriber management scenario, when dynamic VLAN (DVLAN) demux interface is configured on MX Series router, the interface may get in stuck state, and it could be observed that the statistics of demux0 may stop incrementing. This is because Session Database (SDB) may incorrectly calculate the number of subscriber over DVLANS. When the issue occurs, for example, the router may not be able to process any PPPoE Active Discovery Initiation (PADI) packets, and fail to establish the PPPoE session. [PR1054914](#): This issue has been resolved.
- In subscriber management environment, the Berkeley Database (DB) may get into deadlock state. It is brought on by multiple daemons attempting to simultaneously access or update the same subscriber or service record. In this case, due to the access to DB was blocked by device control daemon (dcd), the subscriber management infrastructure daemon (smid) fails to recover the DB. Consequently, the router may stop responding to all the login/logout requests as well as statistics activity. This timing related issue is most likely to occur during login or logout and when the system is busy. [PR1054292](#): This issue has been resolved.
- OpenSSL project has published a security advisory for vulnerabilities resolved in the OpenSSL library on January 8th 2015: CVE-2014-3569, CVE-2014-3570, CVE-2014-3572, CVE-2014-8275, CVE-2015-0204, CVE-2015-0205. Refer to JSA10679 for more information. [PR1055295](#): This issue has been resolved.
- Class 4 (32W) Optics are not supported on MPC4E (2CGE+8XGE). Upon insertion and removal of a Class 4 optic, the TX laser will remain powered-off, even when a supported optic is inserted. [PR1068269](#): This issue has been resolved.
- If a subscriber-facing AE interface has child links which spread over multiple Packet Forwarding Engines on a single FPC, when subscribers attempt to login, "LUCHIP Congestion Detected" error messages will be seen periodically and there might be some potential forwarding issues for subscribers. [PR1069292](#): This issue has been resolved.
- In subscriber management environment, changing the system time to the past (for example, over one day) may cause the daemons (for example, pppoe, and autoconfd) that are using the time to become unresponsive. [PR1070939](#): This issue has been resolved.
- In scaled subscriber management environment (for example, 3.2K PPPoE subscribers), after heavy login/logout, the session setup rate keeps decreasing and also PAP-NAK messages are sent with "unknown terminate code". This continues till Broadband Network Gateway (BNG) does not accept PPP sessions and all newly incoming sessions

are stuck in PAP Authentication phase (No PAP ACK received). [PR1075338](#): This issue has been resolved.

- On MX Series routers, the CLI command “set interfaces interface-name speed auto-10m-100m” is not supported. [PR1077020](#): This issue has been resolved.

High Availability (HA) and Resiliency

- When LACP is configured in 'periodic fast' mode, the traffic loss of more than 30 seconds will be seen during unified ISSU. The workaround is to change LACP to 'periodic slow' mode before unified ISSU. [PR1059250](#): This issue has been resolved.

Interfaces and Chassis

- Multicast traffic may not be forwarded to the "Downstream Neighbors" as reported by the command "show pim join extensive". There can be occasions where this traffic is blackholed and not forwarded as expected. Alternatively, there may be an occasion where multicast traffic is internally replicated infinitely, causing one or more of the "Downstream Neighbors" to receive multicast traffic at line rate. [PR944773](#): This issue has been resolved.
- On standalone T Series router or TX platform, during Routing Engine rebooting, a bad (or busy) I2C device on Switch Interface Board (SIB) might cause Switch Processor Mezzanine Board (SPMB) to crash. Please note the TXP platform might also experience same issue due the bad I2C, and it has been addressed in another PR, which has been fixed on Junos OS Releases 13.1R5, 13.2R6, 13.3R1, 13.3R4, 14.1R3, 14.2R1, and 15.1R1. [PR1010505](#): This issue has been resolved.
- On MX240 or MX480, the SNMP_TRAPs may not be generated when Exhaust temperature for an FRU exceeds threshold. [PR1012497](#): This issue has been resolved.
- In case the IQ2 or IQ2E PIC is working in tunnel-only mode, rebooting the tunnel PIC while the traffic is passing through the tunnel might cause the tunnel PIC to not transfer traffic any more. [PR1041811](#): This issue has been resolved.
- jpppd daemon ran out of memory as subscribers login failed due to missing CoS parameters. Following logs will be seen in messages when the subscribers login fail.
Nov 16 12:19:21 jtac-host jpppd: Semantic check failed for profile=PPPoE-1-QoS, error=301
Nov 16 12:19:21 jtac-host jpppd: dyn_prof_send_request: add pre_processing failure, error=301
Nov 16 12:19:21 jtac-host jpppd: Profile: PPPoE-1-QoS variable: \$junos-cos-shaping-rate value: failed semantic check. [PR1042247](#): This issue has been resolved.
- It is observed that the syslog messages related to kernel and Packet Forwarding Engine may get generated at an excessive rate, especially in subscriber management environment. Most of these messages may appear repeatedly, for example, more than 1.5 million messages may get recorded in 2 hours, and there are only 140 unique messages. Besides, these messages are worthless during normal operation and due to the excessive rate of log generation, high Routing Engine CPU consumption (for example, Routing Engine CPU utilization can be stuck at 100% for a long time (minutes or hours), it depends on the activity of subscribers (frequency of logins and logouts) and on the AI scripts used by the customer) by event process (eventd) might be observed on the device. [PR1056680](#): This issue has been resolved.

- In multichassis link aggregation groups (MC-LAGs) environment, the MC-LAG peers have the MAC and port information and can forward the traffic appropriately. If a single VLAN on ICL interface is modified to a different VLAN, and then the administrator rolls back the VLAN configuration to the original one, the remote MAC might be stuck in the "Pending" state and not be installed in the bridge MAC-table, which causes the traffic forwarding to be affected. [PR1059453](#): This issue has been resolved.
- In scaling PPP subscriber environment, when the device is under a high load condition (for example, high CPU utilization with 90% and above), the long delay in session timeout may occur. In this situation, the device may fail to terminate the subscriber session (PPP or PPPoE) immediately after three Link Control Protocol (LCP) keepalive packets are missed. As a result, subscriber fails in reconnect due to old PPP session and corresponding Access-Internal route are still active for some time. In addition to this, it is observed that the server is still sending KA packets after the session timed out. [PR1060704](#): This issue has been resolved.
- Error message is continuously logged every second after a particular copper-SFP [P/N:740-013111] is plugged into a disabled port on MIC. ***** error message *****
mic_sfp_phy_program_phy: ge-*/*/ - Fail to init PHY link mic_periodic_raw: MIC(*/*)
- Error in PHY periodic function PQ3_IIC(WR): no target ack on byte 0 (wait spins 2)
PQ3_IIC(WR): I/O error (i2c_stat=0xa3, i2c_ctl[1]=0xb0, bus_addr=0x56)
mic_i2c_reg_set - write fails with bus 86 reg 29 mic_sfp_phy_write:MIC(*/*) - Failed to write SFP PHY link 0, loc 29 mic_sfp_phy_mdio_sgmiilnk_op: Failed to write: ifd = 140 ge-*/*/; phy_addr: 0, phy_reg: 29 ala88e1111_reg_write: Failed (20) to write register: phy_addr 0x0, reg 0x1d Fails in function ala88e1111_link_init. [PR1066951](#): This issue has been resolved.

Layer 2 Features

- The Layer 2 Control Protocol process (l2cpd) leaks memory when interface config is applied to LLDP-enabled interfaces using 'apply-groups'. Size of the leak is ~700 bytes per commit. [PR1052846](#): This issue has been resolved.
- LACP partner system ID is shown wrong when the AE member link is connected to a different device. This might be misleading while troubleshooting the LAG issues. [PR1075436](#): This issue has been resolved.

Network Management and Monitoring

- SNMP mib walk jnxMac does not return value with et- interfaces on MPC3/MPC4/MPC5/MPC6. [PR1051960](#): This issue has been resolved.
- SNMP queries for LAG MIB tables while LAG child interface is flapping, may cause mib2d to grow in size and eventually crash with a core file. Mib2d will restart, and recover by itself. [PR1062177](#): This issue has been resolved.

Platform and Infrastructure

- For inline BFD over aggregated Ethernet (AE) interface in which member links are hosted on different FPCs, BFD packets coming on ingress line card will be steered to anchor Packet Forwarding Engine through fabric. If FPC reconnect to master Routing Engine (such as Routing Engine switchover operation), the inline BFD session punts the BFD packet to the host. The BFD packet should go through loopback interface filter of VRF on which it is received. But in this case, the BFD packet might hit the incorrect loopback interface filter from wrong routing-instance since the VRF information is not carried across fabric. [PR993882](#): This issue has been resolved.
- In EVPN scenario, MPC may crash with core-file when any interface is deleted and add that interface to an aggregated Ethernet bundle or change the ESI mode from all-active to single-active. [PR1018957](#): This issue has been resolved.
- If several aggregates are configured with shared-bandwidth-policer and those aggregates share the same Packet Forwarding Engine for child member links and one member links flaps, all traffic might get policed and dropped. The traffic dropped might not be on the bundle whose child member link flapped. [PR1035845](#): This issue has been resolved.
- MSDPC-HTTP redirect stops working. [PR1039849](#): This issue has been resolved.
- When IRB interface is configured with VRRP in Layer 2 VPLS/bridge-domain, in corner cases IRB interface may not respond to ARP request targeting to IRB sub-interface IP address. [PR1043571](#): This issue has been resolved.
- Due to a defect in the Junos OS Software, when a telnet user experiences some undefined network disconnect, .perm and .env files under /var/run are left behind. This scenario happens only under certain unknown ungraceful network disconnects. When considerable number of .perm/.env files get accumulated under /var/run, issue is seen with telnet users, that they are not able to perform permitted operations on the router, post-login. [PR1047609](#): This issue has been resolved.
- On the MX Series-based line cards, if inline Network Address Translation (NAT) service, Generic Routing Encapsulation (GRE) tunneling and packets fragmentation are performed on the same Packet Forwarding Engine (specifically, after NAT, the packet go to tunnel and then to fragmentation), the fragmented packets may get dropped by FTP client due to the incorrect TCP checksum of the fragmented packet. [PR1051144](#): This issue has been resolved.
- NTP.org has published a security advisory for multiple vulnerabilities resolved in ntpd (NTP daemon) that have been assigned four CVE IDs. Junos OS has been confirmed to be vulnerable to one of the buffer overflow vulnerabilities assigned CVE-2014-9295 which may allow remote unauthenticated attackers to execute code with the privileges of ntpd or cause a denial of service condition. Refer to JSA10663 for more information. [PR1051815](#): This issue has been resolved.
- Software upgrade might cause firewall filters to redirect packets to an incorrect routing instance. This issue only affects Junos OS Releases 12.3R7, 12.3R7-S1, 12.3R7-S2, 12.3R7-S3, and 12.3R8. [PR1057180](#): This issue has been resolved.

- Customers without the fix cannot use 'then decapsulate gre routing instance <routing instance>' without also another action such as 'then [count | sample]' or 'then decapsulate gre [sample <protocol>] [forwarding-class <fc>] [no-decrement-ttl]'. If they do they could lose all traffic if the board reboots or if the configuration for the routing instance and the filter are both applied at the same time. [PR1061227](#): This issue has been resolved.
- On MX Series routers with MPCs and T4000 routers with Type 5 FPCs, the feature "enhanced-hash-key" is configured to select data used in the hash key for enhanced IP forwarding engines. If "type-of-service" is configured at the [edit forwarding-options enhanced-hash-key family inet] hierarchy level, or "traffic-class" is configured at the [edit forwarding-options enhanced-hash-key family inet6] hierarchy level, the last significant 2 bits of the TOS/TC bytes under the IPv4/IPv6 header are extracted incorrectly as load sharing input parameters, this might cause unexpected load balancing result. [PR1066751](#): This issue has been resolved.
- An FPC with interfaces configured as part of an Aggregated Ethernet bundle may crash and reboot when the shared-bandwidth-policer is configured as part of the firewall policer. [PR1069763](#): This issue has been resolved.
- VPLS filter applied under forwarding-options might drop VPLS frame unexpectedly when it's coming from an lt- interface. [PR1071340](#): This issue has been resolved.
- When inline-sampling is enabled, in race conditions, if packet gets corrupted and the corrupted packet length shows 0, it may cause "PPE_x Errors thread timeout error" and eventually cause MPC card to crash. [PR1072136](#): This issue has been resolved.
- MAC filter ff:ff:ff:ff:ff:ff is cleared from the Packet Forwarding Engine hardware mac table. So arp requests are not forwarded to irb. Not all mac entries pointing to invalid l2 token are candidates for being deleted. Static mac entries are managed by Control plane only. So Packet Forwarding Engine cannot delete these entries. The logic for skipping mac deletion for static mac entries done earlier is not proper Packet Forwarding Engine. Fixed the same. [PR1073536](#): This issue has been resolved.
- On MX Series-based line cards, when the firewall filters with prefixes are configured, the heap memory leak issue might be observed. [PR1073911](#): This issue has been resolved.
- With MSDPC equipped on BNG, there might be a memory leak in ukernel, which eventually causes MSDPC to crash and restart. [PR1085023](#): This issue has been resolved.

Routing Protocols

- When a BGP peer goes down, the route for this peer should be withdrawn. If it happens that an enqueued BGP route update for this peer has not been sent out, issuing the CLI command "show route advertising-protocol bgp <peer addr>" might crash the routing protocol process (rpd). This is a very corner issue and hardly to be experienced. [PR1028390](#): This issue has been resolved.
- Junos OS Multicast Source Discovery Protocol (MSDP) implementation is closing an established MSDP session and underlying TCP session on reception of source-active TLV from the peer when this source-active TLV have an "Entry Count" field of zero. "Entry Count" is a field within SA message which defines how many source/group tuples are present within SA message. [PR1052381](#): This issue has been resolved.

- When running Simple Network Management Protocol (SNMP) polling to specific IS-IS Management Information Base (MIB) with invalid variable, it will cause routing protocol process (rpd) crash. [PR1060485](#): This issue has been resolved.
- When there are a number of secondary BGP routes in inet.0, an SNMP walk of inet.0 by the bgp4 MIB can cause a core if the corresponding primary routes are being deleted. [PR1083988](#): This issue has been resolved.

Services Applications

- On M Series, MX Series, T Series routers with Multiservices 100, Multiservices 400, or Multiservices 500 PICs with "dump-on-flow-control" configured, if prolonged flow control failure, the coredump file might generate failure. [PR1039340](#): This issue has been resolved.
- When the tunnel between L2TP access concentrator (LAC) and L2TP network server (LNS) is destroyed, the tunnel information will be maintained until destruct-timeout expire (if the destruct-timeout is not configured, the default value is 300 seconds). If the same tunnel is restarted within the destruct-timeout expire, the LNS will use the previously negotiated non default UDP port, which might lead to the tunnel negotiation failure. [PR1060310](#): This issue has been resolved.
- A Layer 2 Tunneling Protocol daemon (l2tpd) crash is seen sometimes when the L2TP service interface unit number is configured higher than 8192. A restriction has been added to force unit numbers below 8192. [PR1062947](#): This issue has been resolved.
- Service PIC daemon (spd) might crash with core-files due to CGNAT pool's snmp-trap-thresholds configuration. [PR1070370](#): This issue has been resolved.

Subscriber Access Management

- The authd process memory leak slowly when subscribers login and logout, which eventually leads the process to crash and core-file. [PR1035642](#): This issue has been resolved.

VPNs

- In MVPN RPT-SPT mode, with a mix of local and remote receivers all using (*,g) joins (spt-threshold infinity), the downstream interfaces may not get updated properly and there may be a stuck (s,g) forwarding route. This issue can occur with the following sequence of events: 1. Local receivers are joined 2. Traffic starts, then stops, and the route times out. 3. Remote receiver joins. Both a (*,g) and an (s,g) forwarding route are created. 4. Another local receiver is joined, or an existing one is pruned. 5. In the (*,g) route the downstream interface list reflects the update, but in the (s,g) route the downstream interface list does not. 6. When traffic starts again, the (s,g) route -- which has the wrong interface list -- is used. The traffic flows to the wrong set of receivers. [PR1061501](#): This issue has been resolved.

Resolved Issues in Release 12.3R9**Class of Service (CoS)**

- SNMP get-request for OID jnxCosIngressQstatTxedBytes (ingress queue) might return the value of jnxCosQstatTxedBytes (egress queue). But SNMP walk works fine since it uses get-next-request. [PR1011641](#): This issue has been resolved.
- For ichip based platform, IQ2 pic expects FC index in the cookie from ichip for packet queuing. For Transit traffic, fc index is coming in cookie where are for host outbound traffic, queue number is coming in cookie to IQ2 pic. As IQ2 pic is not aware of whether traffic is transit or host outbound, it treats value received in cookie as FC value and looks into fc_to_q table to fetch queue number. This is causing issue in queueing of host outbound traffic in IQ2 PIC in incorrect queue. This is a day one issue and will come if in FC to Queue mapping, fc id and queue number are not same. [PR1033572](#): This issue has been resolved.

Forwarding and Sampling

- In rare condition, dfwd process might crash during user logging out and in. [PR982477](#): This issue has been resolved.
- When a firewall filter has one or more terms which have MX Series routers with MPCs or MICs match condition or actions, such filters will not be listed during SNMP query. This behavior is seen typically after Routing Engine reboot/upgrade/master-ship switch. Restarting mib2d process will cause to learn these MX Series routers with MPCs or MICs filters: cli > restart mib-process After mib2d restart, SNMP mib walk of firewall OIDs will: - list all the OIDs corresponding this MX Series routers with MPCs or MICs filter - count correctly as configured in the filter Now, despite the SNMP mib walk for firewall OIDs lists all OIDs and appropriate values, messages logs will report the following logs for every interface that has this MX Series routers with MPCs or MICs filter applied. > Jul 8 15:52:09 galway-re0 mib2d[4616]:
%DAEMON-3-MIB2D_RTSLIB_READ_FAILURE: get_counter_list: failed in reading

counter names ae33.1009-i: 288 (No such file or directory) > Jul 8 15:52:09 galway-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE: get_counter_list: failed in reading counter names ae31.1004-i: 257 (No such file or directory) > Jul 8 15:52:09 galway-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE: get_counter_list: failed in reading counter names ae33.1010-i: 289 (No such file or directory) > Jul 8 15:52:09 galway-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE: get_counter_list: failed in reading counter names ae31.1004-i: 257 (No such file or directory) The above two issues are addressed in this PR fix. [PR988566](#): This issue has been resolved.

- On the 32-bit Junos OS, when a very big burst-size-limit value (2147492676 and above) is configured in the ingress interface policer, the kernel may drop Routing Engine destined traffic. [PR1010008](#): This issue has been resolved.
- When an ARP policer is applied to an interface, it appears commented out in the configuration with the following message: "invalid path element 'disable_arp_policer'". [PR1014598](#): This issue has been resolved.

General Routing

- Changing the static route configuration from next-hop to qualified-next-hop might result in static route getting missed from routing table. Restarting routing process can bring back the routes but with rpd core. [PR827727](#): This issue has been resolved.
- In this scenario the CPCD (captive-portal-content-delivery) is configured for HTTP-REDIRECT for Subscriber Management clients using MS-DPC. When services sessions start to redirect the HTTP traffic, the memory-usage consistently increments for MSPMAND on the multi-service PIC. The memory limit then might cause packets loss. [PR954079](#): This issue has been resolved.
- In large scale L3VPN environment (in this case, there are 80K L3VPN routes) with Non-Stop Routing (NSR) enabled. When the L3VPN routes are added and deleted frequently, in rare condition, the Composite Next Hop (cnh) deletion from kernel after backup rpd process learns cnhs with duplicated key but with different nhids, this might lead to rpd process crash on backup Routing Engine. This issue is not reproducible and only happened once. [PR959331](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- In the dual Routing Engines scenario with 8K PPP dual stack subscribers. In rare condition, after Routing Engine switchover, some subscribers are stuck in terminating state forever. [PR974300](#): This issue has been resolved.
- In the dual Routing Engines scenario with large scale nexthops (in this case, more than 1-million nexthops and around 8K VRFs). In rare condition, kernel might crash on backup and/or master Routing Engine due to exhaustion of nexthop index space. [PR976117](#): This issue has been resolved.
- On MX Series router with MX Series linecard. When a faulty Non-Ethernet clear channel OCx MIC is inserted into MPC, the MPC goes offline. After removing the MIC, the MPC

starts going online and offline continuously. These MICs as below belong to SONET/SDH OC3/STM1 (Multi-Rate) MIC: * MIC-3D-8OC3OC12-4OC48 * MIC-3D-4OC3OC12-1OC48 * MIC-3D-8CHOC3-4CHOC12 * MIC-3D-4CHOC3-2CHOC12 * MIC-3D-8DS3-E3 * MIC-3D-8CHDS3-E3-B * MIC-3D-1OC192-XFP. [PR976675](#): This issue has been resolved.

- With nonstop active routing (NSR) enabled, deleting non-forwarding routing instance might result in the rpd process crash on backup Routing Engine. The core files could be seen by executing CLI command "show system core-dumps". [PR983019](#): This issue has been resolved.
- On MX Series, delete an interface A from routing-instance VRF1; then create routing-instance VRF2 and interface A is added to VRF2 with qualified-next-hop configured; finally, delete VRF1. Commit the entire above configuration once, in rare condition, rpd might crash. [PR985085](#): This issue has been resolved.
- 1) Due to a previous fix chassisd on the protocol master Routing Engine and the protocol backup Routing Engine connect to the main snmpd on the protocol master using the following methods. a) Chassisd on the protocol master Routing Engine connects using a local socket since snmpd is running locally. b) Chassisd on the protocol backup RE connects using a TNP socket since snmpd is not local. 2) However this fix changed the way the other daemons connect to snmpd. All important daemons run on the protocol master and should connect to snmpd using a local socket. However the fix changed it so that all daemons that ran on the protocol master (other than chassisd) tried to connect using the TNP socket. SNMPD does not accept these connections. As a fix, in an MX-VC, we made sure that chassisd connects to all processes which run on the protocol master using internal socket while the chassisd process on the protocol backup and protocol linecard connect using TNP socket. [PR986009](#): This issue has been resolved.
- In 6PE scenario, when PE router is sending IPv6 TCP traffic to MPLS core, in rare occasions, the kernel might crash and reboot with a vmcore file generated. [PR988418](#): This issue has been resolved.
- In the dual Routing Engines scenario with NSR configuration, backup peer proxy thread is hogging CPU for more than 1 second if there are multiple updates (>5000) going from master Routing Engine to backup Routing Engine. This leads to FPC socket disconnections. The traffic forwarding might be affected. [PR996720](#): This issue has been resolved.
- When having ECMP routes and multiple levels of route/next-hop recursion, a particular sequence of routes churn may result in rpd process crash and traffic outage. [PR1006523](#): This issue has been resolved.
- MS-DPC memory leak on system service set when HTTP Redirect attempts to process none-HTTP traffic with HTTP ports (80/8080/443). [PR1008332](#): This issue has been resolved.
- When deleting a routing-instance or making changes to the routing-instance, the deletion of the routing-instance to kernel might come before the deletion of the logical interfaces in the routing-instance, resulting in rpd crash. This is a timing issue, hard to reproduce. [PR1009426](#): This issue has been resolved.

- Whenever a FPC goes down suddenly due to hardware failure on T Series router, the data traffic in transit towards this FPC from other Enhanced Scaling (ES) FPCs could be stuck in the fabric queue thereby triggering fabric drops due to the lack of buffers to transmit the data to active destination FPCs. [PR1009777](#): This issue has been resolved.
- On MX Series platforms with ADPC FPCs, M120 or M7i/M10i with Enhanced CFEB each VPLS LSI interface flapping triggers a memory leak in jtree segment 0. There is no memory leak in FPC heap 0 memory. [PR1009985](#): This issue has been resolved.
- Unknown unicast flood is seen with interface flap after router reboot and with static MAC, no-mac-learning, interface-mac-limit configured for a virtual-switch. [PR1014222](#): This issue has been resolved.
- The routing protocol daemon (rpd) might crash continuously with core-dumps upon adding a sub-interface with "disable" configuration to a MC-LAG interface. [PR1014300](#): This issue has been resolved.
- The OpenSSL project released a security advisory on 2014-08-06 that contained nine security issues. The following four issues affect Junos: CVE-2014-5139: Crash with SRP ciphersuite in Server Hello message CVE-2014-3509: Race condition in ssl_parse_serverhello_tlsextr CVE-2014-3511: OpenSSL TLS protocol downgrade attack CVE-2014-3512: SRP buffer overrun See JSA10649 for more information. [PR1016458](#): This issue has been resolved.
- The existing PTSP_SUBS_TIMEOUT_NONE is within the max of 86400 (PTSP_SUBS_PKT_TIMEOUT_INTVL) range. Hence subscriber may not be cleaned from pkt_timeout table when the time interval of 0xFFFF is hit. Defined a separate value to clean the packet idle timeout subscribers. [PR1016896](#): This issue has been resolved.
- The jptspd_debug print JPTSPD_DEBUG_TRACE(JPTSPD_TRACE_SRC) << __JPTSPD_HERE__ << ": No session found " + sid << endl; will cause problems with "+" before sid instead of << with. [PR1016959](#): This issue has been resolved.
- In dynamic subscribers management environment with "maintain-subscriber" feature enabled, when scaling up the logged in subscribers, the demux interface might not be associated with the subscriber and "show auto-configuration extensive" CLI command only print partial output. [PR1017544](#): This issue has been resolved.
- MAC accounting support was added for 40G and 100G interfaces on MPC3 and MPC4 cards. [PR1017595](#): This issue has been resolved.
- Under corner cases, if there are multiple back-to-back Virtual Chassis port (VCP) related CLI commands, Network Processing Card (NPC) core may be observed and FPC hosting the VC ports might reboot. [PR1017901](#): This issue has been resolved.
- In the scenario where router acts as both egress LSP for core network and BRAS for subscribers, RSVP-TE sends PathErr to ingress router due to matching to subscriber interfaces wrongly when checking the explicit route object (ERO), if subscribers are associated with same lo0 address as used by RSVP LSP egress address. [PR1031513](#): This issue has been resolved.
- With an unrecognized or unsupported Control Board (CB), mismatch link speed might be seen between fabric and FPCs, which results in FPCs CRC/destination errors and

fabric planes offline. Second issue is in a race condition, Fabric Manager (FM) might process the stale destination disable event but the error is cleared indeed, it will result in the unnecessary FPC offline and not allowing Fabric Hardening action to trigger and recover. [PR1031561](#): This issue has been resolved.

- In rare cases, the AUTHD daemon may crash and cause a corruption of subscriber dynamic profiles. In-use profiles may be incorrectly marked as not in use. Any subscribers that reference that profile are forced to remain in Terminating state, until the router is rebooted. Daemon restarts and GRES switches are ineffective in working around this situation. [PR1032548](#): This issue has been resolved.
- If a logical interface is used as the qualified-next-hop (which implies the IFL has unnumbered-address configured), and there are changes in the logical interface filter configuration, then the static route might disappear from routing table. To make it reappear, need to delete it from the configuration and add it back. [PR1035598](#): This issue has been resolved.
- When recovering from a split master Virtual Chassis (VC), the line-cards in the new VC-Bm chassis may contain provisioning data out of synchronization versus the master Routing Engine. [PR1036795](#): This issue has been resolved.
- In a subscriber scenario with auto-sensed VLAN configured, after scaled subscribers (in this case, 16K subscribers) login/logout for several times, the subscriber management process might get stuck and not able to restart due to a Session Database (SDB) deadlock issue. [PR1041094](#): This issue has been resolved.
- If the flow routes (flow route is an aggregation of match conditions for IP packets) are active in the kernel, the rpd process might crash after executing command "show route table <x>.inetflow.0 extensive". [PR1047271](#): This issue has been resolved.
- On T Series FPC 1-3 and M320 except E3-FPC with fib-local configuration. If there are multiple FIB local FPCs or the FIB local is a multiple Packet Forwarding Engine FPC, the TCP packets might be out of order, packets re-ordering would occur. It reduces the application level throughput for any protocols running over TCP. [PR1049613](#): This issue has been resolved.

High Availability (HA) and Resiliency

- This issue occurs in rare condition. In the dual Routing Engines scenario, doing interface flap after Routing Engine switchover. If this action is repeated many times, the stale indirect nexthop entry might be seen in kernel, this leads to traffic blackhole. [PR987959](#): This issue has been resolved.

Interfaces and Chassis

- Error message CHASSISD_IPC_DAEMON_WRITE_ERROR is seen in the messages log when there is a Routing Engine mastership change (system reboot, Routing Engine reboot, GRES switchover CLI command), which causes a restart of alarmd, which breaks the IPC connection between alarmd and chassisd. Chassisd does not detect that the IPC connection has been broken, because it is busy processing the mastership change, and then tries to send alarm information to alarmd during this time. So it encounters a write error (broken pipe) and logs the message. [PR908822](#): This issue has been resolved.
- If dynamic VLAN subscriber interface is over a physical interface (IFD), and there are active subscribers over the interface, when deactivate the dynamic VLAN related configuration under the IFD and add the IFD to an aggregated Ethernet (AE) interface which has LACP enabled, the Routing Engine might crash and get rebooted. [PR931028](#): This issue has been resolved.
- Link speed of a LAG bundle may not properly reflect the total bandwidth, when microBFD is enabled on the LAG interface. [PR967046](#): This issue has been resolved.
- In the dynamic-profile environment with preferred-source-address configuration. If subscribers stuck in terminating state, it is impossible to commit changes. [PR978156](#): This issue has been resolved.
- On MX Series router, in rare condition, the kernel might crash and the router will go in db prompt when router reboots. [PR993978](#): This issue has been resolved.
- In L2 circuit, with async notification configured on a client facing interface goes down, then on the remote PE the corresponding CE interface shows up in show interface terse output while in log snmp reports interface down. [PR1001547](#): This issue has been resolved.
- As current Junos OS Multichassis link aggregation groups (MC-LAGs) design, the ARP entry will not sync when learning ARP via ARP request but not Gratuitous ARP/ARP reply, in some specific scenarios (e.g. a host changes its MAC address without sending a Gratuitous ARP), traffic loss might occur. [PR1009591](#): This issue has been resolved.
- VRRP daemon (vrrpd) memory leak might be observed in "show system processes extensive" when VRRP is set with routing-instance and then change any configuration. [PR1022400](#): This issue has been resolved.
- if DPCE 20x 1GE + 2x 10GE X card is present in the chassis, BFD sessions over AE interfaces may not be distributed. [PR1032604](#): This issue has been resolved.

Layer 2 Features

- After configuration change or convergence events, kernel may report `ifl_index_alloc` failures for LSI interfaces and causing KRT queue ENOMEM issue, eventually preventing new IFLs being added to the system. This condition always recovers on its own once convergence is completed. [PR997015](#): This issue has been resolved.
- In the Ethernet ring protection switching (ERPS) environment, once Graceful Routing Engine Switchover (GRES) happens on the ring protection links (RPLs) owner node, there will be a ~30s Ring automatic protection switching (R-APS) message storm in the ring, which in turn cause some VPLS instance flapping. [PR1004066](#): This issue has been resolved.
- If "maintain-subscriber" knob is enabled on the router, DHCPv6 server/relay might be unable to process any packet if deactivate and then activate the routing instance, which means the subscribers can not get the IPv6 addresses. Please note, even with the fix, the results of this scenario are also expected if with "maintain-subscriber" knob enabled. Consider using the workaround to avoid this issue. [PR1018131](#): This issue has been resolved.
- After FPC restart, bridge domain (BD) implicit filters for Ethernet ring protection switching (ERPS) might get reprogrammed with wrong logical interface (ifl) index, which cause ERPS cannot work correctly. [PR1021795](#): This issue has been resolved.
- In a mixed VPLS instance where both LDP and BGP flavors are present with "best-site" knob configured under "site" block, any cli change in that instance will result in rpd crash. [PR1025885](#): This issue has been resolved.
- If a customer is using SNMP and performs an `snmpwalk` on the dhcp binding table, all of the entries might not be displayed. This fix resolves that issue so that bindings for all ip addresses are displayed. [PR1033158](#): This issue has been resolved.

MPLS

- Error "tag_icmp_route:failed to find a chain composite ahead of fwd nh" might be observed when doing traceroute. [PR999034](#): This issue has been resolved.
- When the size of a Routing Engine generated packet going over an MPLS LSP is larger than MTU (i.e. MTU minus its header size) of an underlying interface, and the extra bytes leading to IP-fragmentation are as small as <8 bytes, then that small-fragment will be dropped by kernel and lead to packet drop with kernel message "tag_attach_labels():m_pullup() failed". For example - If SNMP Response with specific size falls into above mentioned condition, then small fragment will be dropped by kernel and eventually the SNMP response will fail. [PR1011548](#): This issue has been resolved.
- TED link information of protocol from highest credibility level is used irrespective of the level at which CSPF is computing. i.e., `cspf-metric` in "show mpls lsp extensive" would have the sum of `te-metric` of IGP with highest credibility at each hop in ERO. This has been corrected and the `cspf-metric` will be sum of `te-metric` of current credibility at each hop. [PR1021593](#): This issue has been resolved.
- When RSVP label-switched-path (LSP) optimize is enabled, RSVP LSP might stay down after a graceful Routing Engine mastership switchover (GRES). To resolve the

problem, the corresponding label-switched-path configuration needs to be deactivated, then, be activated again. [PR1025413](#): This issue has been resolved.

- When configuring point-to-multipoint (P2MP) Label Distribution Protocol (LDP) label-switched paths (LSPs), the labels will never be freed even they are no longer needed. This could lead to the MPLS label exhaustion eventually. To clear the state, the rpd process will restart with core files. [PR1032061](#): This issue has been resolved.

Network Management and Monitoring

- Due to a communication error between the master agent (snmpd process) and the subagent (mib2d process), the device fails to register some MIBs. For example, the following commands do not display any output when you run the command:
user@hostname>show snmp mib walk ifTable user@hostname:~\$ snmpwalk -v 2c -c snmp@exp X.X.X.X ifAlias The following message is displayed: IF-MIB::ifAlias= No Such Object available on this agent at this OID. This means the OID is not registered. [PR978535](#): This issue has been resolved.

Platform and Infrastructure

- If the "Host Loopback" error is triggered after configuration change (e.g. it was seen after just changing the config from XE to AE with scaled IGMP receivers), it will in turn triggers error-reporting infra (CMERROR). Since host-loopback falls under Major error class, default action for Major class is "log". Hence it triggers a TFTP for transferring debug info to Routing Engine. If this tftp write fails due to lower layer issue (i.e. RE-PFE link issue), TFTP would try to send again (max 5 times). Every time lower protocols append headers but when it fails these headers will not get removed. This causes packet corruption and Packet Forwarding Engine would crash while freeing packet. [PR935764](#): This issue has been resolved.
- In a highly congested system (e.g. high multicast traffic rate), traffic/subscribers loss might occur while performing unified In-Service Software Upgrade (ISSU). [PR945516](#): This issue has been resolved.
- When apply-groups are used in the configuration, the expansion of interfaces <*> apply-groups will be done against all interfaces during the configuration validation process, even if the apply-group is configured only under a specific interface stanza. [PR967233](#): This issue has been resolved.
- On MX Series routers with MX Series linecards in a setup involving Packet Forwarding Engine fast reroute (FRR) applications, if an interface is down for more than ARP timeout interval or if ARP entries are cleared by cli commands, then after the interface is up again packet forwarding issues may be seen for traffic being forwarded over that interface. [PR980052](#): This issue has been resolved.
- On MX2020/MX2010 we might see sporadic FO request time-out error reported under heavy system traffic load. This would mean the request returning into a grant took longer then +/-30usec. The packet will still get forwarded through the fabric hence no operational impact. [May 6 18:56:59.174 LOG: Err] MQCHIP(2) FO Request time-out error [May 6 19:33:47.555 LOG: Info] CMTFPC: Fabric request time out pfe 2 plane 6 pg 0, trying recovery. [PR991274](#): This issue has been resolved.

- BFD sessions within default routing-instance are not coming up once inline-services pic is configured and fixed class-of-service forwarding-class is assigned. BFD sessions operating in no-delegate-processing are not affected. [PR999647](#): This issue has been resolved.
- MPLS traffic going through the ingress pre-classifier logic may not determine mpls payload correctly classifying mpls packet into control queue versus non-control queue and expose possible packet re-order. [PR1010604](#): This issue has been resolved.
- When a MIC is inserted into a freshly booted modular MPC3E on MX2020 router, it does not get detected and when an attempt is done to bring it online using CLI command "request chassis mic fpc-slot <> mic-slot <>" it shows the slot is empty. [PR1012004](#): This issue has been resolved.
- The fix was committed for this PR# but it also needs DDOS configuration additional to this fix and it is as below: 1) check the "show ddos-protection protocols statistics terse" 2) For each of the Control plane protocols on the system like ospf/vrrp/pvstp, it is recommended to configure 2X of the rate as give below example along with increasing DDOS rate for virtual-chassis control. Example, ##### set system ddos-protection protocols virtual-chassis control-high bandwidth 20000 set system ddos-protection protocols virtual-chassis control-high burst 20000 set system ddos-protection protocols ospf aggregate bandwidth 1000 set system ddos-protection protocols ospf aggregate burst 1000 set system ddos-protection protocols vrrp aggregate bandwidth 100 set system ddos-protection protocols vrrp aggregate burst 100. [PR1017640](#): This issue has been resolved.
- On the MX2020 platform, the system might fail to replicate multicast packets to the downstream interface located on the FPC slot 12 or above. There is no workaround. [PR1019414](#): This issue has been resolved.
- On MX Series platform with scaled set-up, after deactivate/activate or renaming a bridge domain (BD) which has irb interface associated, the IGMP snooping configured under the BD might not work any more. Please note it happens only when the router is in "network-services enhanced-ip" mode. [PR1024613](#): This issue has been resolved.
- On MX Series based platform, with igmp-snooping enabled and a multicast route with integrated routing and bridging (IRB) as a downstream interface, a multicast composite nexthop is created with a list of L3 and corresponding L2 nexthops. In a rare corner case, the corresponding L2 nexthop to the L3 IRB nexthop is a DISCARD nexthop and will cause the FPC to crash. [PR1026124](#): This issue has been resolved.
- MX Series-based line card might crash when trying to install the composite next-hop used for the next-hop-group configuration related to port mirroring of traffic over IRB to an LSI attached to VPLS instance for a remote host. [PR1029070](#): This issue has been resolved.
- The MX960 Rx Seq # should start at "0", but instead is random and does not reset to zero. Sequence number starts off at what appears to be a "random" value but then it does progress sequentially. The starting value is most likely where the last session ended and the counter was not reset by the control plane. It only occurs in 12.3. [PR1031201](#): This issue has been resolved.

- On MX Series 3D MPC, when there is a congested Packet Forwarding Engine destination, the non-congested Packet Forwarding Engine destinations might experience an unexpected packet drop. [PR1033071](#): This issue has been resolved.
- When the 'enhanced-hash-key services-loadbalancing' feature is used by MX Series based line cards, load balancing of flows across multiple service PICs via the source-address across does not work when internal BGP (IBGP) is used to steer traffic to the inside service-interface. For example the operator will see on the stateful firewall that the same source-address has flows across multiple service interfaces. [PR1034770](#): This issue has been resolved.

Routing Policy and Firewall Filters

- Executing CLI command "show route resolution" and stopping the command output before reaching the end of the database, the rpd process might crash when executing the same command again. [PR1023682](#): This issue has been resolved.

Routing Protocols

- The case occurs in the BGP multipath scenario, in this case, 3 EBGP CEs inject same route to VRF (the origin attribute of these routes are "Incomplete"). One of the routes is selected, based on the order they come in, an older one is selected (not the one with lowest router id). After that, two remote PEs also inject same route which is imported as secondary (the origin attribute of these routes are "IGP"), then the remote PE (secondary) routes become active due to origin attribute ("IGP" is preferred over "Incomplete"). When the remote PE (secondary) routes become inactive, a different EBGP CE path becomes active (lowest router-id). The original EBGP CE path that have the multipath nexthop is never turn down and it still retains the multipath nexthop, whereas the new EBGP CE path does not have any multipath nexthop, that leads to no multipath in dataplane. [PR835436](#): This issue has been resolved.
- When rpf-selection is configured with next-hop specified, if FPC is restarted, in rare condition, the rpd process might crash. [PR915622](#): This issue has been resolved.
- In scaled BGP routes environment, the BGP router has dual Routing Engines, Graceful Routing Engine Switchover (GRES) and Nonstop Active Routing (NSR) is configured, after performing the operation of deactivate/activate BGP groups and commit the configuration, the BGP router might be stuck in "not-advertising" state. [PR961459](#): This issue has been resolved.
- In the multicast environment, when the name of a VRF is changed, if an IGMP interface is trying to be associated with a new multicast instance before the new multicast instance is created, the rpd process might crash. [PR962885](#): This issue has been resolved.
- In the dual Routing Engines scenario with "commit synchronize" configuration. In a corner case, when commit operation is executed, if BGP session flaps before master Routing Engine receives the sync done message from standby Routing Engine, the routing protocol daemon might crash. [PR976184](#): This issue has been resolved.
- In scaled BGP environment, if an NSR enabled router does not have any routing-instance configured, after flapping BGP groups with multiple peers, some BGP neighbors might get stuck in 'not advertising' state. [PR978183](#): This issue has been resolved.

- When all the following conditions are met, if the knob "path-selection always-compare-med" is configured, the rpd process might crash. - routing-instance (VR, VRF) with no BGP configuration - rib-group in default instance with routing-instance.inet.0 as secondary-rib - rib-group applied to BGP in default instance - BGP routes from master tables (inet.0) leaked to the routing-instance table (routing-instance.inet.0) [PR995586](#): This issue has been resolved.
- Abnormal ip6 route-calculation behavior can be seen when ospf3-te-shortcut is configured. [PR1006951](#): This issue has been resolved.
- When the same PIM RP address is learned in multiple VRFs, with NSR configured, rpd on the backup Routing Engine may crash due to memory corruption by the PIM module. [PR1008578](#): This issue has been resolved.
- Under certain circumstances, route validation session might result in rpd process crash. [PR1010216](#): This issue has been resolved.
- When inet.3/inet6.3 is not enabled, BGP group uses inet6.0 table to advertise the routes for both inet6 unicast and inet6 labeled-unicast families. When BGP family is changed, BGP sessions re-establish. When BGP starts to advertise routes to the peer, BGP expects to see route label however if the old inet6 unicast routes are still present (not completely cleaned), then rpd process crashes. The fix is to separate BGP group for inet6 unicast with inet6 labeled-unicast with same rib. The old peers are cleaned up in the old group and new peers are established in new group. Thus, new peer establishment is not delayed by the cleanup of the old peer. [PR1011034](#): This issue has been resolved.
- IsisRouterTable MIB issues, when we do "show snmp mib walk isisRouterHostName/isisRouterTable" we were not getting exact hostname as it is in "show isis hostname" so the actual implementation was not as per RFC-4444, because it was showing only the hostnames of the devices which are immediate neighbors of Dut. Fix: added level info to get sysinfo per each level correctly and filled data(isisRouterTable) correctly. [PR1011208](#): This issue has been resolved.
- Under certain sequence of events RPD can assert after a RPD_RV_SESSIONDOWN event. [PR1013583](#): This issue has been resolved.
- When receiving open message with any capability after the "add-path" capability from BGP peer, the session will be bounced. [PR1016736](#): This issue has been resolved.
- The snmp trap generated when an ipv6 BFD session goes up/down does not contain the ipv6 bfd session address. [PR1018122](#): This issue has been resolved.
- Junos OS implementation of RFC 3107 uses unspecified label (0x000000) when sending route with label withdrawn message. This means Junos OS sends 0x000000 instead of 0x800000 for label withdrawn, which is inconsistent with RFC 3107. [PR1018434](#): This issue has been resolved.
- When BGP add-path feature is enabled on BGP route-reflector (RR) router, and if the RR router has mix of add-path receive-enabled client and add-path receive-disabled

(which is default) client, due to a timing issue, the rpd process on RR might crash when routes update/withdraw. [PR1024813](#): This issue has been resolved.

- When BGP is doing path selection with default behavior, soft-asserts requests are introduced. If BGP route flaps a lot, it needs to do path selection frequently, because of which a great deal soft-asserts might be produced which will cause unnecessary high CPU and some service issues, such as SNMP not responding and even rpd core. [PR1030272](#): This issue has been resolved.

Services Applications

- Softwire tunnel count management is inconsistent and incorrect, thus the output of "show service softwire statistics" might be incorrect. [PR1015365](#): This issue has been resolved.
- With Real Time Streaming Protocol (RTSP) Application Layer Gateway (ALG) enabled, the PIC might crash in case the transport header in status reply from the media server is bigger than 240 bytes. [PR1027977](#): This issue has been resolved.
- In Network Address Translation (NAT) scenario with Endpoint-Independent Mapping (EIM) configured on service PIC, when a new ICMP session is created which matches an existing EIM mapping, the service PIC might crash. [PR1028142](#): This issue has been resolved.
- For T Series or M320 router containing Dynamic Flow Capture (DFC) PIC (either a Monitoring Services III PIC or Multiservices 400 PIC), there are two issues for DFC feature. The first one is the value of "timeout-remaining" for some filters installed on the DFC pic are too huge. The second issue is for some filters, there won't be any flows to which they are attached when forwarding traffic to the content-destination during random DTCP ADDs. [PR1029004](#): This issue has been resolved.
- The cause of the KMD crash is not known. This is not due to SA(Security Associations) memory corruption. The code looks that SA is getting freed without clearing the table entry. [PR1036023](#): This issue has been resolved.

User Interface and Configuration

- CST: chassis core generated while applying group configuration on chassis > FPC [PR936150](#): This issue has been resolved.

VPNs

- For VPLS over VPLS topology, when the VPLS payload has two labels (Customer-VPLS-label and Customer-MPLS-label), the frame might be dropped by the core facing interface hosted on IQ2 PIC with "L2 mismatch timeout" error. This particular scenario is fixed. But there are some other worse scenarios which might hit this issue again due to the system architecture limitation, which are not fixed but need to avoid: * Addition of VLAN tags on Service provider's or CE's VPLS payload e.g. configuring QinQ. * Addition of MPLS tags on Service provider or CE's VPLS payload. * Enabling VPLS payload load balancing on Service provider's PE router. [PR1038103](#): This issue has been resolved.
- In NG MVPN, after ipv6 VRF RP config change, we may hit ipv6 data loss for a short period of time. [PR1049294](#): This issue has been resolved.

Resolved Issues in Release 12.3R8**Class of Service (CoS)**

- On MX Series routers with both MX Series linecards (in this case, MPC and MPCE on the box) and other type linecard (DPCE on the box), when the Default Frame Relay DE Loss Priority Map is configured and committed, all FPCs are getting restarted with core files. [PR990911](#): This issue has been resolved.

Forwarding and Sampling

- DPC crashed after deactivate/activate [routing-instances TPIX bridge-domains IX bridge-options]. [PR983640](#): This issue has been resolved.

General Routing

- RPD crashed on backup Routing Engine when trying to compare gateways of two different types of nexthops, like table next hop which is installed in kernel for one route, router next hop which is selected in backup RPD. [PR828797](#): This issue has been resolved.
- Changing the redundancy mode of rlsq interface from "hot-standby" to "warm-standby" on the fly might lead kernel crash and the router will go in db> prompt. [PR880451](#): This issue has been resolved.
- 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 out. [PR954033](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- On M7i/M10i with enhanced 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](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- Issuing the CLI command "restart packet-triggered-subscribers" might cause sessions to be out-of-sync between the MX SAE and SRC (external policy manager), which results in new subscribers not being able to be created. [PR990788](#): This issue has been resolved.

- 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](#): This issue has been resolved.
- The routing protocol daemon (rpd) might crash continuously with core-dumps upon adding a sub-interface with "disable" configuration to a MC-LAG interface. [PR1014300](#): This issue has been resolved.

Interfaces and Chassis

- 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](#): This issue has been resolved.
- Due to a bug in Packet Forwarding Engine microkernel driver for MX Series MICs, slight variations of the readings from a built-in DC-DC converter may cause ports of the MIC to go down, or even result in MPC crash, with the following message logged: `Ixchip(0): pio_handle(0x4b1ea7d8); pio_read_u32() failed: 1(generic failure)!`
`ix_inq-addr=00200a30` Only MIC-3D-2XGE-XFP (750-028380) cards are affected by this software defect. This issue is resolved in 11.4R10, 12.1R9, 12.2R7, 12.3R5. [PR919618](#): This issue has been resolved.
- Queue stats counters for AE interface will become invalid after deactivating ifl on the AE interface. [PR926617](#): This issue has been resolved.
- If there is an IRB interface configured for "family inet6" in a bridge-domain on an MX Series router, the Packet Forwarding Engine may 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](#): This issue has been resolved.
- Demux Subscriber IFLs might show the interface as 'Hardware-Down' even though the underlying ae bundle and its member link show up. [PR971272](#): This issue has been resolved.
- Temperature Top and Bottom are swapped in show chassis environments output for Type3/Type4 FPCs of T Series. [PR975758](#): This issue has been resolved.
- 1GbE SFP(EX-SFP-1FE-LX) output optical power is restored after reseating by manually removal/insert of SFP although the IF is disabled. [PR984192](#): This issue has been resolved.
- SNMP OID VRRP-MIB::vrrpAssolpAddrRowStatus returns only one Ip address when the interface ifl has been configured with two virtual-addresses under two vrrp-groups. [PR987992](#): This issue has been resolved.
- CFMD may crash after configuration change of an interface in a logical system which is under OAM configuration for a l2vpn instance. [PR991122](#): This issue has been resolved.
- On MX Series router with MX Series linecard 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](#): This issue has been resolved.
- As current Junos OS Multichassis link aggregation groups (MC-LAGs) design, the ARP entry will not sync when learning ARP via ARP request but not Gratuitous ARP/ARP

reply. In some specific scenarios (e.g. a host changes its MAC address without sending a Gratuitous ARP), traffic loss might occur. [PR1009591](#): This issue has been resolved.

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](#): This issue has been resolved.
- In the large-scaled VPLS scenario with LSI (label-switched interface) or VT (virtual tunnel) interface, in race condition, when the LSI or VT interface is deleted, the kernel might crash. [PR990269](#): This issue has been resolved.

Layer 2 Ethernet Services

- 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) packet with multiple management address TLV, memory leak might occur resulting in l2cpd process crash. [PR986716](#): This issue has been resolved.
- When "system no-redirect" is configured, l2 descriptor destination MAC address gets overwritten and causes "DA rejects" on next-hop router. [PR989323](#): This issue has been resolved.
- jnxLacpTimeOut trap may show negative values and incorrect values for jnxLacpifIndex and jnxLacpAggregateifIndex. [PR994725](#): This issue has been resolved.
- In race condition, when FPC gets rebooted or reset, link(s) from this FPC which are part of aggregated 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 aggregated 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](#): This issue has been resolved.

MPLS

- In LDP signaling L2circuit or VPLS scenario, if L2-smart-policy is configured under LDP protocol, and the LDP interface flaps, RPD might crash. [PR899810](#): This issue has been resolved.
- When a First hop LSR is sending Resv Message with non-directly connected IP as nexthop (in Resv HOP object), Junos OS on head end will try to install this in forwarding table. As the nexthop to be used is a non-directly connected address, forwarding table

update will fail with following KRT_Q_STUCK message: RPD_KRT_Q_RETRIES: Route Update: Invalid argument. [PR920427](#): This issue has been resolved.

- 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](#): This issue has been resolved.

Platform and Infrastructure

- With Junos OS 12.3R7, high CPU utilization (100%) will be seen on MPC if inline sampling/Jflow is enabled. The high CPU will affect the MPC's normal tasks. [PR671136](#): This issue has been resolved.
- The Routing Engine and FPCs are connected with an internal Ethernet switch. In some rare case, the FPCs might receive a malformed packet from the Routing Engine (e.g. packet gets corrupted somewhere on its way from Routing Engine to FPC), then the toxic traffic might crash the FPC. [PR938578](#): This issue has been resolved.
- The issue might come when a non-template filter gets deleted (but does not get completely cleaned up) and the same filter index gets reassigned to a template filter. This could be considered as a timing issue given it comes with a very specific sequence of events only. [PR949975](#): This issue has been resolved.
- On MX Series based line card, VPLS traffic might get blocked for about five minutes (timer of MAC address aged-out) after re-negotiating control-word. [PR973222](#): This issue has been resolved.
- If a router has DDoS protection enabled, when DDoS attack happens, the router might not process a DDoS event properly, which results in kernel crash. [PR987193](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- Issues in shared bandwidth policer handling may manifest itself as NPC core, very low call setup rate and various error messages, especially when configuration includes VLAN demux over AE interface. [PR989240](#): This issue has been resolved.
- "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](#): This issue has been resolved.
- Packet dropped with ipv6 reject route is currently subjected to loopback ipv6 filter processing on MX Series-based line cards. As a result, the packet dropped by a reject route might be seen from the "show firewall log". [PR994363](#): This issue has been resolved.
- On MX Series router with MX Series 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 packet loss. [PR999029](#): This issue has been resolved.

- In the IRB interface environment with "destination-class-usage" configuration, if the bridge domain ID is the same as Destination Class Usage (DCU) ID (bridge domain ID and DCU ID are generated by system), the firewall filter might match wrong packets, and the packet forwarding would be affected. [PR999649](#): This issue has been resolved.
- On M7i, or M10i equipped with Enhanced Compact Forwarding Engine Board (CFEB-E). When a MPLS LSP flaps, the CFEB-E is unable to recover 8 bytes of JTREE memory per event. [PR100385](#): This issue has been resolved.
- MS PIC may reset after GRES in case of excessive resolve traffic. [PR1001620](#): This issue has been resolved.
- The non-first IP fragments containing UDP payload may be mistakenly interpreted as PTP packets if the following conditions are met: - the byte at the offset 9 in the IP packet contains 0x11 (decimal 17) - UDP payload - the two bytes at the offset 22 in the IP packet contain the value 0x01 0x3f (decimal 319; byte 22=0x01 and byte 23=0x3f) - PTP protocol The mis-identification of the packet as PTP will trigger the corruption of the fragment payload. [PR1006718](#): This issue has been resolved.
- MPLS traffic going through the ingress pre-classifier logic might not determine mpls payload correctly classifying mpls packet into control queue versus non-control queue, and expose possible packet re-order. [PR1010604](#): This issue has been resolved.

Routing Protocols

- The case occurs in the BGP multipath scenario, in this case, three EBGP CEs inject same route to VRF (the origin attribute of these routes are "Incomplete"). One of the routes is selected, based on the order the router come in, and an older one is selected (not the one with lowest router id). After that, two remote PEs also inject same route which is imported as secondary (the origin attribute of these routes are "IGP"), then the remote PE (secondary) routes become active due to origin attribute ("IGP" is preferred over "Incomplete"). When the remote PE (secondary) routes become inactive, a different EBGP CE path becomes active (lowest router-id). The original EBGP CE path that has the multipath next hop is never turned down, and it still retains the multipath next hop, whereas the new EBGP CE path does not have any multipath next hop, which leads to no multipath in dataplane. [PR835436](#): This issue has been resolved.
- In inter-AS option-C Layer 3 MPLS VPN scenarios, routing protocol process (rpd) might crash when multi hop EBGP is configured between the two autonomous system boundary routers (ASBRs) and static LSPs are configured between the ASBRs to resolve the indirect next hops. [PR869488](#): This issue has been resolved.
- In PIM-SM network with "bootstrap routing" RP selection mechanism used, it is observed that some bootstrap messages (BSMs) generation and forwarding behavior of Junos OS does not conform to RFC standard, specifically in the sections 3.2 (Bootstrap message generation), 3.3 (Sending Candidate-RP-Advertisement Messages) and 3.4 (Creating the RP-Set at the BSR). [PR871678](#): This issue has been resolved.
- In Protocol Independent Multicast (PIM) scenario, if interface gets 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](#): This issue has been resolved.

- In scaled BGP routes environment (global table ~1.5 million routes). First flapping one BGP session (e.g. change the BGP authentication method can get it), after that deleting another BGP session that holds the active routes, might lead to routing protocol daemon (rpd) scheduler slips. [PR928223](#): This issue has been resolved.
- Performing CLI command "clear multicast bandwidth-admission interface <int>" on 64-bit Junos OS results in the rpd to crash. The command should be used without the interface qualifier on the impacted releases. [PR949680](#): This issue has been resolved.
- 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. [PR97501](#): This issue has been resolved.
- On MX Series platforms with IGMP snooping enabled on an IRB interface, some transit TCP packets might be incorrectly considered as IGMP packets, causing packets to be dropped. [PR979671](#): This issue has been resolved.
- There are two receivers joined to same (S,G) and IGMP immediate-leave is configured. When one of the receivers sends the leave message for (S,G), another receiver is not receiving the traffic for 1-2 minutes. [PR979936](#): This issue has been resolved.
- In Inter-AS FEC-129 VPLS scenario with route-reflector (RR) in each AS, if the RR router is also a VPLS PE, it stops reflect VPLS routes received from MP-EBGP peer to its clients. [PR980834](#): This issue has been resolved.
- Ppmd filter is not programmed properly which is resulting Routing Engine to absorb BFD packets instead of Packet Forwarding Engine. [PR985035](#): This issue has been resolved.
- 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 might be seen on the router. [PR986444](#): This issue has been resolved.
- In V4 RG, member site receives traffic from both serving sites for few sources upon withdraw/inject routes for 30 seconds. [PR988561](#): This issue has been resolved.
- In the P2MP environment with OSPF adjacency established, one router's time is set to earlier date than another router. OSPF adjacency might not come up when one router goes down and comes up. [PR991540](#): This issue has been resolved.
- When IS-IS is configuring for traffic engineer (TE), after removing family mpls from the interface and removing the specific interface from [edit protocols RSVP] and [edit protocols mpls] hierarchy levels, corresponding link is not removed from the TED as expected. [PR1003159](#): This issue has been resolved.
- When there are more than 65535 "flow-spec" routes existing in the routing table, the rpd process might crash because it exceeds the current maximum supportable scaling

numbers (Current scaling numbers are in the range of 10K~16K). [PR1004575](#): This issue has been resolved.

- Under certain sequence of events, RPD can assert after an RPD_RV_SESSIONDOWN event. [PR1013583](#): This issue has been resolved.

Services Applications

- 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](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- 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 OS 11.4R7 and all more recent releases without this fix. [PR986974](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- The redundant services interfaces (rsp) configured with "hot-standby" mode might flap upon commit. [PR1000591](#): This issue has been resolved.
- The following messages are being logged at ERR not DEBUG severity: mspd[3618]: mspd: No member config mspd[3618]: mspd: Building package info This PR sets the correct severity. [PR1003640](#): This issue has been resolved.

Subscriber Access Management

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

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](#): This issue has been resolved.

VPNs

- In NG-MVPN scenario with S-PMSI tunnel has been triggered, performing Routing Engine switchover right after deleting the MVPN routing instance might cause rpd process on master Routing Engine to crash. [PR941160](#): This issue has been resolved.
- The S-PMSI tunnel might fail to be originated from ingress PE after flapping the routes to customer multicast source. [PR983410](#): This issue has been resolved.

- In MVPN scenario, a multihomed ingress PE might fail to advertise type-4 after losing routes to local sources. [PR984946](#): This issue has been resolved.
- 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](#): This issue has been resolved.

Resolved Issues in Release 12.3R7

Class of Service (CoS)

- If any of the schedulers have an ID of zero, cosd process might crash following a commit. [PR953523](#): This issue has been resolved.
- Sometimes the cosd process generates the corefile when add/delete child interface on the LAG bundle. [PR961119](#): This issue has been resolved.
- Applying a scheduler with transmit rate below 65,535 bps and rate-limit option fails the commit if the associated interface is a non-existing interface or a virtual interface. [PR964647](#): This issue has been resolved.
- On MX Series router with non-Q DPC (in this case, DPCE 40x 1GE R), when the "interface-set" is configured on a non-Q DPC, then execute the command "show interfaces interface-set queue <interface-set-name>", the DPC might crash. [PR979668](#): This issue has been resolved.

Forwarding and Sampling

- VPLS mac-table doesn't get populated with mac of previous lt interface after replacing the lt interface in the configuration, which might cause CE connected to the lt interface to get isolated. [PR955314](#): This issue has been resolved.
- When port-mirroring or sampling is configured, if a lot of route updates are happening in the system, the routing protocol convergence time might be long and packet loss might be observed. [PR963060](#): This issue has been resolved.
- Under rare circumstances, when a forwarding table filter (FTF) is configured and applied as a routing-instance (i.e. VPLS) in one same configuration change (i.e. in a same commit operation), if the Packet Forwarding Engine receives the forwarding-table binding message before the firewall filter content from the Routing Engine, then the MPC may restart. [PR963584](#): This issue has been resolved.
- In the large scaled DHCP subscribers setup (e.g. 54,000 dual-stack DHCP subscribers), dynamic firewall daemon (dfwd) memory leak occurs during DHCP subscribers login/logout. [PR967328](#): This issue has been resolved.
- DPC crashed after deactivate/activate [routing-instances TPIX bridge-domains IX bridge-options]. [PR983640](#): This issue has been resolved.

General Routing

- MPC might crash with corefiles due to watchdog timer expired. [PR593444](#): This issue has been resolved.
- When gr- interface is disabled, the DECAP-NH also needs to be deleted / set to discard. [PR791277](#): This issue has been resolved.
- The ingress family feature unicast Reverse Path Forwarding (uRPF) check execution order was invalidated when Filter Based Forwarding (FBF) was enabled on MX Series routers with MPCs or MICs. This solution repositions uRPF just prior to FBF, so that both actions are compatible and applicable. This applies to both IPv4 and IPv6. [PR805599](#): This issue has been resolved.
- RPD crashed on backup Routing Engine when trying to compare gateways of two different types of nexthops, like table next hop which is installed in kernel for one route, router next hop which is selected in backup RPD. [PR828797](#): This issue has been resolved.
- When dynamic profile is used in a Dual-Stack network, the IPv4 routes of the subscriber might unexpectedly get deleted after deleting IPv6 routes. Same thing happened when delete IPv4 subscriber routes. [PR864696](#): This issue has been resolved.
- MXVC /kernel: rts_ifstate_client_open: Number of ifstate clients have reached threshold, current = 63 maximum = 63. [PR894974](#): This issue has been resolved.
- The chassisd process is getting stalled and later on restarts with core. The stalling might be caused by debug code which is used to log the errors in the counters when i2c driver or kernel bug shows. [PR912990](#): This issue has been resolved.
- Leak in /mfs/var/sdb/iflstatsDB.db. [PR924761](#): This issue has been resolved.
- On MX Series routers containing multiple Packet Forwarding Engines such as MX240/MX480/MX960/MX2010/MX2020, with either MPC3E or MPC4E cards (MPC3 Type 3 3D/MPC4E 3D 2CGE+8XGE/MPC4E 3D 32XGE), if multicast traffic or Layer 2 flood traffic enters the router via these MPC3E or MPC4E line cards, these line cards may exhibit a lockup, and one or more of their Packet Forwarding Engines corrupt traffic towards the router fabric. [PR931755](#): This issue has been resolved.
- In the MX-VC scenario, have chassis fabric redundancy mode set to increased bandwidth (root@user# set chassis fabric redundancy-mode increased-bandwidth). Then configure the "offline-on-fabric-bandwidth-reduction" for any slot (root@user# set chassis fpc <slot> offline-on-fabric-bandwidth-reduction). After that execute commit, the commit check failed and chassisd crashed with core-dumps. [PR932356](#): This issue has been resolved.
- In Junos OS versions later than 11.2 where IFL localization is enabled, Routing Engine mastership switchover could lead to IFL indexes inconsistency in Ichip FPCs when graceful Routing Engine switchover (GRES) is configured. This inconsistency could gradually lead to IFL index overlaps and traffic blackholing. [PR940122](#): This issue has been resolved.
- When configuring "no-readvertise" flag to existing static route, then this static route will not be exported to other VPN routing and forwarding (VRF) tables from onwards

which is expected, however, because the static route has already been exported to other VRF tables before "no-readvertise" configuration, no deletion event occurs. Also, the "rt-export" bit still set for the static route which is exported to other routing tables after "no-readvertise" configuration. [PR950994](#): This issue has been resolved.

- On MX Series router with MX Series linecard or T4000 router with type5 FPC, and NAT is configured when random size packets are sent, some packets don't translate correctly. [PR951232](#): This issue has been resolved.
- If rpd ACK feature is enabled through command "indirect-next-hop-change-acknowledgements", when a route being added and a quick route change happens on the same route, high routing protocol process (rpd) CPU utilization might be seen and stays high (above 90%) until rpd is restarted. [PR953712](#): This issue has been resolved.
- 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 out. [PR954033](#): This issue has been resolved.
- FPC might lose the socket connection to the Routing Engine during the time kernel live-core dump is active. IGP session might get dropped after the socket connection got closed. The FPC will get restarted by the kernel once the live-core dump has finished. [PR954045](#): This issue has been resolved.
- In subscriber management environment, upgrade Junos OS to specific version (include 12.3R6 13.2R4 13.3R2) via ISSU might make subsequence subscribers fail to connect with following error: "jdhcpd_profile_request: Add Profile dhcp request failed for client in state LOCAL_SERVER_STATE_WAIT_AUTH_REQ: error = 301". [PR959828](#): This issue has been resolved.
- On MX Series router with dynamic vlan scenario, when improper sort order data is sent to dynamic vlan on the Packet Forwarding Engine, the Modular Port Concentrator (MPC) might crash and generate core files. [PR961645](#): This issue has been resolved.
- On all Junos OS platforms, if there are some events that cause Packet Forwarding Engine to restart (e.g. changing system mode on a QFX5100 switch), service might be interrupted because the stale interface index is not deleted successfully. [PR962558](#): This issue has been resolved.
- In the initial router configuration, if static routes are configured over GRE interface and OAM is enabled, then the static routes may remain active while the GRE tunnel is down. [PR966353](#): This issue has been resolved.
- The symptom of the issue: When the PIC reboots, new subscribers are not coming up. [PR967070](#): This issue has been resolved.
- Support for Layer 3 VPN localization has been deprecated in the Junos OS releases and platforms listed below. This affects the following CLI command: "set routing-instances [instance-name] routing-options localize" Junos OS releases: - 12.3R7 (CLI command is hidden) - 13.1R5 (CLI command is hidden) - 13.2R5 (CLI command is hidden) - 13.3R3 (CLI command is removed) - 14.1 (CLI command is removed) - 14.2 (CLI command is removed) Platforms: - M320 Series router - MX Series routers (all) - T Series routers (all). [PR967584](#): This issue has been resolved.

- PPP over ATM transit traffic was not being fragmented correctly by ATM MIC. The changes allow the fragmentation of the transit traffic to work properly. [PR976508](#): This issue has been resolved.
- On T Series router with FIB Localization enabled, if reboot the Routing Engine while scaled traffic running, the FIB-remote FPC might crash. [PR979098](#): This issue has been resolved.
- In the high scale P2MP LSP environment, heap memory leak might occur when the LSP flaps. Then some P2MP LSPs might be not installed, so the traffic will lose. [PR979211](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- In the BFD environment with static route, the BFD session is established between two routers. When disable the subinterface on one router, the BFD AdminDown packet will be sent out from the router (this is not expected). But according to RFC 5882, another router receives the AdminDown packet, and the static route will never be deleted on it. That might cause traffic packets to be dropped. [PR982588](#): This issue has been resolved.
- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX with DPC. In a race condition, the Dense Port Concentrator (DPC) may crash when ifls get added to an ifl-set while that same ifl-set gets 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](#): This issue has been resolved.
- MX Series creates ptsp sessions with large delay. [PR986905](#): This issue has been resolved.
- 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](#): This issue has been resolved.

High Availability (HA) and Resiliency

- In a scenario where backup Routing Engine is resyncing, there is time window between the backup Routing Engine is marked as graceful Routing Engine switchover (GRES) ready and it factually ready. Then graceful Routing Engine switchover (GRES) might fail if it is performed during that time window. [PR958453](#): This issue has been resolved.

Infrastructure

- On RE-S-1800 family of Routing Engines, after an intensive writing to SSD, the immediate rebooting might cause SSD to be corrupted. [PR937774](#): This issue has been resolved.

Interfaces and Chassis

- When using vrrp inheritance, the vrrp stat could stick in bring-up state if it uses wrong vrrp parent, which will affect all new added vrrp even with a correct parent. As immediate recovery, it is suggested to correct or remove these two interfaces: xe-1/0/0.67 and xe-1/0/0.1066 and restart the vrrp process with the command: restart vrrp. [PR820298](#): This issue has been resolved.
- If the "tunnel-destination" address of a Generic Routing Encapsulation (GRE) interface is placed in one instance and the GRE interface is placed in another routing-instance, the lookup for the GRE tunnel destination is done on inet.0 instead of the appropriate routing instance's inet.0 table. The similar issue could happen on IP-over-IP or Automatic Multicast Tunneling (AMT) tunnels too. [PR851165](#): This issue has been resolved.
- In scenario when CCM has been running for a while and user issues the following CLI command: "show oam ethernet connectivity-fault-management interfaces interface-name extensive", the initial value reported for CCMs sent is wrong and then the command is executed immediately again and the value is correct. [PR880615](#): This issue has been resolved.
- Traffic which uses MPLS next-hops enters bridge-domain via IRB interface and if forwarding next-hop moves from non-aggregate interface to aggregate interface (MAC move), the MPLS next-hops are not correctly programmed in the Packet Forwarding Engine and dropped. The child next-hop of the aggregate interfaces are missing. Once IRB MPLS next-hop moves from aggregate interface to non-aggregate interfaces are not affected. IPV4 traffic will not trigger traffic drop upon mac move. The second symptom is a possible kernel core-dump on the new backup Routing-Engine after mastership switch. This applies to an IRB mac move for ipv4,ipv6 and mpls next-hops. [PR924015](#): This issue has been resolved.
- Queue stats counters for AE interface will become invalid after deactivating ifl on the AE interface. [PR926617](#): This issue has been resolved.
- "Too many I2C Failures" alarm happens when a FRU (in this case: PWR-MX960-4100-AC-S) experienced six consecutive i2c read/write failure. While the PEM still providing power to the chassis, chassisd daemon cannot read/write information from the PEM until it is reseated. In recent investigation, engineering team has come up some enhancements for this MX960 HC-AC and HC-DC PEM: 1. PEM i2c bus hang avoidance 2. Junos OS recovery from a hung i2c bus 3. noise reduction This

Junos OS eliminates the need for the PEM FW upgrade, and at the same time is 100% compatible with those PEMs which have been upgraded. [PR928861](#): This issue has been resolved.

- Digital Optical Monitoring MIB jnxDomCurrentRxlaserPower gives wrong value in 12.3R3-S6. [PR946758](#): This issue has been resolved.
- When Connectivity Fault Management (CFM) is configured, if maintenance domain intermediate point (MIP) session associated to default maintenance domain (MD) is inactive, a deletion of interface can not delete MIP session structure, hence might causing memory leak. This crash could also be seen if delete more than one Virtual private LAN service (VPLS) routing instance with no neighbor configuration. [PR947499](#): This issue has been resolved.
- On MX Series routers (platforms) with Small Form-factor Pluggable (SFP). There are two connected ports. Both of them are configured with auto-negotiation, and one port is configured speed with 10M. So both the ports auto negotiate to 10m. After removing the SFP from one port and reinserting it, then the port is not coming up. [PR953518](#): This issue has been resolved.
- Kernel crash may happen when a router running a junos install with the fix to PR 937774 is rebooted. This problem will not be observed during the upgrade to this junos install. It occurs late enough in the shutdown procedure that it shouldn't interfere with normal operation. [PR956691](#): This issue has been resolved.
- If there is an IRB interface configured for "family inet6" in a bridge-domain on an MXSeries router, the Packet Forwarding Engine may 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](#): This issue has been resolved.
- In the large scaled VPLS environment , during delete routing-instance of type VPLS, the memory is not getting freed. The connectivity-fault management daemon (cfmd) might crash with a core file generated.,The core files could be seen by executing CLI command "show system core-dumps". [PR975858](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- After the following process, we can find MCAE becomes standby/standby status. Even if we set "set interfaces aeX aggregated-ether-options mc-ae events iccp-peer-down prefer-status-control-active" for both routers, we can find this issue. << topology example >> iccp ge-1/0/1 ge-1/0/1 [MX80(router A)]-----[MX240(router B)] \ ae0 ae0 / --active-- \ / --standby-- \ MC-LAG / \ \ \ / ae0(ge-0/0/0)\ /ae0(ge-0/0/1) [EX4200(switch C)] << process >> initial status router A : active router B : standby 1. disable ae0 of router A. 2. disable iccp link of router A. 3. disable ae0 of switch C 4. enable iccp link of router A. (Please wait until iccp status up.) 5. enable ae0 of switch C 6. enable ae0 of router A. [PR982713](#): This issue has been resolved.

Layer 2 Ethernet Services

- When DHCP local server and DHCP relay are both configured on same router, the DHCP relay binding might get lost if a graceful Routing Engine switchover (GRES) is performed. [PR940111](#): This issue has been resolved.
- IP address change of a DHCP relay interface does not get reflected in gateway IP address (giaddr) when maintain-subscribers knob is enabled, which needs to restart DHCP daemon to make it work again. [PR951909](#): This issue has been resolved.
- In L3 Wholesale environment, the DHCP clients might fail to renew their address in DHCP relay scenario. [PR956675](#): This issue has been resolved.
- In the MPLS environment, LDP session protection is configured and LDP session establishes via IRB interface. When disable IRB interface, LDP session protection does not work. The LDP session goes down even if there is an alternate route. [PR959396](#): This issue has been resolved.
- Configuring Ethernet Ring Protection Switching (ERPS), after changing interface's MTU on Ring Protection Link (RPL) owner, all the interfaces on RPL owner change into forwarding state, hence cause a layer 2 loop. [PR964727](#): This issue has been resolved.
- On MX Series platform with Ethernet Ring Protection Switching (ERPS) configuration, after disabling Ring Protection Link (RPL) interface and then moving RPL from west interface to east interface, as a result, the ERPS east and west interface might go into discard state at same time. [PR970121](#): This issue has been resolved.
- When Cisco running in an old version of PVST+, it doesn't carry VLAN ID in the end of BPDU. So Juniper Networks equipment fails to respond Topology Change Notification ACK packet when interoperates with Cisco equipment. After the fix, Juniper Networks equipment will read the VLAN ID information from Ethernet header. [PR984563](#): This issue has been resolved.

MPLS

- When a First hop LSR is sending Resv Message with non-directly connected IP as nexthop (in Resv HOP object), Junos OS on head end will try to install this in forwarding table. As the nexthop to be used is a non-directly connected address, forwarding table update will fail with following KRT_Q_STUCK message: RPD_KRT_Q_RETRIES: Route Update: Invalid argument. [PR920427](#): This issue has been resolved.
- When Packet Forwarding Engine fast reroute (FRR) applications are in use (such as MPLS facility backup, fast-reroute or loop free alternates), a primary path interface flap could be triggered due to Operation, Administration, and Maintenance (OAM) link failure detection or by Bidirectional forwarding detection (BFD). However, this interface flap might lead to a permanent use of the backup path, which means the original primary path could not be active again. [PR955231](#): This issue has been resolved.
- We add timer for all aggregate LDP prefixes but are not deleting it when the timer expires because of a bug. Since the timer is not expiring, we never update the route for any change. This will be sitting in the routing table as a stale entry. [PR956661](#): This issue has been resolved.

- When the Label Distribution Protocol (LDP) feature is enabled and the background job "LDP sync send filtered label job" is running, when shut down the LDP, due to LDP failing to delete a job that didn't exist while shutting down, routing protocol process (rpd) might crash. [PR968825](#): This issue has been resolved.
- In the large scaled MPLS setup with NSR enabled, when restart routing protocol process (rpd) on standby Routing Engine, or reload standby Routing Engine, or reload router, some filtered output label bindings might be missed on the backup Routing Engine, which leads to Label Distribution Protocol (LDP) database between the master and backup Routing Engines to be inconsistent. [PR970816](#): This issue has been resolved.
- In a scaled MPLS environment, whenever fast reroute (FRR) or Link Protection (LP) or Node Protection (NP) is configured, the switchover from the primary LSP to the secondary LSP might cause traffic loss for a few seconds. [PR973070](#): This issue has been resolved.
- In the MPLS environment, when execute the command "show snmp mib walk mplsXCTable" to walk the MPLS cross connect table, the routing protocol daemon (rpd) CPU utilization might reach over 90%, and the rpd process doesn't respond to any CLI show commands. [PR978381](#): This issue has been resolved.
- snmpwalk/snmpgetnext or "show snmp mib walk" fail when polling MPLSLSP OCTETS, MPLSLSP PACKETS, MPLSLSP INFO OCTETS or MPLSLSP INFO PACKETS. [PR981061](#): This issue has been resolved.
- 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](#): This issue has been resolved.

Platform and Infrastructure

- NPC core observed @nh_dfw_get_correct_next_intf_prefix. [PR801607](#): This issue has been resolved.
- Unwanted continuous vc_ifd_is_remote(): is_remote_ifd ptr NULL! syslog messages being printed by the MS-DPC. [PR868273](#): This issue has been resolved.
- When the instance have vlan-id all and adding interface unit with "vlan-tags outer X inner Y" to this instance, traffic from ALL instance VLANs is leaking over that unit tagged with outer tag X and each VLANs own inner tag A,B,C. [PR883760](#): This issue has been resolved.
- On MX Series based line card, for interfaces tagged with VLAN ID same as the native-vlan-id configured on the interface, FPC adds Native VLAN ID to the packets received on the interface and destined to the host. This is irrespective of the packet content. This results in the packets getting doubly tagged when receiving packets which are already tagged with VLAN ID matching the Native VLAN ID, and thus cause ARP resolution failure on Native VLAN. For example, the ARP packets to IRB (on VLAN 101) are tagged with VLAN ID 101 (which is also the native VLAN ID) and are getting additionally tagged. Hence they are dropped by the IRB and this can cause the ARP

request packet to not get resolved on Native VLAN. [PR917576](#): This issue has been resolved.

- When the transit traffic is hitting the router and the destination is a local segment IP which requires ARP resolution, it's mis-classified by the DDOS filter and an incorrect policer is applied. This leads to host queue congestion. [PR924807](#): This issue has been resolved.
- MPC Type 2 3D may crash with CPU hog due to excessive link flaps causing the interrupts to go high . [PR938956](#): This issue has been resolved.
- On I-chip platforms, when forwarding table filter (FTF) is configured for a virtual private LAN service (VPLS) routing instance, the jtree memory corruption might occur if the routing table attached by FTF is destroyed. The route table that is attached by FTF can get destroyed with different events such as interface which is part of the VPLS routing instance flaps or route-distinguisher is changed, etc. [PR945669](#): This issue has been resolved.
- If a PE router is both egress and transit node for a p2mp lsp, the Packet Forwarding Engine may report errors and install a discard state for the fib entry representing the p2mp lsp label with bottom of stack bit set to 0 . This problem does not have any impact since there is no application using the s=0 entry of a p2mp lsp. [PR950575](#): This issue has been resolved.
- MIC-3D-40GE-TX (3D 40x 1GE(LAN) RJ45) restarts with core-dumps repeatedly after configuring "VRRP interface" and "traffic-manager mode ingress-and-egress" on PIC2 or PIC3. [PR950806](#): This issue has been resolved.
- Current display of "cli> request chassis routing-engine hard-disk-test show-status" command for Unigen SSD identified by "UGB94BPHxxxxxx-KCI" is incorrect and can be misleading when use for troubleshooting. For example, attribute 199 is display as "UDMA CRC Error Count" is actually "Total Count of Write Sector". [PR951277](#): This issue has been resolved.
- For Trio-based platforms, traffic unbalance (60:40 or 70:30) can be seen in the output interface of 2nd node in the cascaded topology. Current Junos OS hash-seed implementation on Trio-chipset can be used to protect the hash-cascade problem (unbalance at 2nd node output, 0:100 for example) but it doesn't work very well (60:40 or 70:30 can be seen). This fix made enhancement, it can deliver nearly 50:50 LB performance. [PR953243](#): This issue has been resolved.
- On Trio based platforms, the counter for Network Time Protocol (NTP) protocol of the output of "show ddos-protection protocols ntp" would be always null, even though it is confirmed that there are NTP unicast packets received. The reason for this is that only multicast NTP packets are treated as NTP packet type by DDOS protection policer in current implementation, whereas the unicast NTP packets are not. [PR954862](#): This issue has been resolved.
- * MX2020 FanTray power specification. - zone#1:FT#3 - gets power from zone#1 only - zone#1:FT#2 - gets power from zone#0 in case of no-power in zone#1 - zone#0:FT#1 - gets power from zone#0 only - zone#0:FT#0 - gets power from zone#1 in case of no-power in zone#0 - Critical(Minimum) number for MX2020 operation is 3 If one of zone has no PSM, then it means FAN single-fault in the chassis's point of view. For

example, if zone#1 has no PSM, then the FT#3 doesn't get power as it is local-powered FT. Hence, in this case, the FT#3-LED should show ORANGE to notify the single-fault to user, while FT#2 can show GREEN if it gets enough power from zone#0. In addition, CRAFT-LED for FT#3 should be turned off. * Due to HW-limit(bicolor), it could not show ORANGE color. In current implementation, both CRAFT-LED, FT#3-LED show GREEN. That's problem. * NOTE: Junos OS doesn't support FT double-fault scenario. (MX2020 needs minimum 3 FTs.) If FT#2 gets in trouble in above case(i.e., FT double-fault), the user should see serious cooling-trouble on SFMs within 1 minute.

[PR957395](#): This issue has been resolved.

- When the special character "%" in comment field, execute the command "root@user# show | display xml", management daemon (mgd) might crash. [PR957651](#): This issue has been resolved.
- Unable to modify dynamic configuration database after first commit. [PR959450](#): This issue has been resolved.
- When we set "traffic-manager mode ingress-and-egress" on "MIC-3D-40GE-TX (3D 40x 1GE(LAN) RJ45)", we cannot use ingress queue correctly on PIC2 and PIC3. *Note: We cannot see this issue if we set the above configuration to PICO or PIC1. [PR959915](#): This issue has been resolved.
- In current Junos OS, a PSM shows dc output value even though it is turned off by switch. This cosmetic bug causes miscalculation of actual usage in 'show chassis power'. [PR960865](#): This issue has been resolved.
- Upon the deletion of a routing-instance and subsequent commit, error logs are generated from each Type 1 - 3(non E3) based FPC. These logs are cosmetic and can be ignored. [PR964326](#): This issue has been resolved.
- A defect in L3VPN Make Before Break code was resulting in freeing memory corresponding to old nexthops which is being used by egress Packet Forwarding Engine. This was resulting in memory corruption. [PR971821](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- SNMP alarms/traps could be generated for unpowered fan trays when only one zone is powered. [PR982970](#): This issue has been resolved.

Routing Protocols

- On MX Series routers containing multiple Packet Forwarding Engines such as MX240/MX480/MX960/MX2010/MX2020 routers, with DPC (Dense Port Concentrator) or FPC (Flexible Port Concentrator) or with line cards designated with "3D", RPD may restart when attempting to send a PIM assert message on an interface (whose interface index exceeds 65536). It is likely that RPD restarts repeatedly, since after RPD has restarted and protocols have converged, the same PIM assert will trigger further RPD restarts. [PR879981](#): This issue has been resolved.
- The rpd process might crash when executing the command "show route advertising-protocol bgp <nbr>" without a table option, or with a table that is not advertised by BGP. [PR959535](#): This issue has been resolved.
- With BGP import policy as next-hop peer-address, if the local router receives inet (or inet-vpn) flow network-layer reachability information (NLRI), routing protocol process (rpd) might crash. Junos OS is designed to create a fictitious nexthop for inet flow and inet-vpn flow families as they don't send/expect-to-receive nexthops. So in this case when the import-policy set a non-null next-hop for the received inet (or inet-vpn) flow route, it could not handle properly which might result in rpd crash. [PR966130](#): This issue has been resolved.
- In the dual Routing Engine scenario, after a Routing Engine switchover, the periodic packet management daemon (ppmd) might exit. [PR979541](#): This issue has been resolved.
- Due to some corner cases, certain commits could cause the input and/or output BGP policies to be reexamined causing an increase in rpd CPU utilization. [PR979971](#): This issue has been resolved.
- Forwarding cache limit not properly meeting threshold after configuration change when configured per address family. [PR980578](#): This issue has been resolved.

Services Applications

- SIP call forwarding may fail when NAT is used between parties even though the SIP ALG is in use. [PR839629](#): This issue has been resolved.
- 11.4 Junos OS releases introduce the IKEv2 support and a stricter check on IKE/IPsec SAs proposal parameters. [PR843893](#): This issue has been resolved.
- Any SIP MESSAGE request will be dropped by the SIP ALG. This type of request is unsupported from day one. This is rare type of request that will not prevent more usual SIP operations such as voice calls, but it may affect some instant messaging applications based on SIP. [PR881813](#): This issue has been resolved.
- 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. [PR939832](#): This issue has been resolved.
- DNS multiple queries A and AAAA might cause the Service-PIC to restart. [PR943425](#): This issue has been resolved.

- When sending traffic from Internet end, the softwire count is incorrect. [PR948583](#): This issue has been resolved.
- The dynamic flow control process (dfcd) might core dump when Dynamic Tasking Control Protocol (DTCP) trigger request is same for both the VLAN and DHCP subscriber. [PR962810](#): This issue has been resolved.
- In the context of DS-Lite softwire scenario, where the Address Family Transition Router (AFTR) node performs NAT with Endpoint Independent Filtering (EIF) and Endpoint Independent Mapping (EIM) enabled, the simultaneous arrival of two packets from opposite sides of the NAT will trigger the creation of the same flow, which in a race condition results in the Service-PIC restart. [PR966255](#): This issue has been resolved.
- JI2tpd could crash because of loss of private sdb table delete operations when the mirroring connection bounces on the standby. [PR968947](#): This issue has been resolved.
- When transferring large FTP file, the server might send packets with incorrect Layer 4 checksum. If inline NAT service is enabled on the router, it might transit the packets to client instead of dropping it, which eventually causes the client FTP time out. [PR972402](#): This issue has been resolved.
- Due to a race condition, the Service PIC handling the MX Series DS-LITE tunnels may core during a softwire lookup. [PR978598](#): This issue has been resolved.

Software Installation and Upgrade

- Routing Engine could be brought to DB mode when rebooting after interrupted downgrade. [PR966462](#): This issue has been resolved.

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](#): This issue has been resolved.

VPNs

- In Multicast VPN scenario, when multiple instances exist and tunnel-source is configured in more than one instance, the routing protocol process (rpd) might reinitialize with core file if one of the instances is deleted or Route Distinguisher (RD) ID is changed. [PR877682](#): This issue has been resolved.
- With these high amount of streams, we have higher number of data-mdt-tlvs to process which is becoming a bottleneck. [PR957280](#): This issue has been resolved.
- 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](#): This issue has been resolved.
- Route group member site and regular site may receive data from two serving sites of two groups for the same (S,G). This only happens when in one RG there are no receivers. [PR974245](#): This issue has been resolved.

- In Rosen MVPN environment, if there a two multihomed ingress PEs, when the route to multicast source flaps, the receiver router might keep switching between sender Data MDTs, which result in traffic loss. [PR974914](#): This issue has been resolved.
- In the Rosen MVPN environment, setting the TOS IP control packet bit can avoid the possibility of data-mdt TLV messages being dropped in the core during congestion. But in this case, the TOS field to indicate its IP control packet (0xc0) is not set. This might lead to traffic loss. [PR981523](#): This issue has been resolved.

Resolved Issues in Release 12.3R6

Class of Service (CoS)

- After swapping MPC2E-3D-Q card with MPC2E-3D-EQ card, an interface is still running out of queues with only 32k queues in use. [PR940099](#): This issue has been resolved.

Forwarding and Sampling

- Filter state failed to be present in the kernel and was not created on Packet Forwarding Engine. Added check to retry creating filter state before pushing to Packet Forwarding Engine. [PR937607](#): This issue has been resolved.

General Routing

- When graceful Routing Engine switchover (GRES) and ARP purging is enabled, frequent route flapping, route entry and nexthop fail to sync up between the master Routing Engine and the backup Routing Engine. So when the master Routing Engine would like to add a new nexthop but see the backup Routing Engine has already found a nexthop with same destination. It makes the backup Routing Engine reboot and crash on both the Routing Engines. [PR899468](#): This issue has been resolved.
- RPD on backup Routing Engine might hit out of memory condition and crash if BGP protocol experiences many flaps. [PR904721](#): This issue has been resolved.
- After FPC/MPC is reset or while PPPoA customer login, in rare case, the ppp daemon (jpppd) might get an incorrect value from device control daemon (dcd) which might cause all the new Link Control Protocol (LCP) messages to be ignored and results in static PPPoA sessions can not come up. This problem is seen on MX Series platform products so far, but the problem is mostly common and if other products are using the same version of Junos OS software it might apply to them. [PR912496](#): This issue has been resolved.
- Leak in /mfs/var/sdb/iflstatsDB.db [PR924761](#): This issue has been resolved.
- MX80 routers now support CLI command "show system resource-monitor summary". [PR925794](#): This issue has been resolved.
- High routing protocol process (rpd) CPU utilization is seen and it stays high (above 90%) until the rpd is restarted. [PR925813](#): This issue has been resolved.
- If MX Series router is in increased-bandwidth fabric mode, pulling out one SCB might cause packets loss. [PR934544](#): This issue has been resolved.
- tcp_inpcb buffer leak in ADC and TLB service pics. [PR934768](#): This issue has been resolved.

- When an SNMP walk is performed to query the native VLAN (mib-2.17.1.4.5.1...: dot1qPvid) or the logical type (trunk or access) of the interface (mib-2.17.1.4.3.1.5...: dot1qPortVlan), the SNMP walk might cause a memory leak on the Layer 2 address learning process (l2ald), and the process might crash with a core file generated. [PR935981](#): This issue has been resolved.
- If IPv6 duplicate address is detected, interface can't recover to normal state after flapping interface. Reconfigure IPv6 address will resolve this issue. [PR936455](#): This issue has been resolved.
- Master Routing Engine reboot due to "panic: pfe_free_peer: not in peer proxy process context" Trigger: replacement of backup Routing Engine. [PR936978](#): This issue has been resolved.
- LNS drops the LCP Compression Control Protocol (CCP) packet silently comes from L2TP tunnel. [PR940784](#): This issue has been resolved.
- In subscriber management environment, profile database files at backup Routing Engine get corrupted when the dynamic profile versioning and commit fast-synchronize are enabled in configuration. After graceful Routing Engine switchover (GRES) when the backup Routing Engine become master, all the existing DHCP subscribers stuck in RELEASE State and new DHCP subscribers can't bind at this point. [PR941780](#): This issue has been resolved.
- MP-BGP route withdraw update might not been sent after deletion of a routing-instance configured with resolve import policy. [PR942395](#): This issue has been resolved.
- Egress multicast statistics displays incorrectly after flapping of ae member links on M320 or T Series FPC (M320 non-E3 FPC and T Series non-ES FPC). [PR946760](#): This issue has been resolved.
- With scaled configuration of ATM VCs (~4000 VCs) on a single MIC-3D-8OC3-2OC12-ATM ATM MIC, the MIC might crash. The crash is not seen with lower scale (i.e. less than 3500 VCs per MIC). [PR947434](#): This issue has been resolved.
- CLI command "show interfaces queue" does not account for interface queue drops due to Head drops. This resulted in the "Queued" packets/bytes counter to be lesser than that was actually received and dropped on that interface queue. This PR fixes this issue. Head-drops, being a type of RED mechanism, is now accounted under the "RED-dropped" section of the CLI command "show interfaces queue". [PR951235](#): This issue has been resolved.

High Availability (HA) and Resiliency

- With minimal flow configuration, if graceful Routing Engine switchover is not enabled, routing protocol process (rpd) crashes during shutting down the rpd process due to missing safety checks. The core files could be seen by executing CLI command "show system core-dumps". [PR852766](#): This issue has been resolved.

Interfaces and Chassis

- If there are several logical systems in one router, basically one logical tunnel (lt-) interface needs to work with another lt- interface, which is peer lt- interface. If one of them allocates a MAC address first and the other attempts to allocate a MAC address, then panic happens since it is a reallocation which finally results in the kernel crash. The problem might be seen when deactivating and then activating logical systems or renaming the lt- interface. [PR837898](#): This issue has been resolved.
- The eeprom SFP-Type descriptor has been updated to display different unique values for fixed-rate or tri-rate copper SFPs. Going forward, the model SFP-1GE-T shows as "1000BASE-T Copper SFP" while model SFP-1GE-FE-E-T shows as "Tri Rate Copper SFP". [PR877152](#): This issue has been resolved.
- In scenario when CCM has been running for a while and user issues the following CLI command: "show oam ethernet connectivity-fault-management interfaces interface-name extensive", the initial value reported for CCMs sent is wrong and then the command is executed immediately again the value is correct. [PR880615](#): This issue has been resolved.
- Problem scenario: CFM UP MEP for Bridge/VPLS is configured on MPC with action profile as 'interface down' Problem statement: When the CFM sessions go down due to network outage at the core, action profile is triggered and configured interface is brought down. When the Core network failure is corrected, CFM will not automatically recover because interface will continue to remain down. [PR884323](#): This issue has been resolved.
- When MX Series routers are running with MC-LAG in active-active mode, the layer 2 address learning daemon (l2ald) might crash if a MAC address is being deleted from one port while the same entry is locally learned on a different port. [PR888636](#): This issue has been resolved.
- In Point-to-Point Protocol over Ethernet (PPPoE) scenario, if some PPPoE session was added and deleted, after performing Routing Engine switchover operation, the Broadband Remote Access Server (BRAS) might fail to allocate PPPoE session IDs on interFace Descriptor (ifd). [PR896946](#): This issue has been resolved.
- In Multichassis Link Aggregation (MC-LAG) scenario, when MC-LAG works on Active-Active mode, if the link of MC-LAG flaps repeatedly, the layer 2 address learning daemon (l2ald) might crash with a core file generated. [PR913222](#): This issue has been resolved.
- Problem Statement: OAM Packets do not gets forwarded with UP and Down MEP configured in access and core interfaces of the bridge down respectively along with MIP configured on the BD. The above configuration was resulting in not honor split

horizon forming a loop in core network. This results in packet drop in core network.

[PR925288](#): This issue has been resolved.

- "Too many I2C Failures" alarm happens when a FRU (in this case: PWR-MX960-4100-AC-S) experienced 6 consecutive i2c read/write failures. While the PEM still providing power to the chassis, chassisd daemon cannot read/write information from the PEM until it is reseated. In recent investigation, engineering team has come up some enhancements for this MX960 HC AC PEM: 1. PEM i2c bus hang avoidance 2. Junos OS recovery from a hung i2c bus 3. noise reduction This Junos OS eliminates the need for the PEM FW upgrade, and at the same time is 100% compatible with those PEMs which have been upgraded. [PR928861](#): This issue has been resolved.
- In PPPoE subscriber management environment, when PPP daemon is receiving an LCP packet with an invalid code ID and without any option, jpppd process crashes with a core file dumped. [PR929270](#): This issue has been resolved.
- After APS switchover, duplicate packets might be received from the backup circuit under SONET APS configuration with channelized enhanced intelligent queuing (IQE) interface. [PR930535](#): This issue has been resolved.
- This is a day-1 issue. When a member link was added to or removed from an aggregate bundle like AE on a dual Routing Engine system without graceful Routing Engine switchover (GRES), Kernel in the backup Routing Engine would crash due to assertion failure in the function `rt_pfe_nh_cont_nh_decrement_ack_count`. [PR935729](#): This issue has been resolved.
- Traffic is not flowing over Demux input interface. [PR937035](#): This issue has been resolved.
- PCS statistics counter(Bit errors/Errored blocks) not working on Mammoth PIC(xge). [PR942719](#): This issue has been resolved.

Layer 2 Features

- ===== BACKGROUND ===== A global graceful Routing Engine switchover (GRES), which will cause a master Routing Engine to transition to backup, WILL require all Kernel state to be cleaned so that it can start a fresh resync from the new master. Ksyncd is tasked with cleaning up Kernel state. On cleaning routing tables, if any table has a non-zero reference count, it will return "Device Busy" to the ksyncd. Ksyncd will try 5 successive cleanup attempts after which it will trigger a live Kernel core. ===== PROBLEM ===== In ksyncd's kernel cleanup, the Bridge Domain mapped to a VPLS routing table is deleted AFTER an attempt is made to delete the route table. This is a catch-22 since BDs hold reference counts to the routing table. ===== FIX ===== Cleanup of VPLS routing tables should proceed bottom up in the following order: NextHop Deletes, User Route Deletes, Interface Deletes(ifd,ifl,iff), STP Deletes, Bridge Domain Deletes, Mesh Group Deletes and finally Routing Table delete. This ensures that when we get to routing table delete, all dependencies, that could hold a ref cnt to the routing table, are now gone. [PR927214](#): This issue has been resolved.

Layer 2 Ethernet Services

- In multilink scenario, while polling the multilink statistics, the Packet Forwarding Engine statistics thread might be yielded. This might happen where there are large number of bundles and links. When the statistics thread is yielded, the context switches to the Packet Forwarding Engine manager thread to handle link and bundle delete operations, some pointers used by the statistics thread are freed up and so when the statistics thread regains control it crashes because of the dangling pointers. [PR827326](#): This issue has been resolved.
- In MX Virtual Chassis (MXVC) scenario, under high scale system environment (many Aggregated Ethernet interfaces, many logical interfaces), after performing global graceful Routing Engine switchover by CLI command "request virtual-chassis routing-engine master switch", the Link Aggregation Control Protocol (LACP) state of access Link Aggregation Group (LAG) interface might change and therefore resulting in traffic loss. [PR885013](#): This issue has been resolved.
- In Ethernet ring protection scenario, upon FPC reboots the STP index will get mis-aligned causing traffic drop. when this issue occurs following message can be seen. Before FPC restarts: user@router> show protection-group ethernet-ring vlan Ethernet ring IFBD parameters for protection group Ring1 Interface Vlan STP Index Bridge Domain xe-5/3/0 302 222 default-switch/v302 xe-0/2/0 302 223 default-switch/v302 xe-5/3/0 308 222 default-switch/v308 xe-0/2/0 308 223 default-switch/v308 After FPC restarts: user@router> show protection-group ethernet-ring vlan Ethernet ring IFBD parameters for protection group Ring1 Interface Vlan STP Index Bridge Domain xe-5/3/0 302 245 <<<< default-switch/v302 xe-0/2/0 302 223 default-switch/v302 xe-5/3/0 308 222 <<<< default-switch/v308 xe-0/2/0 308 223 default-switch/v308 [PR937318](#): This issue has been resolved.
- Service accounting interim updates not being sent. [PR940179](#): This issue has been resolved.

MPLS

- When static LSPs are configured on a node, RPD could assert upon committing a MPLS-related configuration change. Example: `router> show system rollback compare 9 8 [edit protocols mpls] interface ae11.0 { ... } + interface as3.0 { + admin-group red; + } [edit protocols ISIS interface as3.0 level 2] ! inactive: metric 2610`; The following error is seen in `/var/log/messages` in-relation to a static lsp, immediately following the above-mentioned configuration change: `rpd[1583]: UI_CONFIGURATION_ERROR: Process: rpd, path: [edit groups STATELESS_ARIADNE protocols mpls static-label-switched-path static-lsp], statement: transit 1033465, static-lsp: incoming-label 1033465 has already been configured by this or other static applications` [PR930058](#): This issue has been resolved.
- In certain circumstance, the Junos OS rpd route flash job and LDP connection job are always running starving other work such as stale route deletion. These jobs are running as LDP is continuously sending label map and label withdraw messages for some of the prefixes under ldp egress policy. This is due to LDP processing a BGP route from inet.3 for which it has a ingress tunnel (the same prefix is also learned via IGP) creating a circular dependency as BGP routes can themselves be resolved over a LDP route. [PR945234](#): This issue has been resolved.
- In a highly scaled configuration the reroute of transit RSVP LSPs can result in BGP flap due to lack of keepalive messages being generated by the Routing Engine. [PR946030](#): This issue has been resolved.
- On IS-IS interfaces configured with point-to-point and ldp-synchronization, after a change of IP address on the interface from the remote router, and if the old LDP adjacency times-out after the new LDP adjacency is up, the ISIS protocol will be notified about old LDP adjacency down event and the LDP sync state will remain in hold-down even if the new LDP adjacency is up. [PR955219](#): This issue has been resolved.
- We add timer for all aggregate LDP prefixes but are not deleting it when the timer expires because of a bug. Since the timer is not expiring, we never update the route for any change. This will be sitting in the routing table as a stale entry. Issue is being fixed in later versions. [PR956661](#): This issue has been resolved.

Platform and Infrastructure

- In the Network Time Protocol (NTP) configuration, if the specified source ip address is not in current routing-instance, the router will use primary address of interface (which will be used to send packet) as source address, Client routers will treat the NTP packets as incorrect packets, and then NTP synchronization failed. [PR872609](#): This issue has been resolved.
- When tagged frames with larger than MTU size are received, some frames are not counted as oversized frames on 20x1GE MIC. [PR879519](#): This issue has been resolved.
- After interface reset, CoS information may not be applied correctly to the Packet Forwarding Engine, leading to inconsistency in scheduling/shaping in Qx Chip. [PR908807](#): This issue has been resolved.

- In a MX-VC environment, in certain situations the inter-chassis traffic may not be equally balanced across all available vcp links after adding extra links. [PR915383](#): This issue has been resolved.
- The system MAC address is not getting saved in a unified in-service software upgrade (ISSU) blob and it is not getting programmed again by the Routing Engine when the Packet Forwarding Engine re-connects. The hash seed is generated by using the system MAC address and since it is not saved in a unified ISSU blob, after an ISSU it is 0 and the hash seed is generated using that. If a FPC reboot, then it will get the correct system MAC address and generate the hash seed based on that. This will cause different FPCs in the system to have different hash seeds and could cause AE multicast traffic loss if the ingress and egress FPCs have different hash seeds. [PR915933](#): This issue has been resolved.

- In subscriber management scenario, memory leak might occur when the firewall fast-update-filter feature is configured, and it will impact any new subscriber login. Such memory leak can be seen with following command, root@router> show chassis fpc

Temp CPU Utilization (%) Memory Utilization (%)

Slot State (C) Total Interrupt DRAM (MB) Heap Buffer

0 Online Absent 8 0 1024 70 << 13

1 Online Absent 8 0 1024 29 13 [PR926808](#): This issue has been resolved.

- Under certain timing conditions the MPC/TFEB can receive the firewall filter configuration before it is fully booted/UP/ONLINE. Because the firewall filters can depend on certain default values which are not yet programmed the MPC/TFEB will crash/core-dump and reboot/restart/reload. [PR928713](#): This issue has been resolved.
- When replacing ichip FPC with MX Series FPC, "traceroute" packets going through a MX Series FPC may experience higher drop probability than when using an Ichip FPC. [PR935682](#): This issue has been resolved.
- On MX Series routers with DPC type FPCs running a 11.4 (or later) Junos OS release disabling family inet with uRPF enabled on a logical interface might result in another logical interface on the router to drop all incoming IPv4 packets. The lookup index is calculated by taking the lower 16 bits of the logical interface index (also called the IFL index). In other words lookup index = IFL index MOD 65536. It is normal, valid and expected to have logical interfaces which share the same lookup index. The problem described in this PR is not the fact that the lookup indexes are the same. Here is an example of two different logical interfaces on two different FPCs which share the same lookup index: Interface ge-1/1/0.0 has an IFL index of 141073 and a lookup index 10001: > show interfaces ge-1/1/0.0 Logical interface ge-1/1/0.0 (Index 141073) (SNMP ifIndex 2318) ^^^^^^ Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Input packets : 0 Output packets: 0 Protocol inet, MTU: 993 ^^^^^ Flags: Sendbcst-pkt-to-re, uRPF ^^^^^ Addresses, Flags: Is-Preferred Is-Primary Destination: 1.1.1.0/30, Local: 1.1.1.1, Broadcast: 1.1.1.3 Protocol multiservice, MTU: Unlimited Flags: Is-Primary And interface ge-2/0/7.1647 has an IFL index of 10001 and a lookup index of 10001: > show interfaces ge-2/0/7.1647 Logical interface ge-2/0/7.1647 (Index 10001) (SNMP ifIndex 20551) Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1647] Encapsulation: ENET2 Input

packets : 0 Output packets: 0 Protocol inet, MTU: 8978 Flags: Sendbcst-pkt-to-re, uRPF, uRPF-loose Protocol multiservice, MTU: Unlimited In the example above if family inet is disabled on ge-2/0/7.1647 then ge-1/1/0.0 will start dropping all incoming packets silently. [PR936249](#): This issue has been resolved.

- On TXP system false "SIB Cell Drop Error" alarm might be raised for LCC-SIB after autohealing CRC errors on corresponding HSL2 channel.

This alarm should be treated as a false one as there are no drops of valid data cells.

This alarm considered false only when it was raised during fabric autoheal. Fabric autoheal log can be checked as:

```
> show chassis fabric errors autoheal 2013-12-04 18:40:52 CET Req: LCC0 plane 4 ln 15 LCC-to-SFC
```

2013-12-04 18:40:52 CET Succeeded: LCC0 plane 4 ln 15 LCC-to-SFC [PR937330](#): This issue has been resolved.

- On front panel display LED status for PSM is incorrect after manually Remove/Insert of PSM. [PR937400](#): This issue has been resolved.
- "Total errors" counter of MAC statistics on MX DPC(ge/xge) is always 0. [PR942183](#): This issue has been resolved.
- TWAMP connection/session will come 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. [PR943320](#): This issue has been resolved.
- In PPPoE subscriber management environment, if the BRAS router is MX Series router with MS-DPC equipped and traffic from the subscribers is NATed on MS-DPC card, when PPPoE subscribers flap, heap memory leak might occur on the MS-DPC. [PR948031](#): This issue has been resolved.

Routing Policy and Firewall Filters

- Policy with Install-nexthop lsp may not work as expected when there is a LSP path change triggering route resolution. [PR931741](#): This issue has been resolved.

Routing Protocols

- When the IPv6 address on fxp0 is active during boot up, the joining of the all-router group causes the kernel to create a ff02::2 route with a private nexthop, which is not pushed to the Packet Forwarding Engine. When a non-fxp0 interface is active later, the private nexthop will be shared by the non-fxp0 interface as well, resulting in packet drops destined to ff02::2 on the non-management interface. - After this PR, the advertising interface should be configured via the following CLI. [edit protocols] + router-advertisement { + interface. <interface_name>; + } [PR824998](#): This issue has been resolved.
- When inter operate with Cisco router, OSPF adjacency might be brought down by Cisco end, if Junos OS CPU is high and LSA ACK is delayed for over 2 minutes. [PR846182](#): This issue has been resolved.

- "show route advertising-protocol bgp <nbr> table foo.mvpn.0" stops working after PR-908199 fix. [PR929626](#): This issue has been resolved.
- On the first hop router if the traffic is received from a remote source and the accept-remote-source knob is configured, the RPF info for the remote source is not created. [PR932405](#): This issue has been resolved.
- If you have fix for PR-929626, Avoid the following show command in a VPN setup "show route advertising-protocol bgp <nbr_addr> table foo.inet.0" Where <nbr_addr> is peer within routing-instance "foo" [PR936434](#): This issue has been resolved.
- In MVPN scenario, while performing CLI command "show route advertising-protocol bgp <neighbor>", the rpd might crash due to a timing issue that BGP rib for bgp.mvpn-inet6.0 table is NULL. [PR940491](#): This issue has been resolved.

Services Applications

- Max number of supported IPSec tunnels might depend on networking activity as well. Under heavy networking activities, while DPD (Dead Peer Detection) is enabled, the maximum number of supported IPSec tunnels can drop to about 1800. [PR780813](#): This issue has been resolved.
- In Carrier Grade NAT scenario, MS-PIC might crash and core dump when Port Block Allocation (PBA) block size is relatively big (8192 ports per block), this issue usually happens when a new block need to be allocated because the block currently is exhausted. [PR874500](#): This issue has been resolved.
- In the Session Initiation Protocol (SIP) Application Layer Gateway (ALG) with port block allocation enabled scenario ("user@root# set services nat pool <pool-name> secured-port-block-allocation block-size <block-size>"), a SIP call to be set up and the ports block are allocated for the media flows. When the SIP media flows time out, the APP mapping starts using another port block. But if no enough port block to be allocated, the services Physical Interface Card (PIC) might crash. [PR915750](#): This issue has been resolved.
- In the IPsec scenario, when all available SAs are expired and the sequence number is wrapping for the IPsec packets, the Physical Interface Card (PIC) will delete the Security Association (SA), however this is not reported back to key management process (kmd). This would cause kmd and the PIC being out of sync regarding the known IPsec SAs, then the traffic blackhole might occur. [PR933026](#): This issue has been resolved.
- No SNMP trap generated when NAT or Flow sessions reach the threshold. [PR933513](#): This issue has been resolved.
- Interim-logging is now supported with NAT64 on microkernel (MS-DPC) platforms. The same pba-interim-logging-interval knob under 'service-options' under the service interface will enable the feature for NAT64 as well. [PR935606](#): This issue has been resolved.
- FW is trying to create a new pair of flows while a drop flow with the same selector is being installed for traffic initiated from the outside by a different CPU. There is a race condition while accessing the flow type field: - CPU1 (installing the drop flow) - creates the flow and adds it to the flow table while holding the corresponding bucket lock. However, the flow type field is filled in later. - CPU2 (installing another flow with the

same selector as the drop flow above) CPU2 will find the entry added by CPU1 but will fail to notice that it corresponds to a drop flow because the type field hasn't been set yet by CPU1. This will lead to checking if there is any softwire info available for the existing flow. The drop flow is installed for outside traffic so no softwire information is available causing the assertion to fail. [PR940014](#): This issue has been resolved.

- Snmp traps are not generating when port utilization threshold is crossed. [PR941931](#): This issue has been resolved.
- During a rare scenario, switchover on another sp interface can crash a service PIC when running a traffic in hairpinning scenario. [PR945114](#): This issue has been resolved.

Software Installation and Upgrade

- In this case, since the high level package (i.e. jinstall) is signed, the underlying component packages are not required to be signed explicitly. However the infra was written such a way to display warning message if the component package is not signed (i.e. jpfe). [PR932974](#): This issue has been resolved.

Subscriber Access Management

- Radius attribute ignore logical-system-routing-instance not ignoring VSA26-1. [PR953802](#): This issue has been resolved.
- Configuration change of the IPv4 address range in address-assignment pool does not always take effect. [PR954793](#): This issue has been resolved.

User Interface and Configuration

- If a configuration file which contains groups related configuration is loaded by command "load replace", a "commit confirmed" operation might fail. When this issue occurs, the new configuration is committed even if you do not confirm it within the specified time limit. [PR925512](#): This issue has been resolved.

VPNs

- Configuration version (child rpd) of rpd generates a core file when doing a commit or commit check. [PR930080](#): This issue has been resolved.
- The issue happens when the virtual routing forwarding (vrf) is configured "no-vrf-propagate-ttl" and the vrf import policy changes the local preference of the vrf route. With "no-vrf-propagate-ttl", BGP will resolve the primary l3vpn route and the vrf secondary route separately. The root cause is overwriting the route parameters of the second vrf route with the route parameters of the primary route. So when changes the local preference of the vrf route might not work. [PR935574](#): This issue has been resolved.
- 'show route table VRF.mvpn.0 extensive|detail' for mvpn VRF routing tables will not show BGP TSI info (which previously contained the MVPN PMSI attribute) for outgoing MVPN route advertisements. Since PR 908199, TSI info for these routes is shown on the copy of the route advertised from the main bgp.mvpn.0 table. 'show route table VRF.mvpn.0 extensive|detail' now shows the MVPN PMSI attribute in the main body of the route output. [PR939684](#): This issue has been resolved.

Resolved Issues in Release 12.3R5

Class of Service (CoS)

- The names "best-effort", "assured-forwarding", "expedited-forwarding", "network-control" are reserved and cannot be currently used in Forwarding Class alias configuration, with several classes mapped to the same queue: `user@router# show class-of-service user@router# set class-of-service forwarding-classes class best-effort queue-num 0 user@router# set class-of-service classifiers inet-precedence test forwarding-class best-effort loss-priority low code-points 000 user@router# commit` check configuration check succeeds `user@router# set class-of-service forwarding-classes class myBE queue-num 0 user@router# commit` check [edit class-of-service classifiers inet-precedence test forwarding-class] 'best-effort' forwarding class undefined: best-effort error: configuration check-out failed [PR827496](#): This issue has been resolved.

Forwarding and Sampling

- After committing some configuration changes (e.g. deactivate an interface), while the Packet Forwarding Engine daemon (PFEd) tries to get statistics of some nodes, it may encounter a NULL ncode, causing PFEd to crash and generate a core file. [PR897857](#): This issue has been resolved.
- When pfed get restarts during a period when pfed is communicating with mib2d, because the communication sockets have been terminated and failed to be re-opened after pfed back up again, mib2d might crash and generate a core file. The core files could be seen by executing CLI command **show system core-dumps**. [PR919773](#): This issue has been resolved.

General Routing

- Only 94 GRE(plain) sessions are in Established state after chassisd restart. [PR801931](#): This issue has been resolved.
- BFD packets sent from FPC (distributed mode) over normal physical interfaces are set with ttl 0 so that it gets decremented by 1 and becomes 255 once it is sent out on the wire. This behavior is not the case when the BFD packets are sent over IPsec routed tunnels where the packets are sent from the corresponding service PIC. In this case, the ttl should be set to 255 as no such decrement action takes place when it is sent from a service PIC. But in the current scenario, the ttl is set to 0 as a result of which the service pic drops the outgoing packet. This was an untested scenario till date. [PR808545](#): This issue has been resolved.
- When the 10x10GE PIC (PD-5-10XGE-SFPP) is configured to run in linerate-mode under [set chassis fpc fpc-number pic pic-number] hierarchy, and an input-scheduler-map with Class of Service (CoS) queues including any of queue 4 to queue 7 is applied to an interface on the 10x10GE PIC, the ingress queues may not map correctly to the internal hardware ingress queues, hence as a result, packet drops may be seen in a higher priority queue than that which is expected. [PR818605](#): This issue has been resolved.
- IPv6 address syntax on rpd log is violation of RFC 5952. For example, 2002:db8:0:0:1:0:0:1 must be logged as 2002:db8::1:0:0:1 in the logs, but it's logged

as 2002:db8:0:0:1::1. 2001:0:0:0:db8:0:0:1 must be logged as 2002::db8:0:0:1 in the logs, but it's logged as 2001:0:0:0:db8::1. [PR840012](#): This issue has been resolved.

- If a router receives the BGP keepalive at time t, the next keepalive is expected at time t+30 secs (+/- 20% jitter). However, right around the time when the next keepalive is expected to be received, the BGP keepalive packet is dropped due to some network issue (e.g. uplink towards peer flaps). During this scenario, retransmission of BGP keepalive message on BGP peer would take long time and the BGP session will be terminated due to hold timer expiry. [PR865880](#): This issue has been resolved.
- SNMP trap is not generated upon Fabric chip failure/offline/online state on MX Series routers. [PR877653](#): This issue has been resolved.
- When syslog feature is configured in firewall filter, one of the Junos OS message creating function has a bug, where the whole string is copied directly with no check for overflow. This could easily overflow and results in no null-termination which causes memory corruption and linecard crash. The core files could be seen by executing CLI command "show system core-dumps". [PR888116](#): This issue has been resolved.
- Backup Routing Engine failed to commit with error "pdb_update_ddl_id: cannot get new id for " dynamic-profiles dynamic-profiles profile-name"", commit full is a workaround. [PR888454](#): This issue has been resolved.
- TLB: Observed a traffic-drd daemon hang once after logging into service PIC and restarting the net-monitord process. This is not an operational procedure, not always reproducible, and the work-around is to restart the traffic-drd daemon using the restart traffic-dird command. [PR889982](#): This issue has been resolved.
- When a bgp route is resolved using a next-hop that is also learned in bgp (i.e. there are multiple levels of next-hop resolution) and bgp multipath is also used, during a route churn next-hop for such a bgp route could be incorrectly programmed. [PR893543](#): This issue has been resolved.
- An MX-VC NSR master switch might put kernel control socket in stale state, and in the subsequent NSR master switch, the kernel will refuse the connection from FPC. As a result, the FPC would be rebooted during the switchover process. [PR896015](#): This issue has been resolved.
- Some ATM interfaces may stay down after flapping the Circuit Emulation MIC. [PR900926](#): This issue has been resolved.
- 100G Ethernet interface (Finisar FTLC1181RDN3-J3) on T4000 type-5 FPC might flap once after bringup. The solution is to change the register bandwidth. [PR901348](#): This issue has been resolved.
- RPD on backup Routing Engine might hit out of memory condition and crash if BGP protocol experiences many flaps. [PR904721](#): This issue has been resolved.
- What was seen was that on certain occasions that the 10GE PHY does not recover from the transition from down->up. This can happen in cases where a link flaps or an SFP+ is inserted. What was also seen that the same set of events did not yield the same set of outcomes i.e. JTAC was not even to replicate this issue in the labs even though this was seen in the field occasionally. Upon further analysis it was found that the PHY was stuck in Freeze state as an explicit command to take it out of that state

was able to recover the link. There was no other issue found except this while all the other parameters of optical power the pluggable optics were all taken into account. [PR905589](#): This issue has been resolved.

- "set system ddos-protection protocol sample aggregate bandwidth" command is not taking effect. This can cause packet loss in ukernel for Routing Engine based sampling if sampling rate exceeds 1000pps. [PR905807](#): This issue has been resolved.
- bootp configuration on TXP platform referencing routing-instance fails to commit. [PR906713](#): This issue has been resolved.
- MX-VC: VC port conversion not working for second set of added VC ports for VCB. [PR906922](#): This issue has been resolved.
- VCMm-power down creates stale vlan demux0 entries at the Packet Forwarding Engine level. [PR908027](#): This issue has been resolved.
- When adding the "no-tunnel-services" knob under vpls protocols of routing-instances, during the processing gap of the new knob, if routing protocol process (rpd) restarts (i.e rpd crashes), logical interfaces with VPLS family do not show up, and there are no logical interfaces available for the corresponding VPLS routing instances. Hence VPLS connections might be down (stuck in LD state) and can not be recovered automatically. [PR912258](#): This issue has been resolved.
- After changing interface description, it doesn't get updated in "show lldp neighbors" output. [PR913792](#): This issue has been resolved.
- 10GbE interface on MIC3-3D-10XGE-SFPP stays up even if far end is disabled and goes down. Since the interface on MIC3-3D-10XGE-SFPP cannot react to remote failure, CCC circuit cannot change the state correctly, if port of MIC3-3D-10XGE-SFPP is configured as CCC end point. [PR914126](#): This issue has been resolved.
- The following note applies to 16x10GE MPC: With respect to this feature, when ISSU is performed from feature non-supporting version (ex. 12.2, 13.1) to feature supporting version (12.3R5, 13.2R3, or 13.3), then 16x10GE FPC needs to reboot in order to use this feature. [PR914772](#): This issue has been resolved.
- A log message "%DAEMON-3: Cannot perform nh operation ADDANDGET nhop 0.0.0.0 type unicast nhindex 0x0 ifindex 0xd3e <interface name> fwd nhidx 0x0 type unicast errno 45 suppressed <number of suppressed> logs" is generated if access-internal route is created during the dynamic interface configuration process. The log message can be permanent or not. Besides this message there were no side effects. [PR917459](#): This issue has been resolved.
- FPC crash can be triggered by a SBE event after accessing a protected memory region, as indicated in the following log: "System Exception: Illegal data access to protected memory!" The DDR memory monitors SBEs and reports the errors as they are encountered. After the syslog indicates a corrupted address, the scrubbing logic tries to scrub that location by reading and flushing out 32-byte cache line containing that location in an attempt to update that memory location with correct data. If that memory location is read-only, it causes illegal access to protected memory exception as reported and resets the FPC. The above-mentioned scrubbing logic is not needed because even if SBE is detected, the data is already corrected by the DDR and CPU has a good copy of the data to continue its execution path. PR/919681 can be triggered on both PTX

and T4000 platforms and can be seen in Junos OS releases 12.1 and 12.3. Fix is available in 12.3R5, 12.3R3-S6, 13.3R1, 13.2R2. Crash signature in the FPC shell shows the following:

```
SNGFPC4(router-re0 vty)# sh nvram System NVRAM : 32751 available bytes, 2477
used, 30274 free Contents: [LOG] Set the IP IRI for table #1 to 0x80000014 [LOG]
IPV4 Init: Set the IP IRI to 0x80000014 [LOG] GN2405: JSPEC V 1.0 Module Init. <..>
Reset reason (0x84): Software initiated reset, LEVEL2 WATCHDOG [Sep 6 17:16:07.231
LOG: Warning] <164>DDR: detected 3 SDRAM single-bit errors [Sep 6 17:16:07.231
LOG: Warning] <164>DDR: last error at addr 0x108d2378, bad
data/mask0x00240401ffffff7/0x000000000000000008 bad ecc/mask=0xbe/0x00
System Exception: Illegal data access to protected memory! <<< Event occurred at:
Sep 6 17:16:07.231087 Juniper Embedded Microkernel Version 12.1X48-D30.1 Built by
builder on 2012-08-23 04:28:12 UTC Copyright (C) 1998-2012, Juniper Networks, Inc.
All rights reserved. Context: Thread (Periodic) Registers: R00: 0x10354900 R04:
0xffffffff R08: 0x00000001 R12: 0x24002084 R16: 0x7040c48c R20: 0xff700000
R24: 0x10de0000 R28: 0x00000003 MSR: 0x00029200 CR: 0x44002048 ESR:
0x00000000 R01: 0x13fcb8c8 R02: 0x13fc99f8 R03: R05: 0x00000001 R06:
0x00000000 R07: R09: 0x00000000 R10: 0x00000000 R11: R13: 0x53564110 R14:
0x210191df R15: R17: 0x05240140 R18: 0x3e023840 R19: R21: 0x447ac8dc R22:
0x10deb408 R23: R25: 0x00000000 R26: 0x10df0000 R27: R29: 0x13fcb8d8 R30:
0x00000000 R31: CTR: 0x00000000 Link: 0x10354908 SP: XER: 0x20000000
DEAR: 0x00000000 PC: K_MSR: 0x00001000 0x108d2360 0x00000001
0x00000001 0x2e105223 0xd855014f 0x10de0000 0x13fcb930 0x108d2378
0x13fcb8c8 0x100425ec Stack Traceback: Frame 01: sp = 0x13fcb8c8, Frame 02: sp
= 0x13fcb928, Frame 03: sp = 0x13fcb958, Frame 04: sp = 0x13fcb988, Frame 05: sp
= 0x13fcb998, Frame 06: sp = 0x13fcb9c8, Frame 07: sp = 0x13fcb9f8, [LOG] syslog
called with interrupts off (caller pc:0x10549988) [LOG] Dumping core-SNGFPC4 to
1 [LOG] syslog called with interrupts off (caller pc:0x105489c0) [LOG] Coredump
finished! [LOG] Set the IP IRI for table #1 to 0x80000014 [LOG] IPV4 Init: Set the IP
IRI to 0x80000014 [LOG] GN2405: JSPEC V 1.0 Module Init. pc pc pc pc pc pc pc =
= = = = 0x10354900 0x100461a0 0x1003c5fc 0x1003c894 0x1003c5fc 0x1003c7b4
0x10030b8c Reset reason: Software reset <..> PR919681: This issue has been resolved.
```

- Following chassisd messages might be observed after executing the "show chassis fabric summary" command, FM: Plane Sate: 1 1 1 2 2 0 0; staggered_pmask: 15 2a 00 00 00 00 00 00 FM: Mux active/trained: 0/0 0/0 0/0 0/0 0/0 0/0 0/0 0/0; Mode:1 act_mask:3f These are non-impacting debug messages. [PR927453](#): This issue has been resolved.
- MS-PIC might crash in IPsec environment after deleting "tcp-mss" knob under IPsec "service-sets" hierarchy. [PR930741](#): This issue has been resolved.
- Polling the OID mib-2.17.7.1.4.3.1.5...: dot1qPortVlan or mib-2.17.7.1.4.5.1...: dot1qPvid might cause a memory leak on the l2ald process, and the process might create core files. [PR935981](#): This issue has been resolved.

High Availability (HA) and Resiliency

- On TX or TXP Line Card Chassis (LCC) with graceful Routing Engine switchover enabled, if a mastership switch is being requested on a LCC whose backup Routing Engine's em0 interface is physically failed (due to hardware failure or driver stops working), this

will cause all FPCs on the LCC disconnect from the old master Routing Engine, but cannot reconnect to the new master one either. [PR799628](#): This issue has been resolved.

- During every failover of redundancy-group 0, the /etc/ssh and /var/db/certs directories are copied from primary node to secondary node. However, the directories are not copied correctly and nested directories such as /etc/ssh/ssh, /etc/ssh/ssh/ssh are created. [PR878436](#): This issue has been resolved.
- In certain systems configured with graceful Routing Engine switchover (GRES), there is the possibility for the master and the backup Routing Engine to reach an inconsistent view of installed state. This fault may be exposed if the master Routing Engine experiences a mastership watchdog timeout at a time when it is not in sync with the backup Routing Engine for a particular piece of state. In practice, this possibility exists only for a short time period after a Routing Engine mastership change. Under such conditions, a replication failure may cause the backup Routing Engine to panic. If the failure is seen, the backup Routing Engine will recover on restart. In 11.4 and 12.1 releases without this fix, the fault may be experienced on any GRES-enabled, non-multichassis configuration on a T Series router. For 12.2 and later releases without this fix, the fault may be experienced on any GRES-enabled, non-multichassis configuration on a T Series or MX Series router. [PR910259](#): This issue has been resolved.

Infrastructure

- If a router receives the BGP keepalive at time t, the next keepalive is expected at time t+30 secs (+/- 20% jitter). However, right around the time when the next keepalive is expected to be received, the BGP keepalive packet is dropped due to some network issue (e.g. uplink towards peer flaps). During this scenario, retransmission of BGP keepalive message on BGP peer would take long time and the BGP session will be terminated due to hold timer expiry. [PR865880](#): This issue has been resolved.
- When multicast is running on a multi-chassis environment, during flapping of 224/4 or ff00/8 pointing to mResolve(NH), the LCC master might get replication error which causes all FPCs going offline. This flapping of resolve route for multicast can occur because of any of the following reasons: enabling or disabling multicast, deletion of resolve route, or routing restart. [PR897428](#): This issue has been resolved.

Interfaces and Chassis

- DCD reports error when configuring hierarchical-scheduler on MX80 with QX chipset. This is cosmetic error and it should not have functional impact. [PR807345](#): This issue has been resolved.
- An MX Series router may cosmetically log "Bottom Fan Tray Unable to Synch". [PR833047](#): This issue has been resolved.
- Tx and Rx Spanning-tree BPDU stopped intermittently during ISSU. [PR849201](#): This issue has been resolved.
- The particularity of logic that DCD daemon crashes when "aggregated-ether-options load-balance" is committed. [PR854207](#): This issue has been resolved.
- M7i Routing Engine crashed with last reboot reason panic:page fault and kernel core, after commit. [PR868212](#): This issue has been resolved.

- "Link down" alarms should never exist on the VC Protocol Backup Routing Engine. They should only be on Protocol Master, if any. The bug is that the "Link down" alarms are not cleared from the Protocol Backup after/during a graceful Routing Engine switchover (GRES) event. Restarting alarmd removes these alarms from the Protocol Backup. [PR886080](#): This issue has been resolved.
- If an AE interface is brought down by protocol CFM/LFM/STP, the interface will go down permanently and can not recover automatically. [PR888728](#): This issue has been resolved.
- While a duplicate interface address (IFA) is configured for two interfaces, software will accept that and generate an error message like this:
%CONFLICT-4-DCD_PARSE_WARN_INCOMPATIBLE_CFG: [edit interfaces ge-0/0/0 unit 0 family inet address x.x.x.x/xx] : Incompatible configuration detected : identical local address is found on different interfaces But at kernel side cannot accept duplicate IFA, and needs to delete the next-hop created for this operation. Due to code problem, the cleanup doesn't remove the duplicated IFA under heavy kernel workload. And it will crash while trying to update this duplicated IFA to Packet Forwarding Engine side. [PR891672](#): This issue has been resolved.
- Following is the document change proposed :- traceroute-ethernet-command :- Source MAC address : MAC address of 802.1ag node responding to the LTM Next-hop MAC address: MAC address of egress interface of the node where LTM would be forwarded show-oam-ethernet-connectivity-fault-management-linktrace-path-database-command :- Source MAC address : MAC address of 802.1ag node responding to the LTM Next-hop MAC address: MAC address of egress interface of the node where LTM would be forwarded The display of next-hop MAC address is incorrect for linktrace path database command. [PR895710](#): This issue has been resolved.
- On MX Series platforms which are running Junos OS Release 12.3R3/R4 and 13.2R1 and Operation Administration and Maintenance (OAM) is activated, periodic packet management daemon (ppmd) might crash after changing the ppm distribution state from distributed to centralized and then restarting connectivity fault management daemon (cfmd) or any action to create ppm interface. It is suggested to deactivate and activate OAM configuration during this configuration change to avoid the problem. [PR905812](#): This issue has been resolved.
- The MX Series router does not always process the first LCP request for a static PPPoE subscriber. [PR908457](#): This issue has been resolved.
- Issue is because of vrrpd not configuring vrrp group id, and state when it's in transition state. In normal scenario when vrrp moves to master it signals dcd to add the VIP. When VIP gets added vrrpd gets a notification and updates state and group id corresponding to that VIP. While updating state vrrpd checks the current state. If state is master it updates state as master and if its backup it updates it as backup. But if vrrp state is in transition it does not do anything. It may not be seen every time and is a timing issue. One can confirm the incorrect mac address by capturing monitor traffic on the affected irb interface specifically one who is master VRPP instance. [PR908795](#): This issue has been resolved.
- When an interface is configured with VRRP protocol, IP address associated with this interface might disappear after deactivating then activating the interface. When this

issue happens, KRT may be getting stuck and never clean up. If the interface belongs to a routing-instance, then deactivate/activate the routing-instance can also trigger the same issue. Issue command 'show krt queue' to verify: root@ABC-re0> show krt queue Routing table add queue: 1 queued ADD table index 37, gf 1 (1377) error 'File exists' [PR912295](#): This issue has been resolved.

- In multicast over AE scenario, if there is a different order of child logical interfaces (logical interface) under parent AE at the master Routing Engine and the backup Routing Engine, then after Routing Engine switchover, multicast traffic might get lost. [PR915440](#): This issue has been resolved.
- For IQ2 PIC, when the setting shaping rate is too high, when configured it with "set chassis fpc 0 pic 1 traffic-manager logical-interface-base-shaping-rate 16" and this will reset the shaping rate to 1Gbps. The corresponding messages are logged in debug level. In the fix, it is corrected into info level. [PR920690](#): This issue has been resolved.
- In MX-VC environment, if LT interface's encapsulation type is ethernet-ccc, after rebooting FPC with LT interfaces or rebooting system, the LT interface might not come up again. [PR922673](#): This issue has been resolved.
- Unified ISSU fails on upgrade to 11.4R5.7 with the following message Logged messages: MIC 4/0 will be offlined (In-Service-Upgrade not supported) MIC 4/1 will be offlined (In-Service-Upgrade not supported) Do you want to continue with these actions being taken ? [yes,no] (no) yes error: /usr/sbin/indb failed, status 0x200 error: ISSU Aborted! Chassis ISSU Aborted ISSU: IDLE Issue happens when a MIC-3D-4OC3OC12-1OC48 card is offline via CLI and removed from the chassis prior to the ISSU. [PR923569](#): This issue has been resolved.
- Traffic which uses MPLS next-hops enters bridge-domain via IRB interface and if forwarding next-hop moves from non-aggregate interface to aggregate interface (MAC move), the MPLS next-hops are not correctly programmed in the Packet Forwarding Engine. The child next-hop of the aggregate interfaces are missing. Once IRB MPLS next-hop moves from aggregate interface to non-aggregate interfaces are not affected. IPV4 traffic is not affected. [PR924015](#): This issue has been resolved.
- The MX960 works as LNS can't accept l2tp session packet from Huawei GGSN. [PR926919](#): This issue has been resolved.

Layer 2 Features

- If STP is configured on AE interface, the l2cpd might be under high utilization and VRRP repeatedly flaps after the VRRP active router reboots. The root cause here is when STP is configured on AE interface, the corresponding Bridge Protocol Data Unit (BPDU) will messages go to Routing Engine instead of processed in Packet Forwarding Engine. [PR882281](#): This issue has been resolved.
- In VPLS environment, while deactivating/activating VPLS routing-instances, in rare conditions, routing protocol process (rpd) tries to free an already used route, then rpd process crashes with core files. [PR908856](#): This issue has been resolved.

- "show snmp mib walk ascii jnxVpnIfStatus" doesn't work for BGP vpls when there is incompleted BGP VPLS instance configuration or LDP VPLS instance. [PR918174](#): This issue has been resolved.
- In BGP autodiscovery for LDP VPLS scenario, as FEC129 VPLS does not support nonstop active routing (NSR), VPLS fails to come up after Routing Engine switchover and traffic will never resume. [PR919483](#): This issue has been resolved.

MPLS

- This message was used to record error condition from nexthop installer. Over time, it becomes common function and same message will be printed in many valid conditions, leading to confusion on these message-logs. [PR895854](#): This issue has been resolved.
- IPv6 traceroute may not show some hops for scenarios where 1) Two LSPs are involved. 2) INET6 Shortcuts are enabled. In such scenarios, hops that are egress for one LSP and ingress for the next LSP in the traceroute do not show up. This was a software issue with icmp error handling for packets with ipv6 payload having a ttl of 1. [PR899283](#): This issue has been resolved.
- RPD might crash under specific conditions and after executing "ping mpls l2vpn interface <interface>" command. [PR899949](#): This issue has been resolved.
- If the maximum-ecmp next-hops under [edit chassis] hierarchy is configured as 32 or 64 (more than the default value of 16), the routing protocol process (rpd) might crash on the new master Routing Engine after performing graceful Routing Engine switchover. The root cause here is while merging nexthops, the Junos OS is iterating over only 16 gateways instead of configured maximum-ecmp number and finally results in an assert. The core files could be seen by executing CLI command **show system core-dumps**. [PR906653](#): This issue has been resolved.
- The output of "show ldp overview" command regarding graceful restart is based on per protocol LDP graceful restart settings. Where graceful restart is enabled by default. So when graceful restart is disabled this command shows it's enabled for LDP. However graceful restart should be enabled globally for LDP graceful restart to operate. [PR933171](#): This issue has been resolved.

Network Management and Monitoring

- Mib2d may get ATM VPI updates before the ATM IFDs are learned. In such cases instead of discarding the updates mib2d has started caching them untill the IFD is learned. [PR857363](#): This issue has been resolved.
- In an IS-IS scenario, with traceoptions enabled under protocol ISIS and syslog level set to debug under routing-options options for a router, if the router has two IS-IS neighbors which have the same router-id configured, after configuring the same ISO system-id on these two IS-IS neighbors, routing protocol process (rpd) on the router will crash with core files dumped. [PR912812](#): This issue has been resolved.

Platform and Infrastructure

- Commit may fail, when a config object is deleted and re-added as transient change from a commit script. [PR814796](#): This issue has been resolved.
- When the egress Next Hop is IRB interface, the Routing Engine to Packet Forwarding Engine output function for the IRB overwrites some info on the packet mbuf to indicate the underlying egress Layer2 interface, with the expectation that the UDP packet needs to be forwarded. However, in the case of this being a traceroute UDP packet with a limited TTL, the TTL limit check decides to drop the packet and issue an icmp response to the sender. Because of the mentioned IRB overwrite, the icmp code could no longer determine how to send a packet to the originator. The fix is to re-arrange the sequence of events so that the overwriting of the mbuf info is done later in the sequence. After all the possible layer3 (IP) TTL checks are done. [PR816202](#): This issue has been resolved.
- Since the AC Power System on the MX2020 is a N+N feed redundant and N+1 PSM redundant, there are two separate input stages per PSM, each connected to one of the two different/redundant feeds. However, only one stage is active at a time. This means, the other input stage (unused input stage) may be bad and system will not know about it till it tries to switch to it in case of a feed failure. This is a pretty bad corner case and needs to be addressed. The way to work around this problem is by testing both stages when the power supply is first powered on. This test is done by the system software and an alarm is raised if any feed failure is detected. [PR832434](#): This issue has been resolved.
- There are two symptoms covered this issue: If there is a mix of high and low priority fabric traffic as can be seen by checking "show class-of-service fabric statistics", the following error messages can be seen when there are bursts of high priority fabric traffic, while low priority fabric traffic is present :- May 6 14:58:41 routename-re0 fpc1 MQCHIP(0) FI Reorder cell timeout May 6 14:58:41 routename-re0 fpc1 MQCHIP(0) FI Cell underflow at the state stage A second symptom with this mix of low and high priority fabric traffic present; if an FPC that is the recipient of this high and low priority fabric traffic restarts, it is possible for the ingress FPC forwarding ASIC to lockup. In this case the following log message might be simultaneously logged :- Jun 5 13:46:50 router fpc4 MQCHIP(0) CPQ Queue underrun error, Qsys1 Queue 42 Jun 5 13:46:50 router fpc4 MQCHIP(0) CPQ Freecnt nearing empty error, Qsys mask 0x2 [PR877123](#): This issue has been resolved.
- High rate of traffic to the Routing Engine may cause control traffic stoppage to the Routing Engine. The indication is the following type of messages: "WEDGE DETECTED IN PFE ... TOE host packet transfer: reason code 0x1 [PR896592](#): This issue has been resolved.
- On MX Series router with MPC, firewall filter counter doesn't count packets when firewall is configured on discard interface. [PR900203](#): This issue has been resolved.
- If there are private sessions in place, it should not abort the effective/revoke of conditional groups. In affected releases, it is not working. [PR901976](#): This issue has been resolved.
- In MX-VC setup using virtual-switch instance type, there can be scenarios where the outer vlan-tag of PPPoE/PADI packets on egress can be stripped off when ingress

interface is a LAG with 2 member links spread across the 2 Chassis members. [PR905667](#): This issue has been resolved.

- Junos OS 12.3R3, 12.3R3S1 and 12.3R3S2, interfaces with interface-mode trunk connected on top PFE[0] and with IRB interfaces, might corrupt forwarding-state on lowest Packet Forwarding Engine of the FPC. This is applicable to system operating with network-services enhanced-ip mode and systems operating in virtual-chassis mode. [PR907291](#): This issue has been resolved.
- Command "show ddos-protection protocols" doesn't report correct Arrival and Max arrival pps rates. One bit of rate value at Packet Forwarding Engine is incorrectly set which results in a wrong ddos rate value. [PR908803](#): This issue has been resolved.
- In MX virtual-chassis (MX-VC) scenario, when the VC-M (master member of VC) reboots and then comes up, the MPC with virtual-chassis port (vcp) configured might crash due to the memory overflowed. [PR910316](#): This issue has been resolved.
- The DDOS classification for Dynamic Host Configuration Protocol (DHCP) "leasequery" message is not working. This message is treated as "unclassified". [PR910976](#): This issue has been resolved.
- IPv6 UDP checksum is implemented, but computed UDP checksum for IPv6 IPFIX export packets gets invalid occasionally. When this issue is seen the following capture would be seen in the collector. 14:05:06.810436 In Juniper PCAP Flags [Ext, no-L2, In], PCAP Extension(s) total length 16 Device Media Type Extension TLV #3, length 1, value: Ethernet (1) Logical Interface Encapsulation Extension TLV #6, length 1, value: Ethernet (14) Device Interface Index Extension TLV #1, length 2, value: 139 Logical Interface Index Extension TLV #4, length 4, value: 71 -----original packet----- PFE proto 6 (ipv6): (hlim 64, next-header: UDP (17), length: 144) xxxx:xxx:ffff:ffff::yy.33068 > xxxx:xxx:0:yyy::yyy.2055: [bad udp cksum 72ff!] UDP, length 136 (IPv6 address masked). [PR911972](#): This issue has been resolved.
- When enhance-route-memory is enabled along with SCU, configuration might cause Jtree Memory corruption on MX Series DPCs. [PR914753](#): This issue has been resolved.
- Description of T4000 midplane changes after "show snmp mib walk jnxContentsDescr" [PR915393](#): This issue has been resolved.
- On MX2020, SNMP traps are generated only for SFB slot 6 and 7 upon graceful Routing Engine switchover (GRES) enabled Routing Engine switchover. [PR915423](#): This issue has been resolved.
- Changing the domain-name doesn't reflect in DNS query unless a Commit full is done. This bug in management daemon (mgd) has been resolved by ensuring mgd propagates the new domain-name to file /var/etc/resolv.conf, so that this can be used for future DNS queries. [PR918552](#): This issue has been resolved.
- Issue observed in inline Jflow during route-record collection. For route-record function in inline-Jflow it is expected that for any aggregated type next hops a child next-hop must be present. This child next-hop info is updated as gateway info for aggregated next-hop. In scenario, where we have valid aggregated next hop id but no child next-hop, system is crashing in inline-jflow during route-record collection. [PR919415](#): This issue has been resolved.

- Without this PR fix, commit script applied configuration may emit the XNM RPC errors when there is XML tag mismatch detected: error: [filename: xnm:rpc results] [line: 771] [column: 7] [input: routing-engine] Opening and ending tag mismatch: routing-engine line 7 and rpc-reply error: [filename: xnm:rpc results] [line: 773] [column: 6] [input: rpc-reply] Opening and ending tag mismatch: rpc-reply line 6 and junoscript error: [filename: xnm:rpc results] [line: 774] [column: 2] [input: junoscript] Premature end of data in tag junoscript line 2 [PR922915](#): This issue has been resolved.
- DDOS_PROTOCOL_VIOLATION alarm shows incorrect timestamps <time-first-detected> and <time-last-detected> on messages. Both fields indicate the same timestamps. Timestamps <time-first-detected> and <time-last-detected> are overwritten. [PR927330](#): This issue has been resolved.
- If port-mirror is used on member-0 interfaces and mirrored on member-0 interfaces, mirrored traffic is incorrectly sent to member-1 and dropped. Fabric drop counters are counted. [PR928315](#): This issue has been resolved.
- Under certain timing conditions the MPC/TFEB can receive the firewall filter configuration before it is fully booted/UP/ONLINE. Because the firewall filters can depend on certain default values which are not yet programmed the MPC/TFEB will crash/core-dump and reboot/restart/reload. [PR928713](#): This issue has been resolved.
- When replacing ichip FPC with MX Series FPC, "traceroute" packets going through an MX Series FPC may experience higher drop probability than when using an ichip FPC. [PR935682](#): This issue has been resolved.

Routing Policy and Firewall Filters

- Install-nexthop lsp-regex does not work as expected when multiple recursive routes share same protocol next hop having different export policy with regular expression option. Route is not updated with correct export forwarding nexthop as same nexthop select handle is calculated for any set of configured export policy with "install-nexthop lspregx" option. [PR863341](#): This issue has been resolved.
- Junos OS releases with a fix for PR/706064 have a regression where the vrf-import policy sanitation logic is faulty. A "# commit check" will fail when the first term references a 'target' community and the second term references an 'origin' community. This should pass the check. [PR911350](#): This issue has been resolved.

Routing Protocols

- With this fix, "jnxBgpM2PrefixesInPrefixesRejected" counter will return the number of prefixes from a BGP peer, that are not eligible to become active. This change makes the variable conform to definition in the specification <http://tools.ietf.org/html/draft-ietf-idr-bgp4-mibv2-03>. There is a new variable "jnxBgpM2PrefixInPrefixesActive" introduced, to return the number of active prefixes from a BGP peer. So the new sequence of variables for the table is as follows:

```
root@root> show snmp mib walk jnxBgpM2PrefixCountersTable
jnxBgpM2PrefixCountersAfi.0.1.1 = 1 jnxBgpM2PrefixCountersSafi.0.1.1 = 1
jnxBgpM2PrefixInPrefixes.0.1.1 = 0 jnxBgpM2PrefixInPrefixesAccepted.0.1.1 = 0
jnxBgpM2PrefixInPrefixesRejected.0.1.1 = 0 jnxBgpM2PrefixOutPrefixes.0.1.1 = 3
jnxBgpM2PrefixInPrefixesActive.0.1.1 = 0
```

PR778189: This issue has been resolved.
- Junos OS label block allocation can only return block size as power of 2 (e.g. 2, 4, 8, 16,...). In inter-as option-b L2VPN scenario, routing protocol process (rpd) core is seen when the ASBR receives a non-power-of-2 label block size from other vendor's device. The root cause here is when rpd requests the non-power-of-2 label block size, an assert occurred. The core files could be seen by executing CLI command **show system core-dumps**. **PR848848**: This issue has been resolved.
- When configuring CAC for a physical interface, the software might enable CAC for unit 0 on that interface, but might not be able to delete it when the configuration is removed. **PR850578**: This issue has been resolved.
- There are improper `</route-family>` tags added to all **multicast route summary** commands when we perform command such as **show multicast route summary | display xml**. **PR859104**: This issue has been resolved.
- On T640/T1600 routers with Enhanced Scaled (ES) FPCs equipped and all MX Series routers with MPC, the Bidirectional Forwarding Detection (BFD) sessions over aggregated Ethernet (AE) interfaces might be down after performing unified In-Service Software Upgrade (ISSU). Note, the problem is only seen on FPC (Packet Forwarding Engine) based BFD (contrasts with Routing Engine based BFD), and the problem is mostly seen on T640/T1600 routers even though the problem affects MX Series routers in principle. **PR859324**: This issue has been resolved.
- The remote discriminator is not reinitialized after bfd session state moves to down (with diagnostic code: control detection time expired) as per RFC 5880 requirement. **PR889970**: This issue has been resolved.
- In a scenario with graceful restart (GR) enabled for BGP between Cisco platform and Juniper Networks platform, Junos OS is helper (default) and Cisco being restarting router, when Cisco restarts BGP process, Juniper deletes all BGP routes due to doesn't receive End Of RIB (EOR) markers for all configured NLRI from Cisco. **PR890737**: This issue has been resolved.
- BGP "accepted-prefix-limit" feature might not work as intended when it is configured together with "damping". Root cause of this issue is that when BGP module counts the maximum routes accepted from BGP neighbor, it doesn't count the accepted BGP routes which are in damping status. So when these damping routes are reused, the

total number of received BGP routes exceeds the configured value for "accepted-prefix-limit" . [PR897124](#): This issue has been resolved.

- In PIM dense mode, if the Assert loser router receives a join/prune (S,G) message with upstream neighbor is the loser router, it should send an Assert(S,G) on the receiving interface to initiate a new Assert negotiation to correct the downstream router's RPF neighbor, but our device will not. [PR898158](#): This issue has been resolved.
- Sometimes "Advertised prefixes" counter for some RIBs may be incorrect for some BGP neighbors. This is a cosmetic issue. Use "show route advertising-protocol bgp <nbr> table <tblname> | match Nexthop | count" to know the right advertised prefixes count. [PR899180](#): This issue has been resolved.
- In multicast scenario with PIM enabled, when you configure both static RP mapping with override knob and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for a given group range, and allow dynamic RP mapping for all other groups, but a software defect cause that RP is selected based on dynamic RP mapping address, instead of accounting for this static override knob. [PR912920](#): This issue has been resolved.
- DR sends a delayed ACK to the LSA on the interface on which the LSA is flooded. This leads to BDR sending only directed ACK to DR, DR-Other is therefore not receiving this ACK and is hence retransmitting the LSA to BDR. [PR914803](#): This issue has been resolved.
- When the interface goes down, the direct route for that peer address is removed from the routing table before BGP processes interface down event and brings down the session. When BGP calculates multipath routes, since the knob "accept-remote-nexthop knob" is configured, BGP needs to determine whether we can reach the nexthop address (ebgp peer address) directly. BGP did not find direct route for this nexthop address and so asks for route nexthop resolution. In this case, the first BGP path from the peer with up interface has direct router nexthop, the second path is set to have indirect nexthop due to the down interface, BGP passed a wrong mixed multipath nexthop, which caused RPD crash. [PR917428](#): This issue has been resolved.
- When NSR is configured and path-selection is changed, there might be a non-functional impacting softcore generated during the commit process. [PR928753](#): This issue has been resolved.
- "show route advertising-protocol bgp <nbr> table foo.mvpn.0" stops working after PR-908199 fix. [PR929626](#): This issue has been resolved.
- If you have fix for PR-929626, avoid the following show command in a VPN setup "show route advertising-protocol bgp <nbr_addr> table foo.inet.0" Where <nbr_addr> is peer within routing-instance "foo" [PR936434](#): This issue has been resolved.

Services Applications

- NAPT: Packet Forwarding Engine side report port range starts from 512 because napt mib counter wrong. This fix make the port range in Packet Forwarding Engine start from 1024. [PR828450](#): This issue has been resolved.
- Any port or IP address value set in SIP VIA header for 'rport' and 'received' attributes will not be checked or translated by the SIP ALG. There is usually no impact from this to a voice call. The contact address inserted by the client in future requests will be the external one but this will not disrupt the SIP ALG. Some rare clients however may have some unexpected reaction that causes problem such as trying to register two IP addresses, the internal one AND the public one, in the same register message which is unsupported by the ALG and causes the message to be dropped. [PR869725](#): This issue has been resolved.
- When a snmp query is running that accesses information from service PIC, and during that an MS-DPC or service PIC restarts, then the adaptive services process (spd) may hang. As a result the thread will never complete, and we will never be able to update and delete the routes through RPD (e.g. routes that point to NAT pool ranges are marked as dead routes because they still point to old logical interfaces). This can result in routes in incorrect state and black-holing of traffic. [PR874347](#): This issue has been resolved.
- MIB-NJX-L2TP syntax errors, commas missing on line 401 and 930. [PR881423](#): This issue has been resolved.
- The jpppd crash on LNS happened because the size of the udp based l2tp packet exceeded the buffer length available. The modification was done to discard the packet instead of creating core. [PR888691](#): This issue has been resolved.
- In rare conditions with large number of traffic flows (like NAT and IPsec flows), the Service PIC may get stuck or crash as a result of prolonged flow-control assertions towards the Packet Forwarding Engine. In order to trigger this issue, many Compute CPUs inside the Service PIC should be overloaded. This will never happen under normal operation, where CPUs can handle large amount of traffic without any issues. [PR900227](#): This issue has been resolved.
- The SIP ALG is unfit for EIM due to standing limitations, hence, SIP and EIM is currently unsupported configuration. [PR900412](#): This issue has been resolved.
- In Carrier Grade Network Address Translation (CGNAT) environment, if memory utilization of MS-DPC/service PICs are in the yellow zone and they are configured with cgn-pic knob, there can be a race condition where there are two timers created for the same flow and during the timer processing, the MS-DPC/service PIC may experience a crash and generate a core file. [PR901795](#): This issue has been resolved.
- In an L2TP scenario, after performing an SNMP walk of "jnxL2tpTunnel" or "jnxL2tpSession" MIBs, the SNMP reply message fails to be written because write buffer is exceeding MTU, causing Routing Engine CPU spikes to 100%. [PR905218](#): This issue has been resolved.

- In some cases rtsp data flows will be left without cleanup when rtsp master flow closes. This will cause some conversation data flows left on router with very huge timeout values. [PR909091](#): This issue has been resolved.
- IKE UDP 500 packet is not processed in correct routing-instance. [PR909909](#): This issue has been resolved.
- In a CGNAT environment, active FTP operations fail when there is latency issue in network. When TCP retransmission, FTP ALG is not translating any fields in the Request: PORT command. As a result server tries to establish the data flow to the private IP address and to a wrong TCP port and it fails as expected. [PR916376](#): This issue has been resolved.
- In Carrier Grade Network Address Translation (CGNAT) with high memory utilization environment (Memory is in yellow zone and use CLI "show services service-sets memory-usage" to check), this crash might be seen in hairpinning scenario where Endpoint Independent Filtering (EIF) is enabled and the initial packet of a specific flow that hits the MS-DPC is dropped by an ALG due to various reasons (malformed or non complying packet/headers). [PR918663](#): This issue has been resolved.
- In Carrier Grade NAT (CGNAT) environment, during heavy setup rate of CGNAT flows, High Availability (HA) sync flaps and then keepalive messages are lost, as there is no control flow prioritization configured. HA sync connection keeps disconnecting. After a long period of time PIC silently reboots. Following syslog message might be seen when issue occurs: Sep 7 16:39:29 ROUTER-RE0 (FPC Slot 2, PIC Slot 0) PFEMAN: Lost contact with master routing engine PFEMAN: Forwarding will cease in 4 minutes, 59 seconds Sep 7 16:40:23 ROUTER-RE0 (FPC Slot 3, PIC Slot 1) PFEMAN: Lost contact with master routing engine PFEMAN: Forwarding will cease in 4 minutes, 59 seconds. [PR920723](#): This issue has been resolved.
- In Carrier Grade Network Address Translation (CGNAT) environment whenever an inbound UDP packet did not hit any rule, a check is performed whether the destination ip and port match any SIP registration. If this check is successful and 'learn-sip-register' is enabled (which is the default in the junos-sip application), if packets are counted as SIP ALG parsing errors, no flow is created and the packet will be forwarded without any transformation. In the case of NAT, the destination address will remain within the nat pool and the packet will keep coming back to the service PIC, causing a routing loop and high CPU utilization. [PR923630](#): This issue has been resolved.
- If multiple service sets with different number of NAT rules/pools are configured, Services PIC might crash when SNMP walk is performed on jnxSrcNatStatsTable. [PR928169](#): This issue has been resolved.
- When tcp session is initiated from inside client and three-way handshake is not completed because that client did not ack the syn-ack send from the server, the service pic will send a tcp reset to the server after the timer expires. In this case tcp reset is send in the wrong direction. Instead of sending in the outbound direction to the server, the service pic will send it in the inbound direction. This PR fixes this issue. No service impact is seen because of this. [PR931433](#): This issue has been resolved.

Subscriber Access Management

- Due to some timing issues, MX Series was generating incorrect LLPDF logs "LLPDF: llpdf_client_connection: Unknown session" every 10 seconds. [PR894013](#): This issue has been resolved.
- This netstat core can be generated during certain upgrade/downgrade scenario. The fix must be present in the image being upgraded/downgraded to. The trigger for having this netstat core generated is having secureld configuration present on the chassis. When the validate phase of "request system software add" runs, the netstat core may be generated. [PR911232](#): This issue has been resolved.
- Authd - UserAccess log events not sent to syslog host when destination-override is used. [PR931975](#): This issue has been resolved.

VPNs

- In an FEC129 VPLS scenario, VPLS pseudowire (PW) processing might hit an assert, causing rpd process to crash with a core file. [PR843482](#): This issue has been resolved.
- In affected releases, the C-PIM Assert mechanism is not working correctly in a Multicast VPN environment. A typical scenario includes an access VLAN with four routers (CE1, CE2, PE1 and PE2) which are C-PIM neighbors of each other. If CE1 sends a PIM Join to PE1, and CE2 sends a C-PIM Join to PE2, both PEs start to inject the C-Multicast flow in the access VLAN. This triggers the PIM Assert mechanism, which should result in either PE1 or PE2 (one of them, not both) injecting the traffic. However the following two situations might occur during 1 minute or more: - Both PE1 and PE2 keep injecting traffic in the VLAN. - Both PE1 nor PE2 stop injecting traffic in the VLAN. Releases with the fix work fine regarding the PIM Assert mechanism and do not show this abnormal behavior. [PR880575](#): This issue has been resolved.
- In L2circuit scenario, after L2circuit established, if pseudowire flaps (e.g. interface flapping) while routing protocol process (rpd) processing this change, memory corruption might occur, causing rpd process to crash with core files. [PR900257](#): This issue has been resolved.
- In L2circuit scenario, if there is no MPLS route to neighbor and there is a static route with discard nexthop in inet.3 table as follows:

```
user@router# show routing-options
rib inet.3 {
  static { route 0.0.0.0/0 discard; } }
```

Then the l2circuit connection will use the above static route in inet.3 table to connect its neighbor as follows:

```
user@router# run show route table mpls
mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden) + = Active Route, - =
  Last Active, * = Both 0
*[MPLS/0] 00:01:42, metric 1 Receive 1
*[MPLS/0] 00:01:42, metric 1 Receive 2
*[MPLS/0] 00:01:42, metric 1 Receive 13
*[MPLS/0] 00:01:42, metric 1 Receive ge-0/0/1.601
*[L2CKT/7] 00:01:38, metric2 0
```

Discard In this situation, routing protocol process (rpd) will core dump while walking Simple Network Management Protocol (SNMP) MIB "jnxVpnPwEntry". [PR906519](#): This issue has been resolved.

- This PR enables default advertisement of MVPN from the main BGP routing tables `bgp.mvpn.0` and `bgp.mvpn-inet6.0` instead of VRF routing table `foo.mvpn.0` or `foo.mvpn-inet6.0`. It also removes withdraw suppression for extranets. If extranets are used, `advertise-from-main-vpn-table` is enabled by default for an MVPN NLRI. [PR908199](#): This issue has been resolved.
- In Rosen and NG-MVPN running in `rpt-spt` mode, valid (*,G) forwarding state can be created (it can not be created in `spt-only` mode). If there is `rpf-check-policy` added to MVPN instance and the `rpd` check is associated on the (*,g) forwarding route installation, the `rpd` might crash. [PR915672](#): This issue has been resolved.

Resolved Issues in Release 12.3R4

Class of Service (CoS)

- When 'scheduler-map-chassis derived' configuration is used under class-of-service, interface related configuration changes can lead to `cosd` process crash. [PR863734](#): This issue has been resolved.
- During addition/deletion or just deletion of interfaces with configuration for shared scheduler, some portion of memory is not reclaimed back normally. So continuous addition/deletion of these interfaces results in memory depletion, packet loss and other issues. [PR890986](#): This issue has been resolved.

Forwarding and Sampling

- Possibility of duplicate packets when sampling and interface-style nat are configured. [PR861984](#): This issue has been resolved.
- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 or MX DPC, if there is distributed Bidirectional Forwarding Detection (BFD) running on Aggregated interface and a firewall filter is configured on loopback interface (lo0), the lo0 will bind an implicit filter, after FPC restarts or the Routing Engine switchover, the next hop of the implicit filter is not updated with the corresponding link word to point to CLI filter, causing the CLI filter to be not executed. To resolve the issue, deactivate the firewall filter under loopback interface and then activate it again. Note: The default operational mode of `bfd` for all protocols is distributed mode (runs on pfe), one exception being `ospf v3` which runs on the Routing Engine by default (centralized mode). + + So `ospf v3` is not affected by this issue. [PR864665](#): This issue has been resolved.
- Outbound control traffic is not counted by accounting-profile which applied to logical interfaces of AE (Aggregated Ethernet). [PR866181](#): This issue has been resolved.
- In T4000 platforms with ES-FPC, for IPv6 firewall filters with match conditions on address prefixes longer than 64 bits, in some corner cases, the filter may not be correctly evaluated and packet loss may occur. [PR879829](#): This issue has been resolved.
- `lab@T1600-2_Critical_VZB_Manjit> show services accounting flow-detail destination-prefix 20.1.1.2/32`

Service Accounting interface: sp-2/0/0, Local interface index: 147

Service name: (default sampling) Interface state: Accounting

Protocol Input Source Source Output Destination Destination Packet Byte Time since
last Packet count for Byte count for interface address port interface address port count
count active timeout last active timeout last active timeout udp(17) xe-0/0/3.0 10.1.1.2
whois++ (63) xe-0/0/2.0 20.1.1.2 whois++ (63) 1075917 49492182 00:17:55 1780922
81922412 tcp(6) xe-0/0/3.0 10.1.1.2 0 xe-0/0/2.0 20.1.1.2 0 106479 4898034 00:01:46
1835070 84413220 [PR881629](#): This issue has been resolved.

- In scaled MPLS scenario, when LSP path switchover happens, sample process deletes sampling parameters from the Packet Forwarding Engine and as a result of that Packet Forwarding Engine stops exporting flows to the collector. [PR891899](#): This issue has been resolved.
- When router goes into Amnesiac mode with 'commit failed' due to statements constraint check failed while upgrading Junos OS to 11.4 or later, ARP Replies will be dropped due to incorrect default arp policer on interface even after fixing the commit errors. [PR895315](#): This issue has been resolved.
- After committing some configuration changes (e.g. deactivate an interface), while the Packet Forwarding Engine daemon (PFEd) tries to get statistics of some nodes, it may encounter a NULL node, causing PFEd to crash and generate a core file. [PR897857](#): This issue has been resolved.

General Routing

- When an MPC fails in a specific manner, while failing it continues to send traffic into the switching fabric for a time, the fabric ASICs report errors such as these with large counts: `chassisd[82936]: %DAEMON-3: New CRC errors found on xfchip 0 plane 0 subport 16 xfport 4 new_count 17651 aggr_count 17651 chassisd[82936]: %DAEMON-3: New CRC errors found on xfchip 0 plane 0 subport 17 xfport 4 new_count 17249 aggr_count 17249 chassisd[82936]: %DAEMON-3: New CRC errors found on xfchip 0 plane 0 subport 18 xfport 4 new_count 65535 aggr_count 65535`

This can cause DPC(s) to stall and not send traffic into the switching fabric to other DPCs or MPCs. Messages such as these may be reported by the affected DPC(s) :

```
[Err] ICHIP(1)_REG_ERR:packet checksum error in output fab_stream 4 pfe_id 64 [Err]
ICHIP(1)_REG_ERR:packet checksum error in output fab_stream 6 pfe_id 64 [Err]
ICHIP(1)_REG_ERR:packet checksum error in output fab_stream 8 pfe_id 64
```

This failure on the affected DPCs persists, and will likely affect all traffic destined to the fabric from affected DPCs. The only temporary resolution is to restart the affected DPCs, which will resume fabric traffic from the affected DPCs. [PR856560](#): This issue has been resolved.

- ATM MIC back-to-back, to many logical interfaces(more than 8k) may cause certain logical interfaces down. [PR859165](#): This issue has been resolved.
- When the fxp0 interface on a k2re is administratively disabled, the local end shows the link as down while the far end device displays the status as up. [PR862952](#): This issue has been resolved.

- During a reference clock switch T4 will be switched off. [PR868161](#): This issue has been resolved.
- The 1588v2 BMCA procedure causes a frequency hold-over event in the system under test. [PR868422](#): This issue has been resolved.
- Configuration of Container Interfaces for APS on MX FPCs is not allowed since Junos OS 12.1. If this feature is needed on MX Series legacy FPCs use a release with this PR fixed. [PR869192](#): This issue has been resolved.
- Under high scale, expiry of a Kernel side reconnect timer would cause it to send a non-servicable message to the Packet Forwarding Engine(asking the line cards to restart and resync since reconnect failed). Since there is no ack- to this kernel message, kernel thought it sent the message and untoggles the GRES flag. The Packet Forwarding Engine wasn't expecting anything so it continued along. The EFFECT: The system is permanently not ready for GRES... CLI GRES check will always report: [cmd] request chassis routing-engine master switch check Apr 14 19:03:13 [INFO] warning: Standby Routing Engine is not ready for graceful switchover. [PR873679](#): This issue has been resolved.
- On systems containing XM-based linecards(for example, MPC3, type 5 FPCs), if a member link of an aggregate Ethernet (AE) bundle is repeatedly flapped, the flapped member link may stop transmitting traffic. Traffic isn't getting dropped, as the remaining member-links will pick up the slack. But in some cases (the traffic is large or some members encounter the problem together), traffic loss will happen. [PR875502](#): This issue has been resolved.
- On MX-VC platform, when the master Routing Engine declares GRES ready by CLI command, there is a time window before some FPCs to be actually ready. After performing GRES, these GRES unready FPCs might get rebooted, resulting in traffic loss. [PR877248](#): This issue has been resolved.
- authd reports syntax error, although the syntax is correct, when trying to activate service profile for subscriber and fails to activate the service. [PR883065](#): This issue has been resolved.
- lrmuxd core seen when committing changes related to BD or routing-instance. Below messages appears and commit fails cbecker@raggpoc01# commit error: Check-out pass for logical system multiplexer process (/usr/sbin/lrmuxd) dumped core (0x86) error: configuration check-out failed. This Issue is resolved now. [PR883090](#): This issue has been resolved.
- The Routing Engine might become non-responsive due to the exhaustion of kernel mbufs with following messages. /kernel: Mbuf: High Utilization Level: (Low) Throttling low priority requests (10 ms) /kernel: Mbuf: High Utilization Level: (Medium) Throttle low priority requests (150 ms) /kernel: Mbuf: High Utilization Level: (High) Block low priority requests [PR886083](#): This issue has been resolved.
- RPD might core dump if HFRR (Host Fast Reroute) is enabled on two logical interfaces in the same routing instance for IPv6 and if link-local address is configured on those logical interfaces. The core files could be seen by executing CLI command **show system core-dumps**. [PR886424](#): This issue has been resolved.

- The backup Routing Engine failed to commit with error "pdb_update_ddl_id: cannot get new id for " dynamic-profiles dynamic-profiles profile-name"", commit full is a workaround. [PR888454](#): This issue has been resolved.
- When multiple framed-route(type-22) AVPs are present in Radius access accept message, the router will install only the first route into the routing table. [PR891036](#): This issue has been resolved.
- Following a global GRES event, the new Master(VC-Mm) will expect relayd to reconnect to it in less than 40 seconds. However under high scale, such as with 54k dual-stack(v4v6) or 110k+ single-stack DHCP subscribers, owing either to a slow relayd(relay daemon) control connection to the Kernel, or due to slow Packet Forwarding Engine reconnects to relayd, we are not able to meet the 40 seconds timer requirement causing subsequent FPC reboots and traffic loss. [PR891814](#): This issue has been resolved.
- When performing DSCP rewrite on LMNR or T640-FPC4-ES and PC-OC-768 cards, the ECN bits will get reset to 0. [PR896847](#): This issue has been resolved.
- Interfaces on FPC slot 5 will disappear after offlining of any fpc(not fpc-slot 5) on M40e/M160. [PR898415](#): This issue has been resolved.
- In subscriber management environment, in a rare case, VLAN auto-sensing daemon (autoconfd) might crash and create a core file due to Session Database (SDB) is inaccessible. [PR899747](#): This issue has been resolved.
- Some ATM interfaces may stay down after flapping the Circuit Emulation MIC. [PR900926](#): This issue has been resolved.

Infrastructure

- Unsolicited Neighbor Advertisement is not sent from backup when vrrp switchover is initiated. [PR824465](#): This issue has been resolved.
- Kernel may crash when delete routing instance under the donor and unnumbered address borrower scenario. When the deleting for the donor is before the deleting of the corresponding unnumbered borrower, in this window, the donor interface does not have an address, and arp processing over the borrower interface during this window may trigger the crash. The core files could be seen by executing CLI command **show system core-dumps**. [PR880179](#): This issue has been resolved.
- Every 10 minutes kernel reports "%KERN-6: MTU for 2001:4c0:1:1301:0:1:0:250 reduced to 1500" after reducing MTU once. There is no impact to the system due to this additional log message. [PR888842](#): This issue has been resolved.
- In a multihop IPv6 BGP session scenario, after configuring single-hop BFD session on the multihop IPv6 BGP neighbor, kernel might try to access a NULL pointer, causing kernel to crash and generate a core file. [PR898153](#): This issue has been resolved.
- Checksum error seen on ICMP reply when 'sequence, data' field in request set to '0'. [PR898487](#): This issue has been resolved.

Interfaces and Chassis

- On E1 interface, when interface flaps on CE side of connection, interface will flap a second time on the PE side. [PR690403](#): This issue has been resolved.
- Traffic loss is seen. Multiple inbound and outbound IPSEC tunnels are created for a single SA during tunnel renegotiation after the lifetime expiry. [PR827647](#): This issue has been resolved.
- Planes might go into faulty state during the SCB initialization when the SERDES on the SF chip failed to come up. [PR839509](#): This issue has been resolved.
- IQ2 core is seen after unified ISSU and traffic will be lost for awhile(about 40s). The crash happens during processing of scheduler free message which comes just after unified ISSU complete on IQ2. Then the heap structure is invalid causing panic. The fix is moving the process to unified ISSU sync stage. [PR845257](#): This issue has been resolved.
- On the following MIC-3D-20GE-SFP only, if the 1GE interface is put into loopback mode, all packets larger than 306 Bytes are truncated on the wire. The solution is to bring the interface down once loopback is configured, to prevent truncated packets to be sent out. [PR856892](#): This issue has been resolved.
- Dump-on-flow-control knob might not work correctly for RSP interfaces configured in "warm-standby" mode. After an RSP switchover, either manual or following a crash, the dump-on-flow-control flag might get cleared from the MS-PIC. [PR867394](#): This issue has been resolved.
- snmpwalk of "jnxPPPoEIfLockoutTable" didn't capture pppoe locked out clients. [PR869024](#): This issue has been resolved.
- MC-LAG will no longer change just the LACP System Identifiers directly, but will also remove the "Synchronization, Collecting, Distributing" bits from the Actor State bits advertised in the PDU. [PR871933](#): This issue has been resolved.
- Injecting Enhanced RDI-P(G1 bit5-7:0x2 Payload defect) alarm to a MPC 10GbE WAN-PHY interface causes RDI_P and LCD-PAIS-V alarm on messages. This is due to string typo. RDI_P and LCD-P should be printed on messages. [PR872133](#): This issue has been resolved.
- On MX Series router with MPC with 20port GE MIC, interface stores packets when disabled and transmits stored packets after enabled. [PR874027](#): This issue has been resolved.
- In subscriber management environment, with dynamic-profiles configured for subscribers, if the routing instance returned from radius is not configured on BRAS, dynamic-profile add fails and there are some places the memory not freed, causing device control daemon (dcd) memory leak. The memory usage of dcd process can be observed by following command:

user@router> show system processes extensive | match dcd

PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND 7076 root 1 97 0 1047M 996M select 6:05 2.88% dcd [PR880235](#): This issue has been resolved.

- MX Series router is not passing transit IPv6 traffic received on a RLSQ interface with fib-localization enabled. [PR880245](#): This issue has been resolved.
- VC-Boot loop when installing new local backup Routing Engine. [PR881906](#): This issue has been resolved.
- Problem scenario: CFM UP MEP for Bridge/VPLS is configured on MPC with action profile as 'interface down' Problem statement: When the CFM sessions go down due to network outage at the core, action profile is triggered and the configured interface is brought down. When the Core network failure is corrected, CFM will not automatically recover because the interface will continue to remain down. [PR884323](#): This issue has been resolved.
- On LAG interface gratuitous ARP is neither generated nor sent out upon link up even when gratuitous-arp-on-ifup is configured. [PR889851](#): This issue has been resolved.
- In dynamic PPPoE subscriber management environment, when MS-DPC card is added and "adaptive-services service-package laryer-2" is configured, while PPPoE subscribers log in, kernel might encounter a memory corruption, causing kernel to crash and generate a core file. [PR894440](#): This issue has been resolved.
- The C-LMI (Consortium LMI) is supported on all I-chip based FPC. Support for the MX-FPC 2 and 3 was missing and now added. [PR895004](#): This issue has been resolved.
- On MX Series based platforms, when PIC is configured with traffic-manager mode ingress-and-egress, after PIC offline, PIC detach does not clean up the corresponding entries completely. Subsequent PIC online results in corresponding entries add failure since previous entries are still intact, resulting in interface attach failure at the Packet Forwarding Engine level. Due to interface add failure, protocols on the interface never come up. [PR895305](#): This issue has been resolved.
- IPv6 IIF-index load-balance works unwantedly when IIF-V4 is enabled alone and vice versa. [PR898676](#): This issue has been resolved.
- On front panel display LED status for PSM is incorrect after manually Remove/Insert of PSM. [PR937400](#): This issue has been resolved.

Layer 2 Features

- When VPLS is configured with GRES, the backup Routing Engine responds to certain route replication requests by simulating address learning. If the route being replicated is associated with an LSI or VT interface, the address learning code references a special LSI or VT nexthop. Thus, there is a dependency between that route and that nexthop. This fix is to explicitly enforce this ifstate dependency, ensuring that the special nexthop is seen by the peer before the route. [PR867929](#): This issue has been resolved.
- In Releases 12.1R3, 12.2R3, 12.3R4, 13.1R1, and 13.2R1, for a configuration with bridge domains containing aggregate interfaces, traffic whose destination address is broadcast, multicast, or unknown will not be load-balanced across the member links of such interfaces. Instead, all such traffic will be sent out a single link of the aggregate interface. With this PR change, load-balancing will always be applied to such configurations for traffic whose destination address is broadcast, multicast, or unknown. This change restores the functionality of older releases. [PR888232](#): This issue has been resolved.

Layer 2 Ethernet Services

- jdhcpd interface traceoptions are not saved to the default log file jdhcpd and require an explicit file name. [PR823129](#): This issue has been resolved.
- New knob is provided to set the prefix to compare requested ip and server address. Knob is configured as - [edit system services dhcp-local-server] #set requested-ip-network-match <0-31> For V6 [edit system services dhcp-local-server] #set dhcpv6 requested-ip-network-match <0-127> Default will be 8 for v4 and 16 for v6 (first terms). [PR872145](#): This issue has been resolved.
- When IPv6 is configured on integrated routing and bridging (IRB) interfaces that have AE interfaces as child links, after GRES was enabled and one child link failed or was removed, the kernel crashed. [PR878470](#): This issue has been resolved.
- DHCPv6 Local Server implementation deletes the client on a reconfigure, so that client can reconfigure. DHCPv6 relay is not forwarding the Reply to the client and simply tearing the client down (generating a release to the server). [PR879904](#): This issue has been resolved.
- When executing **show dhcp relay binding** command with high scales of bound subscribers and with several hundred renewing at a given time, DHCP drops the renew packets. [PR882834](#): This issue has been resolved.
- In an IP demux/vlan demux configuration, where the primary address for the loopback is different from the preferred in the dynamic profile, the ACK to the first RENEW will have the primary address in loopback as server ID since RENEW arrives on ip demux interface. The client will send the next RENEW to that server ID, and the router will drop it. The fix is to always use the server ID from the underlying interface. [PR890562](#): This issue has been resolved.

MPLS

- The LDP protocol might use the lowest IP address configured on an interface even if there is another (higher) address that is explicitly configured as primary. This can lead to unexpected LDP session flap if the lowest but non-primary address is being removed from the configuration. [PR858838](#): This issue has been resolved.
- In an RSVP environment with AutoBw, the Bandwidth Adjustment timer for new LSPs added simultaneously is not smeared along with the rest of the existent LSPs when the smearing algorithm is triggered. [PR874272](#): This issue has been resolved.
- When BGP labeled-unicast route has BGP label as null and its indirect next-hop requires adding 2 or more labels, traffic using the BGP label may not be forwarded properly. [PR881571](#): This issue has been resolved.
- The VpnId value contains no information, but was being returned as the empty string, when the MIB requires that it be a length 7 octet string. The value (since it contains no information) is now returned as 7 zeros. [PR882828](#): This issue has been resolved.
- With OSPF overload enabled, the te-metric will be set as 2^{32} , and the Constrained Shortest Path First (CSPF) process ignores the path with metric value 2^{32} , causing the ingress LSPs not to come up. [PR887929](#): This issue has been resolved.

- When a LDP egress router advertises multiple prefixes, by default the prefixes are bound to a single label and aggregated into a single forwarding equivalence class (FEC). If the nexthops of some prefixes in the FEC change (e.g. LDP interface flapping), LDP still tries to bind a single label to all of the prefixes which is incorrect. [PR889585](#): This issue has been resolved.
- LSP metric will be not correctly changed as the new configured one after committed when cspf finds an Explicit Route Object (ERO) different from the current ERO and the Path State Block (PSB) re-signaling fails. This is because a change in metric is a local PSB change, but after a configuration change (for example, the bandwidth requirement was changed), PSB and associated routes used to get this change only after a cspf computation followed by a session refresh or re-signaling. If the re-signaling fails, the configured metric value is not updated in the existing PSB and the route metric. [PR894035](#): This issue has been resolved.
- Changing the preference on an LSP was considered a catastrophic event, tearing down the current path and then re-establishing a new one. This PR makes the preference change minor and only needs a new path to be re-signalled in a make-before-break manner. [PR897182](#): This issue has been resolved.
- With Junos OS Release 12.1R1 or later, any configuration changes in the MPLS stanza, P2MP LSP connection with a single branch, will flap and cause brief traffic drops if allow-fragmentation knob is configured under the MPLS path-mtu stanza. No traffic drops are seen if the P2MP LSP has two or more branches. Any application which is using P2MP RSVP LSP is exposed to this issue, like ccc p2mp-transmit-switch, static route with p2mp-lsp-next-hop, etc. [PR905483](#): This issue has been resolved.

Network Management and Monitoring

- When we do snmp polling via CLI on a big MIB node that has lots of OIDs and huge data like "show snmp mib walk 1.3.6.1.4.1", CLI might not be able to consume data at the rate it was being generated by snmpd, so the snmpd buffer is occupied more and more. Eventually this would cause snmpd to reach its rlimit then crash. [PR864704](#): This issue has been resolved.
- SNMP query from valid client on routing-instance-1 with community string that belongs to routing-instance-2 gets the details of routing-instance-2 instead of blocking such queries based on community. [PR865023](#): This issue has been resolved.
- When you perform the following MIB Walk on interfaces, for some interfaces the ifLastChange value will show a value of zero. show snmp mib get ifLastChange.<SNMP ifIndex> will show a value of zero. ifLastChange.<SNMP ifIndex> = 0. [PR886624](#): This issue has been resolved.
- A memory leak in the cosd process is seen when both of the following conditions are met: - multiple OIDs from jnxCos MIB, that are under the same logical interface hierarchy, are queried in a single SNMP query sent to the device (i.e. in a single PDU) - either "per-unit-scheduler" or "hierarchical-scheduler" configured on the physical interface The following messages will be logged when the cosd process exceeds 85% of its maximum usable memory: Jun 9 13:16:35.475 2013 router-re0 /kernel: %KERN-5: Process (1457,cosd) has exceeded 85% of RLIMIT_DATA: used 1894060 KB Max 2097152 KB [PR893464](#): This issue has been resolved.

Platform and Infrastructure

- RMOPD crash is due to sort of buffer overflow crash and library function being used improperly. It is not caused by RPM scaling, This issue happens randomly and hard to point out the specific trigger. [PR277900](#): This issue has been resolved.
- Junos OS 10.4R8 or later on MX Series platforms, L3VPN application using l3vpn-composite-nexthop when the indirect-next-hop configuration statement is added or removed it might cause traffic drops affecting L3VPN flows. To recover from this condition all the l3vpn prefixes need to get removed and installed new into the forwarding-table, like clearing the bgp peers where the routes are learned from. [PR741646](#): This issue has been resolved.
- When changing configuration repeatedly, in rare conditions, some internal errors may cause CLI process hogs memory and the utilization keeps on increasing due to memory leak. When the memory usage of CLI process increases to around 85% of system limit, the following logs could be seen: /kernel: Process (1383,cli) has exceeded 85% of RLIMIT_DATA: used 62048 KB Max 65536 KB The memory will be released once user logout from the router. [PR813673](#): This issue has been resolved.
- In rare case, after no graceful FPC rebooting (i.e. temporary power failure on egress FPC), fabric ASIC on ingress STFPC can run into temporary problematic status. This will cause temporary large delay on fabric traffic from STFPC to the egress FPC. [PR831743](#): This issue has been resolved.
- Since the AC Power System on MX2020 is a N+N feed redundant and N+1 PSM redundant, there are two separate input stages per PSM, each connected to one of the two different/redundant feeds. However, only one stage is active at a time. This means, the other input stage (unused input stage) may be bad and system will not know about it till it tries to switch to it in case of a feed failure. This is a pretty bad corner case and needs to be addressed. The way to work around this problem is by testing both stages when the power supply is first powered on. This test is done by the system software and an alarm is raised if any feed failure is detected. [PR832434](#): This issue has been resolved.
- On MPC 3D 16x 10GE, MPC3, and MPC4 platforms, if host outbound traffic is set to any forwarding class which may maps to queue numbers 4 through 7, after configuring "max-queues-per-interface 4" (4-queue mode is enabled), then queues 4-7 will not be configured with proper traffic parameters, but queue is enabled with default config. When physical interface gets oversubscribed the queue which is carrying host originated traffic can starve for bandwidth because of no q-rate for the queue. Eventually in the event of steady state oversubscription causes loss of control traffic and hence control session can flap. Especially on Junos OS release 12.3R3, after configuring "max-queues-per-interface 4", the corresponding interfaces will get only just 4 queues, it causes 100% loss of host originated traffic because the queues from 4-7 are not enabled for traffic transmission. [PR868021](#): This issue has been resolved.
- When an MX Series router collects with inline jflow, exported IPv6 UDP packets show UDP checksum is incorrectly set to 0x0000, which might be discarded by received node. 12:19:11.513058 In IP6 (hlim 64, next-header: UDP (17), length: 138) 2001:db8:ffff:ffff::20.33068 > 2001:db8:0:100::101.2055: [bad udp cksum 9652!] UDP, length 130 12:19:11.524964 In IP6 (hlim 64, next-header: UDP (17), length: 138)

2001:db8:ffff:ffff::20.33068 > 2001:db8:0:100::101.2055: [bad udp cksum 2086!] UDP, length 130 12:19:16.509978 In IP6 (hlim 64, next-header: UDP (17), length: 138)
2001:db8:ffff:ffff::20.33068 > 2001:db8:0:100::101.2055: [bad udp cksum 1340!] UDP, length 130 [PR870172](#): This issue has been resolved.

- When check trace route, RSVP-TE Probe status is not shown as success and it is shown as unhelpful. Note: seeing this issue with ip-enhance mode and not seeing this issue without ip-enhance in same setup and same image. [PR871015](#): This issue has been resolved.
- After restart of an FPC, when it comes online the queue block on another FPC becomes locked up and all traffic into the fabric from this Packet Forwarding Engine is dropped. The issue occurs when there is a lot of high-priority traffic, and low-priority traffic gets stuck behind and therefore causes the timeout and queue draining. [PR877123](#): This issue has been resolved.
- This is a regression issue introduced by the fix of PR801982, which causes DOM MIB values for SFP+ "rx power" related statistics to be incorrect. Note that XFP is not affected. [PR878843](#): This issue has been resolved.
- If interface flaps of a bridge-domain with igmp-snooping enabled or multicast snooping routes are pruned due to Designated Router changes, LUCHIP might report traps and EDMEM read errors. These conditions are transient and only seen once the system is operating with enhanced-ip mode. [PR879158](#): This issue has been resolved.
- PHP to PE link with MPLS MTU 1300 allows transit traffic more than 1268 i.e. up to 1272. Note: PE to PHP has default MTU in this case i.e. there is MTU mismatch between PHP and PE link. Max packet size allowed is 1300 - 20 (ip) - 8 (icmp) - 4 (1 label due to PHP) = 1268 [PR879427](#): This issue has been resolved.
- Deactive/delete AE interface when route is flapping might cause MX Series routers with MPCs/MICs based Packet Forwarding Engine to crash. [PR884837](#): This issue has been resolved.
- While configuring a filter with a generic prefix followed by specific one in different terms may lead to incorrect match, which might lead to packet drop. [PR886955](#): This issue has been resolved.
- In I2circuit connection scenario, when the STFPC interconnect with MX Series based FPC, PPP-CCC I2circuit connection will drop the small packets with Ethernet length error. [PR887098](#): This issue has been resolved.
- In L2VPN scenario, on the PE router, if the encapsulation of the PE-CE interface is vlan-ccc and there is a COS filter under the interface, when the interface flaps, it can cause all the traffic to different sites via different outgoing interfaces is forwarded incorrectly through one of the interfaces. Meantime, when manually flap the label-switched paths (LSPs) on the router after the problem occurred, the traffic is forwarded incorrectly still but only the egress interface will change to other one. The way to resolve the problem is manually clearing the LSPs on the PE router. [PR887838](#): This issue has been resolved.
- It is observed that in the setup route nexthop for destination of collector's IP address was of type indexed nexthop. [PR889884](#): This issue has been resolved.

- Because of the hardware limit, the feature "maximum-labels" on FPC can't exceed 3. Whenever maximum mpls label is configured as 4 or 5 on unsupported FPC, the LDP/RSVP session will go down and cause MPLS traffic black hole for couple of minutes. This dark window will remain till the unicast next hops are installed and attached to the egress interface where the label has been configure. After that MPLS traffic will resume. [PR890992](#): This issue has been resolved.
- Traffic may be affected after performing an offline/online sequence on the PIC in a T4000 system. This issue is usually seen when the event is performed on PICs carried in a Type 5 FPC. [PR892548](#): This issue has been resolved.
- High rate of traffic to the Routing Engine may trigger credit overflow within the Traffic Offload Engine (TOE) and prevents further processing of packets destined to or originated by the Routing Engine. High rate of traffic means continued Hardware input drops reported via the **show pfe statistics traffic** command. The following message might be seen in the system message log-file: member1-fpc2 TOE Pkt Xfer:** WEDGE DETECTED IN PFE 2 stream 0 TOE host packet transfer: reason code 0x1 The following Junos OS software release 12.3R3, 12.3R3-S1, 12.3R3-S2 or 12.3R3-S3 are exposed on MX Series based FPCs and T4000-FPC5-3D FPCs. [PR896592](#): This issue has been resolved.
- Scheduler with zero guaranteed rate and excess priority none is an invalid class of service configuration. The packet enqueued in the corresponding queue will not be able to dequeued. [PR900239](#): This issue has been resolved.
- In MX-VC setup using virtual-switch instance type, there can be scenarios where the outer vlan-tag of PPPoE/PADI packets on egress can be stripped off when ingress interface is a LAG with two member links spread across the two Chassis members. [PR905667](#): This issue has been resolved.
- Junos OS 12.3R3, 12.3R3S1 and 12.3R3S2, interfaces with interface-mode trunk connected on top PFE[0] and with IRB interfaces might corrupt forwarding-state on lowest Packet Forwarding Engine of the FPC. This is applicable to system operating with network-services enhanced-ip mode and systems operating in virtual-chassis mode. [PR907291](#): This issue has been resolved.

Routing Protocols

- When "passive" and "disable" knobs are both configured under **[edit protocols isis interface <intf> level <N>]** hierarchy the interface is treated as "passive" instead of being disabled. [PR697553](#): This issue has been resolved.
- BFD triggered local-repair(RLI9007) not initiating immediately. RLI 9007 is applicable from 12.2 onwards. [PR825283](#): This issue has been resolved.
- Junos OS checks for mask-length mismatch for OSPF P2P-over-LAN interfaces, but skips the check if an interface has /32 mask configured. In a scenario with OSPF configured between Juniper Networks platform and other vendors' platform, if a /32 mask IP address is configured on P2P-over-LAN OSPF interface of Juniper platform and a non /32 mask IP address is configured on the peer, the OSPF neighbor can establish but Kernel Routing Table (KRT) queue gets stuck. [PR840122](#): This issue has been resolved.

- In BGP scenario, the initial peer flaps and goes down, then a new peer is established, which might cause an rpd core. [PR840652](#): This issue has been resolved.
- Memory leak after deleting a single BFD session. Observed in **show heap** command. [PR840672](#): This issue has been resolved.
- Multicast packets coming with source address as 0.0.0.0 may cause the RPD to crash. [PR866800](#): This issue has been resolved.
- If the SNMP MIB for BGP is walked, the AFI=1, SAFI=5 entries are missing. If an SNMP "get" is performed, the values can be retrieved. [PR868424](#): This issue has been resolved.
- In an IS-IS scenario, when a large number of routes are distributed into IS-IS, IS-IS overload bit will be set due to maximum LSP fragment exhaustion, this is correct. Then delete the IS-IS export policy, after that, the IS-IS overload bit should be cleared. But the number of exported prefix might be incorrect even though the number of export prefix is zero actually. This can cause overload bit to be set always. This is because local-data for prefixes is not freed up and leads to some memory leak. [PR874015](#): This issue has been resolved.
- If a static route is configured and exported into OSPF, and if the static route has the same subnet as an OSPF interface address, then committing configuration changes (even unrelated to OSPF, such as a device's hostname) results in the removal of the static route related to OSPF type-5 link-state advertisement (LSA) from the OSPF database. [PR875481](#): This issue has been resolved.
- Returned attribute values are not in the defined value range of the mib `bgp4PathAttrASPathSegment`. [PR882407](#): This issue has been resolved.
- RPD CPU utilization keeps 100% due to "BGP resync" task when BGP is configured with no neighbor and NSR is configured. `id@router> show configure routing-options nonstop-routing; id@router> show configure protocols bgp { group bgp-group { type internal; inactive: neighbor 1.0.0.1; } }` [PR884602](#): This issue has been resolved.
- RPD may crash on the new master Routing Engine after Routing Engine switchover. The issue is NSR related, and it happens due to the bad BGP route data structure on the backup Routing Engine. [PR885305](#): This issue has been resolved.
- When used JUNOScript to run command 'get-pim-neighbors-information instance=' (with NULL instance name), which triggered core dump even though there are no routing-instances with pim enabled. It won't trigger core file if JUNOScript command includes any instance name. [PR887070](#): This issue has been resolved.
- In a scenario with graceful restart (GR) enabled for BGP between Cisco platform and Juniper Networks platform, Junos OS is helper (default) and Cisco being restarting router, when Cisco restarts BGP process, Juniper deletes all BGP routes due to doesn't receive End Of RIB (EOR) markers for all configured NLRI from Cisco. [PR890737](#): This issue has been resolved.
- The downstream PE router's RPF_neighbor(S) on the MDT reverts back to mRIB.next_hop(S) rather than the Assert(S,G)Winner when their PPT expires. Bug identified in the code and is fixed. [PR896898](#): This issue has been resolved.

Services Applications

- Memory leak in key management daemon (kmd) causes some IPsec VPN tunnels to be dropped and don't get re-negotiated for over 10 minutes. Before issue happens, the following logs could be observed: /kernel: Process (1466,kmd) attempted to exceed RLIMIT_DATA: attempted 131080 KB Max 131072 KB /kernel: Process (1466,kmd) has exceeded 85% of RLIMIT_DATA: used 132008 KB Max 131072 KB [PR814156](#): This issue has been resolved.
- In L2TP subscriber management environment, after issuing CLI command "commit full", jl2tpd process (l2tp daemon) deletes all tunnel profiles and brings down all L2TP subscribers. Even though there are no configuration changes. [PR834504](#): This issue has been resolved.
- MIB module in file "mib-jnx-sp.txt" contains a coding error, which may lead to a loop. [PR866166](#): This issue has been resolved.
- If RSP1 and RSP10 interfaces are configured on the same box, issuing the "request interface switchover rs1" or "request interface revert rsp1" causes both RSP1 and RSP10 to switchover or revert. [PR877569](#): This issue has been resolved.
- In a CGNAT environment when sp interfaces, which are underlying rsp interface, are present in the configuration, sp interfaces service-options may incorrectly overwrite rsp interfaces service-options and syslog stopped working and inactivity-timeout values were reset to the default values. [PR881792](#): This issue has been resolved.
- AAPID list configuration not copied to Backup Routing Engine // 12.3R2.5 [PR885833](#): This issue has been resolved.
- The jl2tpd process generates a core file as follows:
"././src/bsd/lib/libc/stdlib/abort.c:69." [PR887662](#): This issue has been resolved.
- SIP ALG - Service PIC might crash when SIP flows are cleared. [PR890193](#): This issue has been resolved.
- Output interface shown as 'Unknown' under show services accounting flow-detail.issue has been analyzed RCA;-At the time when a flow is created in PIC memory, if the route to the destination IP(in the flow) is not known, we set a flag indicating that there is no route to Destination IP in the flow structure. When the flows are queried using "show service accounting flow-detail", picinfo daemon inspects this flag for each flow and prints the Output interface as "Unknown" if this flag is set. Now, after route record for that flow is downloaded to the Service PIC, the flow structure is updated to reflect the corresponding output interface, but, the above flag is NOT UNSET. So, picinfo daemon continues to print the output interface as "unknown" whenever "show services accounting flow-detail" is executed. [PR890324](#): This issue has been resolved.
- L2TP session on MS-PIC may fail and following error is observed
"L2TPD_RADIUS_SERVER_NOT_FOUND" after a test access profile <ppp-profile> is issued. [PR898872](#): This issue has been resolved.
- When the 'learn-sip-register' knob is enabled for the SIP ALG (it is by default), for a SIP request in slow path implicitly denied by the firewall or NAT rules, a look up is done to see if the SIP request has a target that corresponds to any current registration state, in which case the corresponding reverse flows get created. While service PIC creating

the corresponding reverse flows, an internal error may occur, causing service PIC to crash and generate a core file. [PR899195](#): This issue has been resolved.

- In a L2TP scenario, after performing an SNMP walk of "jnxL2tpTunnel" or "jnxL2tpSession" MIBs, the SNMP reply message fails to be written because write buffer is exceeding MTU, causing Routing Engine CPU spikes to 100%. [PR905218](#): This issue has been resolved.

Subscriber Access Management

- In DHCP/PPPoE subscriber management environment, after terminating subscribers, authd process might crash and generate a core file due to an invalid pointer being used. [PR821639](#): This issue has been resolved.
- In situation when CoA message includes both LI attributes and CoA attributes, authd process fails to respond to CoA. [PR821876](#): This issue has been resolved.
- DTCP - First 127 triggers are applied. [PR873013](#): This issue has been resolved.
- PPPoE was not supported for the 802.1ad 0x88a8 TPID on the outer tags of dual-tagged VLANs:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
set tag-protocol-id [0x88a8]
```

[PR874603](#): This issue has been resolved.

- The authdlib logout/terminate release notify request might experience a processing loop. [PR888281](#): This issue has been resolved.

User Interface and Configuration

- In an aggressive provisioning scenario using scripts or automated tools, we recommend that you do not use rollback immediately after a successful commit. [PR874677](#): This issue has been resolved.

VPNs

- In this release Ngen-MVPN does not support NSR. But the commit check when Ngen-MVPN and NSR is configured does not fail. In previous releases this commit would fail. The commit check not failing for this configuration is expected to be fixed in release 12.3 R4. In Release 12.3 R3 config with NSR and Ngen-MVPN configuration should not be committed. Doing this commit can lead to routing application crashes (like PR 864439) as it is an unsupported feature. [PR827519](#): This issue has been resolved.
- Wrong data type for MIB object "mplsL3VpnVrfRteXCPointer". [PR866259](#): This issue has been resolved.
- If SNMP "get" tries to retrieve local and direct routes from mplsL3VpnVrfRteTable, they are not found. SNMP walk does walk the local and direct routes. [PR874365](#): This issue has been resolved.
- In VPLS scenario, if CE facing interface is aggregated Ethernet with multiple member ports (more than two members), BUM (broadcast, unicast unknown, and multicast) traffic from MPLS core will be replicated on all child link of aggregated Ethernet

interface, and BUM from CE will be replicated at sending out from MPLS core facing interface. The problem is specific to M10i and M7i routers running with I chip based CFEB. [PR880422](#): This issue has been resolved.

- RPD might experience software exception during clear pim join on routing-instance. Typically seen in scenario where PIM load balancing is implemented over eibgp sessions. [PR891586](#): This issue has been resolved.

Resolved Issues in Release 12.3R3

Class of Service

- A few memory leaks have been fixed in the class of service process. [PR811613](#): This issue has been resolved.
- This cosmetic issue is specific of 3D linecards, based on MX Series routers with MPCs/MICs. In these cards, the logical interfaces with family mpls do not have any EXP rewrite rule applied by default. In other words, EXP value is copied from the previous codepoints: for example, from IP Precedence in IPv4->MPLS next hops. However, the command "show class-of-service interface" still shows the exp-default rule as if it was applied (in fact, it isn't): user@router> show class-of-service interface ge-2/3/1.204 | match rewrite Rewrite exp-default exp (mpls-any) 33. [PR824791](#): This issue has been resolved.
- When 'scheduler-map-chassis derived' configuration is used under class-of-service, interface related configuration changes can lead to cosd process crash. [PR863734](#): This issue has been resolved.

Forwarding and Sampling

- There is always a chance to see this issue if any daemon adds a blob size which comes closer to 65520 (after IDR encoding). [PR700635](#): This issue has been resolved.
- Memory leak could happen to pfd, dcd, cosd, cfmd and dfcd processes if user frequently and repeatedly executes "show interface extensive" command from multiple telnet sessions under the conditions below. 1. Set screen-length value to small value. Screen length can be changed by the command "set cli screen-length <n> ". 2. User enters "show interface extensive" command simultaneously from multiple telnet sessions. And cancel the output of the command with "q" as soon as "---(more)---" shows up at the end of the output. [PR843145](#): This issue has been resolved.
- MPLS forwarding table filter (ftf) not getting linked in JTREE after router or FPC reboot. [PR851599](#): This issue has been resolved.
- When committing a firewall filter with a "then decapsulate" action, the router might throw the following errors Feb 19 11:20:59 user@host dfwd[45123]: DFWD_FW_PGM_READ_ERR: Read of segment 0/0 in filter 2 failed: Unknown error: 0 Feb 19 11:21:01 user@host dfwd[45123]: DFWD_CONFIG_WRITE_FAILED: Failed to write firewall filter configuration for FILTER idx=2 owned by CLI. Error: Message too long This issue happens on an MX Series router that has at least one I-chip board (MX with DPC).

This happens because the firewall daemon fails to properly update the Packet Forwarding Engine firewall configuration. [PR857708](#): This issue has been resolved.

- In T4000 platforms with ES-FPC, for IPv6 firewall filters with match conditions on address prefixes longer than 64 bits, in some corner cases, the filter might not be correctly evaluated and packet loss might occur. [PR879829](#): This issue has been resolved.

General Routing

- Prior to this change, the L2TP sessions with cos/ firewall attachments fail to come up when the L2TP Access Concentrator (LAC) is reachable over a unicast nexthop. [PR660208](#): This issue has been resolved.
- The 'RL-dropped' lines of "> show interfaces queue" will be missing when the PIC is bounced. [PR749283](#): This issue has been resolved.
- If dynamic profile versioning is configured and In-service software upgrade (ISSU) is performed from 11.4x27.35 (GA build) to 11.4x27.38(Nov-2012), exiting subscribers might either lose traffic or might get terminated. [PR817018](#): This issue has been resolved.
- VPLS traffic gets flooded back over the ingress interface on the local PE as the split-horizon gets disabled upon interface flap. [PR818926](#): This issue has been resolved.
- In a race condition where multiple interrupts are asserted, timer tick might not get well handled and remain asserted. This caused panic and core. [PR828496](#): This issue has been resolved.
- RPD on the backup Routing Engine might crash when it receives a malformed message from the master. This can occur at high scale with nonstop active routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare and backup RPD will restart and the system will recover without intervention. [PR830057](#): This issue has been resolved.
- Since the AC Power System on the MX2020 is a N+N feed redundant and N+1 PSM redundant, there are two separate input stages per PSM, each connected to one of the two different/redundant feeds. However, only one stage is active at a time. This means, the other input stage (unused input stage) might be bad and system will not know about it till it tries to switch to it in case of a feed failure. This is a pretty bad corner case and needs to be addressed. The way to work around this problem is by testing both stages when the power supply is first powered on. This test is done by the system software and an alarm is raised if any feed failure is detected. [PR832434](#): This issue has been resolved.
- Memory leak is observed in authd process, with a churn of 1000 subscribers over 3 min period. [PR835204](#): This issue has been resolved.
- It is possible that RPD's higher priority tasks (HPTs) are scheduled to run such that lower priority tasks (LPTs) might not be able to complete until HPTs are completed. [PR836197](#): This issue has been resolved.
- Enabling PIM - Bidir feature (possibly pim rp with 224.0.0.0/4 group) and rpd restart triggers this issue is hit during regression test for PIM bidir. 2) HW type of chassis/linecard/Routing Engine. If it affects all, just say ?all?. =>all. 3) Suspected

software feature combination. (If customer turns on feature X along with Y, they may hit, etc) =>PIM - Bidir feature (rp configured) and rpd restart is causing the issue. 4) Describe if any behavior/ change to existing function =>None. [PR836629](#): This issue has been resolved.

- After a Routing Engine switchover with graceful Routing Engine switchover enabled, and then deactivate and activate a routing-instance, 4xOC48 IQE PIC might reboot unexpectedly. This is caused by a problem in channel allocation for the 4xOC48 PIC logical interfaces in kernel. [PR841822](#): This issue has been resolved.
- LMNR Chipset type FPC generates a core file with copy-plp-all enabled when adding link to existing AE interface, which is part of downstream interface list of a multicast route. [PR842046](#): This issue has been resolved.
- When MX Series router running with DPC is upgraded by ISSU, some of interface might show incorrect input packet/byte count. And the incorrect count is also seen to the related interface MIB. The value will be a large number. Physical interface: xe-3/1/0, Enabled, Physical link is Up Interface index: 138, SNMP ifIndex: 5449, Generation: 141 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback: Local, Source filtering: Disabled, Flow control: Enabled Device flags : Present Running Loop-Detected Interface flags: SNMP-Traps Internal: 0x4000 Link flags : None CoS queues : 8 supported, 8 maximum usable queues Hold-times : Up 0 ms, Down 0 ms Current address: 00:24:dc:9c:7c:30, Hardware address: 00:24:dc:9c:7c:30 Last flapped : 2013-01-13 14:36:25 JST (02:07:52 ago) Statistics last cleared: Never Traffic statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps ~snip~ Logical interface xe-3/1/0.0 (Index 196614) (SNMP ifIndex 5450) (Generation 140) Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Traffic statistics: Input bytes : 3867797326912475 Output bytes : 0 Input packets: 15108583308733 Output packets: 0 0 Local statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Transit statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps Protocol inet, MTU: 1500, Generation: 160, Route table: 0 Flags: Sendbcst-pkt-to-re Addresses, Flags: Is-Preferred Is-Primary Destination: 10.3.1/24, Local: 10.3.1.1, Broadcast: 10.3.1.255, Generation: 141 Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0 Policer: Input: __default_arp_policer__ gladiolus:Desktop\$ grep .5449 mib_value_after_issu.txt ifName.5449 = xe-3/1/0 ifInMulticastPkts.5449 = 0 ifInBroadcastPkts.5449 = 0 ifOutMulticastPkts.5449 = 0 ifOutBroadcastPkts.5449 = 0 ifHCInOctets.5449 = 3867797326912475 ifHCInUcastPkts.5449 = 0 ifHCInMulticastPkts.5449 = 0 ifHCInBroadcastPkts.5449 = 0 ifHCOctets.5449 = 0 ifHCOUcastPkts.5449 = 0 ifHCOMulticastPkts.5449 = 0 ifHCOBroadcastPkts.5449 = 0 gladiolus:Desktop\$ grep .5450 mib_value_after_issu.txt ifName.5450 = xe-3/1/0.0 ifInMulticastPkts.5450 = 0 ifInBroadcastPkts.5450 = 0 ifOutMulticastPkts.5450 = 0 ifOutBroadcastPkts.5450 = 0 ifHCInOctets.5450 = 3867797326912475 ifHCInUcastPkts.5450 = 15108583308733 ifHCInMulticastPkts.5450 = 0 ifHCInBroadcastPkts.5450 = 0 ifHCOctets.5450 = 0 ifHCOUcastPkts.5450 = 0 ifHCOMulticastPkts.5450 = 0 ifHCOBroadcastPkts.5450 = 0. [PR847106](#): This issue has been resolved.
- It is possible for RPD core when the following conditions are met: - VRF with multipath knob configured - static routes with next-hops which are indirect type and needs further

resolution - the numerically lowest (smallest IP) next-hop of indirect type becomes unreachable RPD core is NOT triggered in either of the following scenarios: - no multipath under VRF - if there is no static route entry - static route whose next-hops are indirect type requiring further resolution multipath under VRF is supported only for BGP configurations. multipath in other conditions are not supported, and a bug in this detection phase is fixed in this PR. [PR847214](#): This issue has been resolved.

- mlfr/mlppp interface are not reachable after restart FPC (primary MSPIC) followed by deactivate and activate R.I or GRES followed by deactivate and activate R.I. This is because link FPC does not have the interfaces programmed towards the bundle. [PR847278](#): This issue has been resolved.
- In certain graceful Routing Engine switchover scenarios, with IPv6 address configured on at least two interfaces, Solicited node multicast addresses (SNMA) and link local addresses with same prefix might be created on the two interfaces. There is a possibility that there could be inconsistency in the Next Hop database between Master and Backup Routing Engines. When the Backup becomes Master in these scenarios, it'll try to program the Packet Forwarding Engines with the bad Next Hop data. This might cause undesired forwarding behavior on the Packet Forwarding Engines. [PR850625](#): This issue has been resolved.
- Ptsf failed to append policy with multi-rules since 'msg over size limit'. [PR852224](#): This issue has been resolved.
- FPC or PIC connects to the Routing Engine Kernel for the first time when it comes up or reconnects during connection trip. After the connection is established with the Routing Engine, if FPC/PIC does not respond kernel for 300 seconds, a timer is triggered to disconnect the Routing Engine from FPC/PIC. In a particular race condition between kernel processing received data on the connection and the fired timer trying to close the connection, kernel crashes and creates a core file. FPC/PIC's slow response may be attributed to high traffic or a faulty hardware. Before kernel crash, the following logs could be seen: fpc3 LCHIP(3): 1 new Lin SIF ins eoqe errors fpc3 LIN(3): PIC HSR is not OK, LCHIP(3) <- PIC 3 HSR 1. [PR853296](#): This issue has been resolved.
- If routing-instance is popping the mpls label through vt tunnel interface and the egress interface MTU of the vrf needs fragmentation and the dont-fragment bit is set in the ipv4 header, the egress vrf interface might stop forwarding traffic. The following syslog message will be reported fpc4 LCHIP(3): 1 new errors in LSIF To recover from this condition you can either bring the interface down via disable knob or deactivate/activate the interface from the configuration. The following platforms are exposed to this condition: M320 (excluding E3 FPCs), T/TX systems (excluding ES FPCs and FPC Type 5). [PR854806](#): This issue has been resolved.
- In the T4000 Type 5 FPC platform, aperture management can lead to a collision between the sched tick timer and asic driver interrupt handlers, which will result in FPC crashes. [PR857167](#): This issue has been resolved.
- In a virtual chassis environment in the event power is loss on the Master virtual chassis the standby chassis has potential to experience slot resets during transition period. [PR859717](#): This issue has been resolved.
- BOOTP request packets might get dropped because of the DDOS protection feature in old MX Series router with MPCs/MICs. In this case, the bootp packets is coming with

1 byte option. So the length of bootp become 241 which is larger than 240. Then Packet Forwarding Engine will identify it not as BOOTP as per the current DDOS algorithm, and tries to parse it as DHCP. Since the packet lacks the options fields which need for DHCP, then `pfe_nhdb_dhcpv4_msg_type()` mark it as DHCPNOMSGTYPE. [PR862206](#): This issue has been resolved.

- When a prefix next-hop address resolution requires a recursive lookup, the next-hop might not be updated correctly after an egress interface is disabled. [PR862989](#): This issue has been resolved.
- When using BGP Flow Spec with rate-limit option, even though the value is in Bytes/second, the value being programmed is in bits/second. [PR864496](#): This issue has been resolved.
- Output of **show subscribers physical-interface aex** displays multiple AE links. [PR864555](#): This issue has been resolved.
- "set chassis fru-poweron-sequence " configuration is not supported for T4000 platform in Junos OS Release 12.3R2. [PR868035](#): This issue has been resolved.
- On T Series platforms with ES-FPC equipped, while adding and deleting source-class usage (scu) or unicast Reverse path forwarding (uRPF) configuration, Jtree memory leak and the following error messages could be observed: `fpc0 nh_jtree_fe_posthandler: RNH_TABLE 1 missing ext rnh`. [PR869651](#): This issue has been resolved.
- In subscriber management environment, with scaling subscribers login (110K DHCP and 20K PPPoE), after restarting one of the line cards which has subscribers, autoconf process might crash and generate a core file due to memory corruption or memory double free. Only 11.4X27.45 is affected by this issue. [PR870661](#): This issue has been resolved.
- In a scenario with scale Routing Instances (RIs) configured, after deactivating/activating two RIs, routing protocol process (rpd) might try to free a specific pointer pointing to an incorrect structure that is actively in use. Then rpd process crashes and generate core files. [PR870683](#): This issue has been resolved.
- When configuration stanza: `[protocols router-advertisement]` starts as: `## ## inactive: protocols router-advertisement ## interface ge-0/0/1.1 { virtual-router-only; }` Then perform the following actions: Step 1 - activate protocols router-advertisement Step 2 - deactivate protocols router-advertisement interface ge-0/0/1.1 Step 3 - set protocols router-advertisement interface ge-0/0/1.2 After issuing "commit check", there are no problems. But after issuing "commit", routing protocol process (rpd) crashes and generates core with following logs: `rpd[1422]: RPD_RA_CFG_UNKNOWN_ACTION: Unknown configuration action 3 received`. [PR871359](#): This issue has been resolved.
- Adding a routing-instance with "/" in its name will cause the router not to boot properly if logical-systems were previously configured. [PR871392](#): This issue has been resolved.
- Under high scale, expiration of a Kernelside reconnect timer would cause it to send a non-serviceable message to the Packet Forwarding Engine (asking the line cards to restart and resync since reconnect failed). Since there is no ack- to this Kernel message, Kernel thought it sent the message and untoggles the GRES flag. The Packet Forwarding Engine wasn't expecting anything so it continued along. The EFFECT: The system is permanently not ready for GRES... CLI GRES check will always report: `[cmd] request`

chassis routing-engine master switch check Apr 14 19:03:13 [INFO] warning: Standby Routing Engine is not ready for graceful switchover. [PR873679](#): This issue has been resolved.

- On MX Series routers with DPC (I-Chip based) type FPCs running a 11.4 (or newer) Junos OS release, disabling uRPF on a logical interface might result in another logical interface on the router to drop all incoming packets. This problem happens only when the following conditions are met concurrently: a) 2 different logical interfaces share the same lookup index b) both logical interfaces have uRPF enabled c) these 2 different logical interfaces belong to 2 different FPCs d) at least one of the logical interfaces belongs to a DPC (ICHIP based) type FPC The lookup index is calculated by taking the lower 16 bits of the logical interface index (also called the IFL index). In other words $\text{lookup index} = \text{IFL index} \text{ MOD } 65536$. It is normal, valid and expected to have logical interfaces which share the same lookup index. The problem described in this PR is *_not_* the fact that the lookup indexes are the same. Here is an example of 2 different logical interfaces on 2 different FPCs which share the same lookup index: Interface ge-0/1/0.945 has an IFL index of 1774 and a lookup index 1774: `user@router-re1> show interfaces ge-0/1/0.945` Logical interface ge-0/1/0.945 (Index 1774) (SNMP ifIndex 1635) ^^^^^^^^^^ Flags: Device-Down SNMP-Traps 0x4000 VLAN-Tag [0x8100.945] Encapsulation: ENET2 Input packets : 0 Output packets: 0 Protocol inet, MTU: 4462 Flags: Sendbcst-pkt-to-re, uRPF, uRPF-loose Addresses, Flags: Dest-route-down Is-Preferred Is-Primary Destination: 52.3.168.216/29, Local: 52.3.168.217, Broadcast: 52.3.168.223 Protocol multiservice, MTU: Unlimited And interface xe-2/2/0.0 has an IFL index of 198382 and a lookup index of $198382 \text{ MOD } 65536 = 1774$: `user@router-re1> show interfaces xe-2/2/0.0` Logical interface xe-2/2/0.0 (Index 198382) (SNMP ifIndex 698) ^^^^^^^^^^^^^^ Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Input packets : 381 Output packets: 376 Protocol inet, MTU: 1500 Flags: Sendbcst-pkt-to-re, uRPF, uRPF-loose Addresses, Flags: Is-Preferred Is-Primary Destination: 155.154.153.0/30, Local: 155.154.153.1, Broadcast: 155.154.153.3 Protocol multiservice, MTU: Unlimited In the example above if uRPF is disabled on ge-0/1/0.945 then xe-2/2/0.0 will start dropping all incoming packets due to RPF failure. When this condition occurs the only way to recover is to disable, commit and re-enable uRPF on the broken interface. When this is done the following error messages are generated: Apr 15 16:02:53 router-re1 fpc2 rt_iff_generic_topo_handler: jtree error Not found for disconnect on iff-post-src Apr 15 16:02:54 router-re1 fpc2 RT(rt_rpf_jtree_drt_remove_ifl): Unable to remove logical interface 198382 from drt(4) Apr 15 16:02:54 router-re1 fpc2 RT(rt_rpf_jtree_drt_remove_ifl): Unable to remove logical interface 198382 from loose(7). [PR873709](#): This issue has been resolved.
- The default setting for the sysctl "net.pfe.relayg_merge_enabled" is 0 (off), this results in a support limit of 16 line-cards within the VC. Even with the group merge disabled, line-cards may have been grouped at system start-up only presenting an issue after they restart. [PR874791](#): This issue has been resolved.

High Availability and Resiliency

- The backup Routing Engine sends Arp 128.0.0.6 to the Packet Forwarding Engine, then they are counted as "unknown" on show pfe statistics traffic. [PR830661](#): This issue has been resolved.

- This issue is seen on IQ2 PICs during ISSU on TX platform. When upgrading to 12.3R2 from releases prior to 12.3R2 through ISSU, IQ2 PICs will report error. This error is due to IQ2 PICs not able to download image during ISSU. [PR855661](#): This issue has been resolved.

Infrastructure

- The root cause of the problem was IFADDR change in VRRP context was not replicated to GRES backup. [PR790485](#): This issue has been resolved.
- Kernel fails to generate ICMP ttl expired when IP packet len is a multiple of 256. [PR829567](#): This issue has been resolved.
- Aggregate Bundle interface with IPV6 Interface stuck in Tentative state. Trigger was deactivation/activation of ae-interface. [PR844177](#): This issue has been resolved.
- Delay in bringing online an FPC after it is inserted into the chassis. [PR853304](#): This issue has been resolved.
- TCP is mistakenly enabling re-transmit timer for pure ACK's which is causing the FPC to reboot. [PR858489](#): This issue has been resolved.
- With nonstop active routing (NSR) enabled, while performing graceful Routing Engine switchover, the Junos OS fails to restore BGP peers' TCP connections on the new master Routing Engine's replicated socket due to it is not able to find the BGP peer address's route, causing BGP peers to flap with following logs: /kernel: jsr_sdr_merge: PSRM merge failed 65 rpd[xx]: RPD_BGP_NEIGHBOR_STATE_CHANGED: BGP peer a.b.c.d (Internal AS X) changed state from Established to Idle (event TcpSocketReplicationError). [PR862796](#): This issue has been resolved.
- When a sonet interface with PPP encapsulation is used as forwarding next hop for the IPv6 remote router loopback address on IPv6 BGP sessions, if the sonet link is down, the IPv6 BGP session might flap at same time although there is valid route via other interface. [PR863462](#): This issue has been resolved.
- After enabling firewall filter of IPv6 on Aggregated Ethernet (AE) interface to block Micro BFD Packets (Dst Port 6784), kernel crashes continually on Master and the backup Routing Engine due to double free of memory. [PR864112](#): This issue has been resolved.
- IPv6 Neighbor discovery (ND) failed after multiple GRES. Nexthop getting stuck in hold state forever. We also see that the neighbor state is in NO_STATE and it is on ND timer queue. In this condition, on ND timer expiry it never sends neighbor solicitation (NS) out and it never transitions to known ND states. [PR864133](#): This issue has been resolved.

Interfaces and Chassis

- When MAC address filters are configured on an AE, MAC filters might not be programmed on the child link of the AE if and only if the following sequence of events occur: AE is disabled via a configuration change, a graceful Routing Engine switchover occurred and AE is subsequently enabled on the new master Routing Engine. [PR561106](#): This issue has been resolved.
- There can be a mismatch between the ifIndex value on IF-MIB-ifName and the ifIndex value on SONET-APS-MIB-apsMapGroupName and apsMapEntry. [PR771877](#): This issue has been resolved.
- This issue is specific to the M120 hardware since there are two independent FRU's from where the PIC needs to be detached/attached. This IPC messages goes out-of-order due to the additional control-plane messages related to routing-change as a result of PIC restart which happens in this case due to the buffer configuration change. When PIC needs to be detached and at the same time there are still a lot of protocol information which should be process as well, the detached messages will NOT be able to be delivered in time. After PIC restarts it request to be attached again but obviously this action failed because from other FRU's perspective the PIC has NOT been detached at all. [PR773081](#): This issue has been resolved.
- With LSQ interface, the MLPPP fragments cannot use the egress queue 4 to 7 on the MLPPP member links. There is no workaround. [PR805307](#): This issue has been resolved.
- Incorrect Detection timestamp in **show chassis fabric reachability**. [PR811846](#): This issue has been resolved.
- Faulty SCG causes continuous interrupts to HCFPC making its CPU Utilization 100% and unusable for any service. As a fix the monitoring mode for the SCG is changed to polling status of SCG device rather than interrupts based awake and monitoring system. [PR827489](#): This issue has been resolved.
- In Integrated Routing and Bridging (IRB) interface over Aggregate Ethernet (AE) interface scenario, if there is an MAC Move event or an L2 IFL change event with IRB, the Junos OS will remove the IRB nexthops on the backup Routing Engine and Packet Forwarding Engines first and then remove it from the master Routing Engine. During this phase, if an logical interface change event of the underlying AE interface occurs, the Junos OS might try to access a stale pointer which was freed already and cause memory corruption. In some conditions, the memory corruption occurs in kernel, hence cause kernel crash and generate a core file. [PR829093](#): This issue has been resolved.
- The kernel "devbuf" memory leaks when fxp0 interface is in down state (admin up). [PR829521](#): This issue has been resolved.
- A request (like snmp query) for collecting input ipv6 stats of ae logical interface on abc chipset is not working properly. [PR831811](#): This issue has been resolved.
- Removing IP address on ATM interface after adding another IP address from the common subnet can lead to a race condition. New IP address configured on the interface still referring to shared broadcast-nexthop. Then when TCP/IP access this broadcast-nexthop kernel panic might happen. [PR833015](#): This issue has been resolved.

- When packet has to be forwarded over NH topology unilist->indirect-indexed and when the packet size is greater than egress interface MTU w/ DF set, then we might log the following message and not send the message back to source indicating "frag needed and DF set". fpc0 NH: Can not find logical interface for nh 1048590 fpc0 NH(nh_get_mtu_iff) : get unilist mtu failed. [PR844987](#): This issue has been resolved.
- In a scenario of PPP sessions over L2TP tunnels, on L2TP network server (LNS), if authentication is none or if authentication is enabled but radius does not return any Framed-IP-Address/Framed-Pool, jpppd process is not setting the IP address key of subscriber to "255.255.255.254" thereby resulting in address allocation failure in authd process. Then the L2TP tunnels can not be established, hence subscribers can not login. When issue happens, the following logs of authd process could be seen: client type jpppd client type REQUESTING: OldStyle 0 OldStyleFilled 0 hint null network null client pool name. [PR849191](#): This issue has been resolved.
- Whenever tunnel interface -pe/-pd got created using the MS-DPC instead of the MPC, it was not able to process register messages. Because of MPC and MS-DPC have different multicast architecture and they are incompatible if chassis is configured in "enhanced-ip" mode, this issue will be seen. Necessary changes have been made to code so that these interfaces will not be created on MS-DPC. [PR853995](#): This issue has been resolved.
- SDG : After rebooting both the Routing Engines together, the FPCs and MS-DPCs might come online, go offline (with "Chassis connection dropped" and "Chassis Manager terminated" error messages) and come back online again automatically. This issue is seen only when both Routing Engines are rebooting at once. There is exactly one additional reboot of the FPCs when this happens, and the FPCs come back up online, and system stabilized by itself within 2 to 3 additional minutes. [PR854519](#): This issue has been resolved.
- In certain topology set up such as multiple trunks are used on a PE with P and the CE-PE interface is MLFR, and enhanced-ip and MS-DPC route-localization are configured, if the active trunk FPC is offlined, VRF traffic from PE towards CE using the mlfr interface might get blackholed. [PR854623](#): This issue has been resolved.
- Multilink Frame-relay (MLFR) stuck in ready state after restarting FPC and then graceful Routing Engine switchover (some of the MLFR bundles will show "ready" although the interfaces are in up/up state which causes data loss). [PR857648](#): This issue has been resolved.
- The backup Routing Engine might log the following often in chassisd: Feb 17 12:40:01 CB:1 need not to sync information Feb 17 12:40:21 CB:1 need not to sync information Feb 17 12:40:41 CB:1 need not to sync information Feb 17 12:41:01 CB:1 need not to sync information. This is a harmless message that can be ignored. [PR857698](#): This issue has been resolved.
- In PPPoE subscriber management environment, PPPoE daemon might crash and generate a core file in following two scenarios: 1 - Firewall Filter/Policer is not configured on Broadband Remote Access Server (BRAS) side, and AAA pushes the filter name in "Ingress Policy Name/Egress Policy Name" which will expire the lockout timer waiting to create required dynamic interface, and eventually causes pppoe process crash. 2 - When IPv6 only capable modem is trying to connect and the configuration does not

contain IPv6 dynamic configuration; i.e. under PPPoE dynamic profile/family inet6 stanza; PPPoE dynamic profile/protocols/router-advertisement, this will again expires lockout timer waiting for dynamic interface creation, which crashes pppoe process.

[PR859000](#): This issue has been resolved.

- Interface hold-time-down is not working properly for PIC type 10x10GE(LAN/WAN) SFPP. [PR859102](#): This issue has been resolved.
- Enables maximum-links CLI knob which specifies the maximum number of links in an aggregated ethernet bundle. This can take a value of 16, 32 or 64 depending on the platform. [PR860152](#): This issue has been resolved.
- ISSU does not support VRRP. [PR862052](#): This issue has been resolved.
- MX Series router is sending RADIUS Acct-Start, in spite of the fact that IPCP/IPv6CP is not established. [PR867084](#): This issue has been resolved.
- The chassisd crashes when enable route-localization with MPC2E. [PR872500](#): This issue has been resolved.
- If both "startup-silent-period" and "delegate-processing" are configured under protocols vrrp, both vrrp routers keep backup-backup state until "startup-silent-period" expires when trying to revert. [PR873488](#): This issue has been resolved.

Layer 2 Ethernet Services

- In the rare case of one GRES is performed after another GRES, without logging out and logging in subscribers, some ipv4 access routes will not be reinstalled. This will result in traffic loss for the affected dhcp v4 subscribers. [PR808932](#): This issue has been resolved.
- DHCPv6 fails for clients using DUID type 2 (Vendor-assigned unique ID), the software was using the DUID to extract MAC address information. This behavior is fixed and tested. [PR838404](#): This issue has been resolved.
- MXVC-DHCP bindings stuck in a "RELEASE(RELAY_STATE_WAIT_AUTH_REQ_RELEASE" state. [PR850187](#): This issue has been resolved.
- For MXVC, the derivation of the dhcp server-id has changed from using hardware serial number to lacp mac addr. The reason is that the lacp mac address is guaranteed to be reflected across the chassis so upon GRES, the same dhcp server id can be built. However, upon ISSU, the old software will derive server-id from hardware serial number and the new software will derive it from lacp mac address and they will not match. After the ISSU, DHCP packets may be dropped by a dhcp server because the serverid in the client packet will not match that of the server. This will only happen when transition to the new method of building the serverid. Once that has happened, all future ISSU should work as before. [PR853329](#): This issue has been resolved.
- In DHCP subscriber management environment, while DHCP subscribers login, in rare conditions, system calls of these subscribers fail, due to only on success does system free the memory, resulting in a memory leak for the jdncpd process. If memory usage of jdncpd process goes to its limit, no new DHCP subscribers can login. When issue happens, high weighted CPU usage of jdncpd process and following logs could be observed. /kernel: %KERN-5: Process (31403,jdncpd) has exceeded 85% of

RLIMIT_DATA: used 2825132 KB Max 3145728 KB jdhcpcd:
%USER-3-DH_SVC_RTsock_FAILURE: Error with rtsock: rtslib: ERROR Failed to allocate new block of size 16384 jdhcpcd: %USER-3-DH_SVC_RTsock_FAILURE: Error with rtsock: rtslib: ERROR Failed to allocate new block of size 16384 jdhcpcd:
%USER-3-DH_SVC_RTsock_FAILURE: Error with rtsock: rtslib: ERROR Allocation Failure for (16384) bytes authd[1822]: %DAEMON-3:
../../../../src/junos/usr.sbin/authd/plugin/radius/authd_plugin_radius_module.cc:1090
Failed to get SDB snapshot for session-id:3549005. [PR856024](#): This issue has been resolved.

- DHCP relay functionality over IRB performing dhcp v4 relay functionality, and configured with both inet and inet6 address families. The removal of the IPv6 address family configuration from the IRB can cause the IPv4 dhcp relay functionality on that IRB to break. This happens regardless of whether the 'family inet6' is configured directly under the IRB or applied through a 'apply-group' configuration. In versions that do not have the fix for this PR, the workarounds to get the dhcp relay functionality working again over the IRB are *either* of the following:
 - Deactivate/activate the IRB configuration.
 - Restart the dhcp process using the **restart dhcp-service** command.



NOTE: This workaround has to be applied everytime any configuration change (as explained in the trigger) is applied that could potentially get the dhcp-relay functionality to break.

[PR870543](#): This issue has been resolved.

- When IPv6 is configured on integrated routing & bridging (IRB) interfaces which has AE interfaces as child links, after GRES was enabled and one child link failure or removal, the kernel get crashed. [PR878470](#): This issue has been resolved.

MPLS

- In a RSVP P2MP crossover/pass-through scenario, more than one sub-LSP can use the same PHOP and NHOP. If link protection is enabled in the above mentioned scenario, when a 'primary link up' event immediately followed by a Path Tear message, disassociation of the routes/nexthops are sequential in nature. When the routes/nexthops dissociation is in progress if a sub-LSP receives a path tear/PSB delete will lead to this core. [PR739375](#): This issue has been resolved.
- The cleanup procedures may leave transient inconsistent references when the interface address of an MPLS enabled GRE or IPIP tunnel is being deleted or the action taken implies an internal reconfiguration of the interface address (for example MTU change). During this period, if these references are being reused by a particular task, the kernel may report an invalid memory access and restart. [PR844790](#): This issue has been resolved.
- The routing protocol process (rpd) might leak memory when there are MPLS LSP changes, the memory leak could eventually cause rpd process to crash. [PR847354](#): This issue has been resolved.

- There appears an unsupported feature warning missing for mLDp+NSR while doing ISSU. [PR849178](#): This issue has been resolved.
- RPD generates a core file on the backup Routing Engine with `rsvp_mirror_telink_attempt_resolve`. [PR859602](#): This issue has been resolved.
- ASBR might not rewrite EXP correctly for egress MPLS packets on the Inter-AS link for the eBGP-LU LSP if the eBGP session is a multihop BGP session. [PR868945](#): This issue has been resolved.
- In a scenario when scaled MPLS tag labels exist, configure and delete ospf overload configuration. After committing the configuration changes, routing protocol process (rpd) might crash and generate a core file due to system tries to delete an already freed MPLS tag label Element. [PR878443](#): This issue has been resolved.

Network Management and Monitoring

- When snmp unknown PDUs are received, the appropriate counter in (show snmp statistics) is not incremented. [PR865121](#): This issue has been resolved.
- Polling an snmp oid that was excluded from the snmp view in configuration might trigger an increase in CPU load related to SNMP and RPD processes. [PR866541](#): This issue has been resolved.

Platform and Infrastructure

- XML tags for get-software-information output missing some elements of the new Junos OS service release naming convention. [PR783653](#): This issue has been resolved.
- Due to a bug in logical interface localization, a DPC restart/offline may cause a removal of legitimate CCC routes on other DPC's. This can also be triggered by removal of an unrelated family CCC logical unit. [PR835216](#): This issue has been resolved.
- When a junoscript get-configuration RPC query, by default the query is done on candidate DB, a MGD process is spawned to handle this request. Now at the same time via another session if the configuration is deleted it is possible for the above spawned MGD process performing the junoscript query to crash. MGD process crashes while accessing a NULL parent which contained an object previously which was deleted. The fix addresses this by not exporting the object which has no parent. [PR844795](#): This issue has been resolved.
- On a device that is in the configuration private mode, when you attempt to deactivate a previously defined VLAN members list and then commit the change, the mgd process creates a core file. [PR855990](#): This issue has been resolved.
- Packet dropped with reject route is currently subjected to loopback filter processing on MPCs, as a result the packet dropped by a reject route might be seen in the output of "show firewall log". This behavior will be changed so that this traffic is no longer subjected to loopback filter processing to bring it in line with other line cards. [PR858511](#): This issue has been resolved.
- Once ingress queuing is enabled on MX Series routers L2 control traffic had no default classifier assigned and used best-effort queue. Under queue congestion, L2 control traffic like IS-IS might get behind and trigger an adjacency flap. L3 control traffic and MPLS control traffic are not affected. [PR858882](#): This issue has been resolved.

- In IPFIX context: 1. In an IPv6 single stack environment, when exporting Data and Template records for family IPv6, the Template records sequence number is not initialized and is always == 0 for all records. This is because the Template sequence numbers are blindly copied from family IPv4 and if this is not configured for IPFIX, then the Sequence Number is always 0. 2. In an IPv4 + IPv6 dual stack environment, since the Template records sequence numbers will be identical for both families, we will get Data and Template records sequence numbers being interleaved when exported. This could confuse the Flow Collector and mislead it into reporting random missing flows. [PR859169](#): This issue has been resolved.
- On MX Series routers with MPCs/MICs, error message "LUCIP(x) has no shadow data for IDMEM[0x00xxxx]" might be seen. [PR859424](#): This issue has been resolved.
- In some corner cases SPMB can stuck in READY state. Restarting the SPMB does not help to recover from the problem state. [PR866127](#): This issue has been resolved.
- The mgd crashed with core-dump after executing **show configuration | display rfc5952**. [PR869650](#): This issue has been resolved.
- On MX Series based line cards, after repeated firewall filter delete/change operations (which might occur with interface flaps, e.g.), memory might leak which can cause ASIC memory exhaustion, causing MX Series based line cards to crash and generate a core file. [PR875276](#): This issue has been resolved.
- When we are deleting a configuration hierarchy which has no groups applied, the corresponding group object hierarchy is also marked as changed in commit script view. [PR878940](#): This issue has been resolved.
- While configuring a filter with a generic prefix followed by specific one in different terms might lead to incorrect match, this might lead to packet drop. [PR886955](#): This issue has been resolved.

Routing Policy and Firewall Filters

- If RPF and/or SCU is enabled then any change to an ingress firewall table filter will trigger RPF/SCU reconfiguration for every prefix in the routing table. This may cause transient high CPU utilization on the fpc which may result in SNMP stats request being time out. [PR777082](#): This issue has been resolved.

Routing Protocols

- If you have configured PIM nonstop active routing (NSR), a core file might be created on an upstream router because of high churn in unicast routes or a continuous clearing of PIM join-distribution in the downstream router. To prevent this possibility, disable NSR for PIM. [PR707900](#): This issue has been resolved.
- On a device that is running Protocol Independent Multicast (PIM) and with nonstop active routing (NSR) enabled on the device, if a PIM corresponding interface flaps continuously, a PIM thread might attempt to free a pointer that has already been freed, causing the routing protocol process (rpd) to crash and create a core file. [PR801104](#): This issue has been resolved.
- With OSPFv3, PIMv6 or LDP configured, the periodic packet management daemon (ppmd) takes responsibility for these protocols' adjacencies. In a rare condition, kernel

might send an invalid packet with a null destination in the message header to ppm process, causing ppm process to crash and generate a core file. [PR802231](#): This issue has been resolved.

- In subscriber management environment, routing protocol process (rpd) might crash and create a core file due to snmpwalk fails at mplsL3VpnVrfRtlnetCidrDestType when a subscriber access-internal route in a VRF has a datalink nexthop (such as when DHCP subscriber connects into a VRF). When issue happens, the following behaviors could be observed: user@router> show snmp mib walk ascii mplsL3VpnVrfRtlnetCidr | no-more Request failed: Could not resolve 'mplsL3VpnVrfRtlnetCidr' to an OID user@router> show snmp mib walk ascii mplsL3VpnVrfRtlnetCidrDest | no-more Request failed: General error. [PR840323](#): This issue has been resolved.
- In IS-IS scenario, with graceful Routing Engine switchover and nonstop active routing (NSR) enabled, after Routing Engine switchover, in very rare case, routing protocol process (rpd) might crash and generate a core file on new master (old backup) Routing Engine. This crash happens upon IS-IS lsp generation due to memory corruption. [PR841558](#): This issue has been resolved.
- Under certain conditions moving a link that has BFD clients can cause stale BFD entry for the old link. [PR846981](#): This issue has been resolved.
- The upstream interface of multicast rpf not matching multicast route in Inter-AS PIM. [PR847370](#): This issue has been resolved.
- In multicast environment with PIM configured, in RP-on-a-stick scenario (aka one-legged RP), if the rendezvous point (rp) receives multicast traffic but there are no receivers, RP's kernel will keep sending resolve requests to routing protocol process (rpd). These resolve requests might get stuck in resolve queue delaying other (S,G) resolves and thereby multicast traffic will be blackholed. [PR851210](#): This issue has been resolved.
- When an import-policy change rejects a BGP-route previously contributing to BGP-Multipath formation, the Peer Active-route-counters in "show bgp neighbor" may not get updated correctly. [PR855857](#): This issue has been resolved.
- If an invalid PIM-SSM multicast group is configured on the routing device, then when you issue the "commit" or "commit check" command, a routing protocol process (rpd) core file is created. There is no traffic impact because the main rpd process spawns another rpd process to parse the corresponding configuration changes, and the new rpd process crashes and creates a core file. When this problem occurs, you might see the following messages: user@router# commit check error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86) error: configuration check-out failed user@router# commit error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86) error: configuration check-out failed. [PR856925](#): This issue has been resolved.
- Routing protocol process (rpd) crashes and generates core files when non-bgp routes (e.g. static route) being advertised as add-path route. [PR859307](#): This issue has been resolved.
- RPD generates a core file. [PR863148](#): This issue has been resolved.

- Multicast packets coming with source address as 0.0.0.0, might cause the RPD to crash. [PR866800](#): This issue has been resolved.
- In VPLS multi-homing environment, with same route-distinguisher configured for the VPLS primary PE and backup PE, routing protocol process (rpd) may crash and generate a core file in each of following two scenarios: 1 - On VPLS backup PE, enable "advertise-external" knob, then rpd process crashes and generates a core file on backup PE. 2 - On VPLS primary PE, enable "advertise-external" knob, after disabling the VPLS interface, rpd process crashes and generates a core file on primary PE. When issue happens, the following behavior could be observed: user@router> show bgp neighbor error: the routing subsystem is not running user@router> show vpls connections error: the routing subsystem is not running. [PR869013](#): This issue has been resolved.
- In Release 12.1 MPLS OAM programs BFD, it does not provide the source address (no change in behavior). In BFD before programming PPMD it queries kernel for the source address matching the prefix of the destination address on a interface. BFD programs PPMD with this source address. PPMD will construct BFD packet with BFD provided source address in the IP header. [PR870421](#): This issue has been resolved.
- In inter-AS Option-B L2VPN scenario, the ASBR might create a L2VPN cloned transit route incorrectly due to a cloned route is a Juniper specific mpls.0 route which the Junos OS creates on the penultimate hop router. Then in a rare case, routing protocol process (rpd) tries to delete the L2VPN cloned transit route (in mpls.0 table) multiple times. After this, routing protocol process (rpd) crashes and generates a core file. [PR878437](#): This issue has been resolved.

Services Applications

- Extensive CLI requests associated with l2tp (show services l2tp < switch >) might result in l2tpd process crash. [PR755948](#): This issue has been resolved.
- Only 94 GRE(plain) sessions are in Established state after chassisd restart. [PR801931](#): This issue has been resolved.
- Memory leak in key management daemon (kmd) causes some IPSec VPN tunnels to be dropped and don't get re-negotiated for over 10 minutes. Before issue happens, the following logs could be observed: /kernel: Process (1466,kmd) attempted to exceed RLIMIT_DATA: attempted 131080 KB Max 131072 KB /kernel: Process (1466,kmd) has exceeded 85% of RLIMIT_DATA: used 132008 KB Max 131072 KB. [PR814156](#): This issue has been resolved.
- The jnxNatSrcNumPortInuse counter is not refreshing when polling the jnxNatSrcNumPortInuse OID via SNMP after RSP switchover. [PR829778](#): This issue has been resolved.
- MAC Flow-control asserted and MS-DPC reboot is needed. [PR835341](#): This issue has been resolved.
- 1) corrected the log to state 4 bundles per tunnel to have been exhausted. 2) change the log level from INFO to DEBUG 3) Add more context to previous log: New IPSec SA install time 1356027092 is less than old IPSec SA install time 1356027092 new log = Tunnel:< tunnel-id > < Local_gw,Remote_gw >: < local-gw-ip-addr,remote-gw-ip-addr > New IPSec SA install time 1356027092 is less than old IPSec SA install time

1356027092 4) added more context to previous log: SA to be deleted with index 3 is not present new log = SA to be deleted with index 3 is not present < Local_gw, Remote_gw >: < local-gw-ip-addr, remote-gw-ip-addr > 5) added a counter to show the number of times each of these messages occur per tunnel. [PR843172](#): This issue has been resolved.

- Syslog is not sent to remote host when rsp interface is used. [PR849995](#): This issue has been resolved.
- When allocate the memory from shared memory for bitmaps used in port blocks , the Junos OS requests as many bytes as the size of the block. If customers assign like 10K block size for deterministic nat or PBA then the Junos OS allocates 10K bytes for that bitmap. However, it only needs 10K/8 bytes as one byte can represent 8 ports. These huge allocations are leading to memory depletion when many source addresses are behind the NAT, and port blocks are big. [PR851724](#): This issue has been resolved.
- The jnxNatSrcNumSessions SNMP OID is broken in 11.4R6-S1 release. [PR851989](#): This issue has been resolved.
- Defining an application with destination-port range starting at 0 can cause TCP handshake to fail through NAT. As a workaround, specify the application with destination-port range starting at 1 instead of 0. [PR854645](#): This issue has been resolved.
- The number of terms per NAT rule cannot exceed 200 for the inline-service si- interface. This constraint check is not applicable for other type of service interfaces like sp-, AMS and ms- etc. Following error message will be displayed when there are more than 200 terms per NAT rule: regress@aria# commit [edit services] 'service-set ss8' NAT rule rule_8 with more than 200 terms is disallowed for si-0/0/0.8 error: configuration check-out failed. [PR855683](#): This issue has been resolved.
- MS-DPC might crash in certain scenarios when using CGNAT PBA and junos-rsh, junos-rlogin, junos-rpc-services-udp and junos-rpc-services-tcp ALGs (either one) in combination with EIM. [PR862756](#): This issue has been resolved.
- When DHCP subscribers log in and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: rts_gencfg_dependency_ifstate(): dependency type (2) is not supported." [PR864444](#): This issue has been resolved.
- Service PIC might crash in corner cases when SIP ALG media flows are deleted. [PR871638](#): This issue has been resolved.
- The issue is seen because of receiving malformed LCP configure-request packet with bad option length from PPP client. In this case when router tries to generate configure-nak it crashed. As a fix, check is added to discard such malformed configure-request packets. [PR872289](#): This issue has been resolved.

Subscriber Access Management

- Subnet mask option is not returned to DHCP client when framed-ip-address is used with dhcp-local-server. [PR851589](#): This issue has been resolved.
- Some requests internally sent to AUTHD process experience a timeout state which may cause the subscribers to remain as either release or terminated. [PR853239](#): This issue has been resolved.

- Authd core experienced when multiple DHCP subscriber connection attempts require SRC for subscriber authentication. [PR862037](#): This issue has been resolved.
- Fixed the misbehavior of 'accounting-stop-on-failure' configuration knob. [PR865305](#): This issue has been resolved.
- PPPoE subscribers do not always get disconnected after the client-session-timeout expires. [PR869559](#): This issue has been resolved.

User Interface and Configuration

- The blank set command while indicating the configuration by **show | display inheritance | display set**. [PR816722](#): This issue has been resolved.
- If a commit sync error occurs for a commit performed in "edit private" mode and later it is followed by another commit in global mode (without private or exclusive mode), the configuration file may remain unzipped after the global commit is complete. [PR823555](#): This issue has been resolved.

VPNs

- Deleted logical interfaces may not be freed due to references in MVPN. [PR851265](#): This issue has been resolved.
- When "multicast omit-wildcard-address" is configured on a route-reflector for the MVPN address families, Leaf-AD route NLRIs are not reflected correctly in the newer, standardized format. The Leaf-AD routes transmitted from the RR in the new format will have invalid Leaf-IP fields in the NLRI set to 0.0.0.0. As a result, ingress PEs may fail to properly identify all egress PEs and thus fail to update provider-tunnel state to deliver traffic to those egress PEs. [PR854096](#): This issue has been resolved.
- While l2circuit/l2vpn is not configured, if user requests for PW object info through mib, l2circuit/l2vpn is creating invalid job, which can lead to rpd crash. The fix exists in: 12.3R3, 11.4R8, 13.1R2, 12.2R5, 12.1R7 and later releases. [PR854416](#): This issue has been resolved.
- When the egress PEs are on a NGMVPN, which then leads on to the assert being silently ignored when dual forwarders are setup over the PE-CE segment. Eventually duplicate traffic being delivered by PE routers onto the ethernet where receiver is connected. [PR862586](#): This issue has been resolved.
- RPD can crash when a cmcast leave is received after disabling the internet-multicast. [PR864304](#): This issue has been resolved.
- Sample topology: multicast +---+ CE_R +---+ PE_R +---MPLS core---+ PE_S +---+ C-BSR +---+ C-RP +---+multicast receiver source With the NG MVPN setup, when RP failed, there could be a delay on RP timeout between PE_S (multicast traffic ingress) and PE_R (multicast traffic egress). And suppose that PE_S removed RP from the PR list and PE_R still learned RP. Under the condition above, when RP came back and BSR informed RP info with generating bootstrap message, PE_R would advertise type 6 routes to PE_S across MPLS core via MPBGP. If a RP is learned on PE_S after PE_S

receives the type 6 routes from the core, PE_S neither creates PIM (*G) join nor sends the join to C-RP. [PR866962](#): This issue has been resolved.

- If a logical interface is taken out of VPLS or L2VPN Pseudowire Routing Instance and placed in protocol l2circuit, after the above configuration changes are done in one commit, routing protocol process (rpd) crashes and generates a core file. [PR872631](#): This issue has been resolved.

Resolved Issues in Release 12.3R2

Class of Service

- When rate limit is enabled and disabled on port cos scheduler configuration leaves rate limit configuration on queues in effect. This causes the rate limit feature in effect even after rate limit is removed. This PR addresses this issue in lieu with PR 843603. [PR833431](#): This issue has been resolved.
- Traffic-control-profile-remaining is not working for logical interface in interface-set. [PR835933](#): This issue has been resolved.
- In PPPoE/DHCP subscriber management environment, with "burst-size \$junos-cos-shaping-rate-burst" configured in subscriber dynamic-profiles, while logging in/out subscribers, the class-of-service daemon (cosd) memory leak due to cosd process doesn't free up memory used for parsing burst attributes of a traffic-control-profiles (tcp) guaranteed rate. The memory usage of cosd process can be monitored by following CLI command: user@router> show system processes extensive | match "PID | cosd" (Note: The "RES" field means "Current amount of resident memory, in kilobytes")
PID USERNAME THR PRI NICE SIZE RES STATE TIME
WCPU COMMAND 1326 root 1 96 0 14732K 4764K select 0:01 0.00% cosd [PR846615](#): This issue has been resolved.
- This seems to be hard to reproduce and noticed only once after GRES. When the cosd restarts (due to the GRES test you performed), cosd reconciles the configurations pushed to the Packet Forwarding Engine with config read from CLI and tries to reuse the object ID. In this case, it was trying to insert the same ID twice. [PR848666](#): This issue has been resolved.
- Commit throws an error "Invalid rewrite rule rule-name for logical interface <ifl-name>. Ifd <ifd-name> is not capable to rewrite inner vlan tag 802.1p bits" even though there is no rewrite configuration related to inner-vlan tag. [PR849710](#): This issue has been resolved.
- Configuring Classifiers under groups might result in Class-of-service daemon to core. Work-around is to avoid configuring Classifiers under groups. [PR863109](#): This issue has been resolved.

Forwarding and Sampling

- With more than four archive-sites configured under [system archival configuration archive-sites] hierarchy, after committing the configuration changes, pfe process crashes and generates a core file due to memory corruption or double free. The core files could be seen by executing the CLI command **show system core-dumps**. [PR849465](#): This issue has been resolved.

General Routing

- Prior to this change, the L2TP sessions with cos/ firewall attachments fail to come up when the L2TP Access Concentrator (LAC) is reachable over a unicast nexthop. [PR660208](#): This issue has been resolved.
- Reconfiguring a deleted interface with BFD sessions can take up to 20 minutes for the BFD sessions to initialize. [PR786907](#): This issue has been resolved.
- With l3vpn composite next-hops configured and 3 or more odd number of core uplinks every l3vpn route deletion will syslog the following error messages. [LOG: Err] JTREE: (jt_mem_free) size 0 for addr 1595452, seg 1, inst 0 [LOG: Emergency] Multiple Free :jt_mem_free There is no operational impact. An even number of core-uplinks will not trigger such error logs. [PR786993](#): This issue has been resolved.
- MPLS LDP traceroute does not work if you have a default route 0/0 pointing to discard on the Egress router with DPC cards. [PR790935](#): This issue has been resolved.
- On T1600-FPC4-ES, T640-FPC3, T640-FPC3-E and T640-FPC3-E2 platforms which have multiple Packet Forwarding Engines, with auto-bandwidth enabled on LSPs where CoS-based forwarding (CBF) is configured, auto-bandwidth might trigger minor changes on LSP nexthops. After this, flapping corresponding interface or any nexthop changes may result in FPC crash and create a core file. The core files can be seen by executing CLI command **show system core-dumps**. This issue will be seen with auto-bw configuration where there will continuous minor/major changes on LSP nexthops based on traffic conditions. When issue happens, the following logs could be seen: fpc3 PDP(pdp_free): %PFE-3: Invalid PDP 0x4e01d7d0 fpc3 PDP(pdp_free): %PFE-3: Error while removing PDP (0x4df4c068) fpc3 PDP(pdp_free): %PFE-3: Error while removing PDP (0x525b3f78) fpc3 PDP(pdp_free): %PFE-3: Invalid PDP 0x4de522b0' [PR818021](#): This issue has been resolved.
- icmp redirects are not disabled even after configuring no-redirects on irb interface. [PR819722](#): This issue has been resolved.
- When an MS-DPC PIC reboots due to a crash or manual intervention, it might get stuck in a booting loop if the MS-DPC up-time is more than 49 days and 17 hours. After 5 consecutive boot failures, the MS-DPC PIC will go offline automatically and gives the following error message: [15:21:22.344 LOG: Err] ICHIP(0): SPI4 Training failed while waiting for PLL to get locked, ichip_sra_spi4_rx_snk_init_status_clk [15:21:22.344 LOG: Err] CMSPC: I-Chip(0) SPI4 Rx Sink init status clock failed, cmsdpc_spi4_init [15:21:22.344 LOG: Err] CMX: I(0) ASIC SPI4 init failed [15:21:22.379 LOG: Err] Node for service control logical interface 68, is already present [15:21:23.207 LOG: Err] ASER0 SPI-4 XLR source core OOF did not go low in 20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 stop src train failed! [15:21:23.208 LOG: Err] ASER0 XLR SPI-4 sink core DPA

incomplete in 20ms. [15:21:23.208 LOG: Err] ASER/XLRO spi4 sink core init failed! [15:21:24.465 LOG: Err] ICHIP(0): SPI4 Stats Unexpected 2'b 11 Error, isra_spi4_parse_panic_errors [15:21:24.465 LOG: Err] ICHIP(0): SPI4 Tx Lost Sync Error, isra_spi4_parse_panic_errors In order to recover from this state the whole MS-DPC needs to be rebooted. [PR828649](#): This issue has been resolved.

- PPPoE sessions cannot be established as rpd is unable to read or access profile database during access-internal route creation via "dynamic-profile->routing-instances->routing-options->access-internal" stanza [PR830779](#): This issue has been resolved.
- An FPC may reboot when a live-core is requested and the /var partition does not have sufficient space to store the live-core. [PR835047](#): This issue has been resolved.
- On T4000 systems where the following conditions are met: - the "forwarding-options sampling input maximum-packet-length" knob is configured to a non-zero value - packets are sent to be sampled from a Type 5 FPC to an ES-Type FPC housing the Multiservices PIC used for sampling then an incorrect format of the notification header sent to the destination ES-Type FPC will trigger a packet loss in the packets sent to be sampled. The following message will be logged in the syslog on the destination FPC: [Jan 17 12:43:25.388 LOG: Err] SRCHIP(0): 1 Bad packets on p1 [Jan 17 12:43:25.389 LOG: Err] SRCHIP(0): 1 SONN errors on p1 The outcome is that the respective packets will be dropped and they will not be sampled. [PR839696](#): This issue has been resolved.
- When you configure tunnel interface in MXVC, the tunnel interface is set to harddown, unfortunately, there is no workaround at this point >High level problem description of the problem Problem: tunnel interface is set to harddown in MXVC >When does it occur When a tunnel interface is configured, it is always set to harddown >Is there a workaround, and if yes, what it is... Unfortunately, there is not workaround to bring this interface up. A fix is planned for R2. [PR839784](#): This issue has been resolved.
- When the transit traceroute packets with ttl=1 are received on the LSI interface, you may retrieve the Source Address from the LSI interface to reply ICMP. As LSI does not have any IFA, it will use first the IFA in routing-instance to reply. So Source Address used was the first IFA added in VPN routing-instance. As a workaround, if the incoming interface is LSI, then retrieve Source Address from the logical interface which is having the Destination IP Address. This will make sure we reply with Source Address from CE-facing the logical interface. [PR839920](#): This issue has been resolved.
- Dynamic arp or routing does not work when using ether-over-atn-llc in the new PIC. [PR840159](#): This issue has been resolved.
- mlfr/mlppp interface are not reachable after restart FPC (primary MSPIC) followed by deactivate and activate R.I or GRES followed by deactivate and activate R.I. This is because link FPC does not have the interfaces programmed towards the bundle [PR847278](#): This issue has been resolved.
- Maximum power required for SFBs is changed from 250W to 220W. Maximum power required for 172mm Fan Trays is increased from 1500W to 1700W. The power requirement for MX2010's upper fan trays is not changed. It is still 500W. With this change, the Reserved Power for critical FRUs (CB/Routing Engine, SFB and FanTrays) changes from 7000W to 7360W for MX2020 and from 6500W to 6660W for MX2010. [PR848358](#): This issue has been resolved.

- Distributed protocol adjacencies (LFM/BFD/etc) may experience a delay in keepalives transmission and/or processing due to a prolonged CPU usage on the FPC microkernel on T4000 Type 5-3D FPCs. The delay in keepalive transmission/processing may result in a mis-diagnosis of a link fault by the peer devices. The issue is seen several seconds after an Routing Engine mastership switch with NSR enabled and the fault condition will clear after a couple of minutes. [PR849148](#): This issue has been resolved.
- FPC/PICs usually have high response time when there are loaded with high traffic. Kernel cores when a PIC/FPC stops responding kernel after a new connection or a reconnection. [PR853296](#): This issue has been resolved.
- After configuring MX as ingress for RSVP LSP's, all the FPC's start throwing the error message "TOPO_FLAVOR_IFF_HW_OUT before family (1) : 0x0", this is a cosmetic issue with no impact to any protocol functionality. [PR854499](#): This issue has been resolved.

High Availability and Resiliency

- On TX Matrix routers with four LCCs and IQ2 PICs, in-service software upgrade (ISSU) from 12.3R1.7 to a newer release results in traffic loss and a FRU upgrade error. [PR768502](#): This issue has been resolved.

Infrastructure

- A kernel crash may occur on routers running 10.4 or higher (which does not have fix for this PR), with "targeted-broadcast" knob configured on a broadcast interface. If this knob is configured, MAC address will be learned for subnet broadcast IP (configured on that interface). When this ARP table entry gets timed out, it corrupts an internal data structure, leading to kernel crash. This MAC learning will happen with one of the following : 1. Mismatched IP subnet is configured on one of the connected devices 2. A malformed packet (ARP request to subnet broadcast IP) is received on that interface
NOTE: - MAC address learned for the subnet broadcast IP can't be seen using "show arp" command. - This issue is platform independent. [PR814507](#): This issue has been resolved.

Interfaces and Chassis

- Under certain circumstances, MX80 may crash when using the command "request system snapshot". [PR603468](#): This issue has been resolved.
- Kernel can cache a high incorrect value for stats and is rejecting the correct subsequently stats coming from the PIC. The fix consists in checking if the difference of what is cached in kernel and what is reported by the PIC is less than an acceptable value. If the answer is not kernel does not gets stuck permanently and recovers while fetching stats next time. [PR806015](#): This issue has been resolved.
- "show interfaces redundancy" may display secondary as down upon following sequence: deactivate R.I.(that contains entire mfr logical interfaces)-->restart fpc(that holds secondary MS pic)--> activate the R.I. back [PR816595](#): This issue has been resolved.
- Warning message added is syslog when external sync is not supported. [PR817049](#): This issue has been resolved.

- Hash Key configuration not programmed in the Packet Forwarding Engine correctly after system reboot. [PR818035](#): This issue has been resolved.
- Prior to this PR, the speed of a GE interface capable of working at FE speeds was set to 'auto' in the Packet Forwarding Engine level. This causes a problem when manually setting the speed on the Routing Engine. Now the behavior is to set the speed to '1 g' in the Packet Forwarding Engine. For automatic speed detection the interface should be set to 'speed auto' in the configuration. [PR821512](#): This issue has been resolved.
- MX Series chassis-control interrupt storm may be falsely reported when a Field Replaceable Unit (FRU) is removed, inserted, or FPM button pushed. A FRU may not be recognized/booted, resulting in chassis operational failure. [PR823969](#): This issue has been resolved.
- IEEE 802.3 ah LFM stats counter "OAM current frame error event information" is not cleared correctly by CLI operation. [PR827270](#): This issue has been resolved.
- If per-unit-scheduler is configured under a physical interface(ifd), and trying to delete this ifd and its sub-interfaces (logical interface) in one single commit, ksyncd may core in the backup Routing Engine which will cause GRES malfunction. [PR827772](#): This issue has been resolved.
- Although physical interface is disabled, reseating 1GbE SFP on MPC/MIC restores its output optical power, hence the opposite router interface turns Up(Near-end interface is still down). Only 1g-SFP on MPC/MIC has the problem, but 1g-SFP on DPC/MX, EX Series and 10G-XFP on DPC/MX don't have the problem. When the sfp is reseated, then the sfp periodic is going ahead and enabling the laser irrespective of the fact that interface has been enabled or disabled. Driver needs to store the state for each sfp link and enable laser based on that. This software problem is fixed in 11.4R7, 12.1R6, 12.2R4, 12.3R2 and later release. [PR836604](#): This issue has been resolved.
- Configuring 100-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning. The "optics-options" alarm/warning "low-light"; the syslog action was not taking effect on T1600 and T4k for 100 GE PICs. This was fixed as part of this PR. [PR836709](#): This issue has been resolved.
- The Logical Interfaces are marked with 0 (null) after deactivate system commit synchronize and deactivate chassis redundancy which result backup Routing Engine to core. [PR840167](#): This issue has been resolved.
- ERA events are not credited back by jpppoed. ERA has a purge timer of 10 minutes which reclaims stale events so new connections are allowed after the purge timer fires. In a high scaled scenario this can lead to slow PPPoE connections. [PR842935](#): This issue has been resolved.
- The device configuration daemon (dcd) may crash when a partial demux subinterface configuration is attempted to be committed. There is no impact to traffic forwarding but before the configuration can be committed, it must provide a valid 'underlying-interface' for the demux subinterface. [PR852162](#): This issue has been resolved.

Layer 2 Ethernet Services

- It can happen that when changing an interface framing from lan-phy (default) to wan-phy and back a few times, the interface doesn't show up any more in "show interfaces terse". [PR836382](#): This issue has been resolved.
- DHCPv6 relay terminates the client if DHCPv6-REPLY message from server contains status-code option. [PR845365](#): This issue has been resolved.
- In certain cases when MX Series is configured as DHCPv6 server and servicing DHCPv6 clients through LDRA relay it may send advertisements with UDP port 546 instead of 547. [PR851642](#): This issue has been resolved.

Network Management and Monitoring

- The default maximum log file size depends on the platform type for TX Matrix or TX Matrix Plus routers it is expected to be 10 MB. However, due to a software defect, this file size was only 1 MB. [PR823143](#): This issue has been resolved.
- On a router with interfaces with Frame Relay encapsulation a SNMP WALK operation will cause a MIB daemon (mib2d) crash and will generate a mib2d core-dump. The crash itself does not cause any impact on the router as the MIB daemon is restarted automatically. The only effect is that a SNMP WALK will never complete successfully.
user@router-re1> show snmp mib walk 1 | no-more sysDescr.0 = Juniper Networks, Inc. mx480 internet router, kernel JUNOS OS 11.4R6.5 #0: 2012-11-28 21:57:12 UTC
builder@evenath.juniper.net:/volume/build/junos/11.4/release/11.4R6.5/obj-i386/bsd/kernels/JUNIPER/kernel Build date: 2012-11-28 21:39:15 UTC Copyright (c sysObjectID.0 = jnxProductNameMX480 sysUpTime.0 = 339594 sysContact.0 < > dot3OutPauseFrames.942 = 0 dot3OutPauseFrames.943 = 0 dot3OutPauseFrames.953 = 0 dot3OutPauseFrames.954 = 0 frDlcmilIndex.153 = 153 frDlcmilIndex.512 = 512 frDlcmilIndex.513 = 513 frDlcmiState.153 = 6 Request failed: General error user@router-re1> show log messages Dec 20 09:23:20 router-re1 clear-log[8240]: logfile cleared Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3-BAD_PAGE_FAULT: pid 7382 (mib2d), uid 0: pc 0x810fe09 got a read fault at 0x7c, x86 fault flags = 0x4 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Trapframe Register Dump: Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: eax: 00000000 ecx: bfbeda88 edx: 00000000 ebx: bfbeda7c Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: esp: bfbeda60 ebp: bfbeda98 esi: 089de834 edi: 089fb680 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: eip: 0810fe09 eflags: 00010297 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: cs: 0033 ss: 003b ds: bfb003b es: 003b Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: fs: 003b trapno: 0000000c err: 00000004 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Page table info for PC address 0x810fe09: PDE = 0x42e60067, PTE = 5290c425 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Dumping 16 bytes starting at PC address 0x810fe09: Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: 8b 40 7c 89 04 24 e8 5a 3f 2f 00 89 45 ec 8b 55 Dec 20 09:23:40.787 router-re1 init: %AUTH-3: mib-process (PID 7382) terminated by signal number 11. Core dumped! Dec 20 09:23:40.787 router-re1 init: %AUTH-6: mib-process (PID 8247) started Dec 20 09:23:40.809 router-re1 mib2d[8247]: %DAEMON-5-LIBSNMP_SA_IPC_REG_ROWS: ns_subagent_register_mibs: registering 88 rows Dec 20 09:23:41.595 router-re1 mib2d[8247]: %DAEMON-6-LIBSNMP_NS_LOG_INFO: INFO:


```

ns_subagent_open_session: NET-SNMP version 5.3.1 AgentX subagent connected Dec
20 09:23:43.533 router-re1 dumpd: %USER-5: Core and context for mib2d saved in
/var/tmp/mib2d.core-tarball.0.tgz Dec 20 09:23:43.793 router-re1 mib2d[8247]:
%DAEMON-6-SNMP_TRAP_LINK_UP: ifIndex 5, ifAdminStatus up(1), ifOperStatus
up(1), ifName dsc < ..... > user@router-re1> show system core-dumps
/var/crash/*core*: No such file or directory -rw----- 1 root field 680417 Dec 20 09:23
/var/tmp/mib2d.core-tarball.0.tgz /var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory /tftpboot/corefiles/*core*: No such file
or directory total 1 PR835722: This issue has been resolved.

```

Platform and Infrastructure

- When using configure private with large group definition and high number of groups the commit process can spend a lot of time to merge the configuration change with the global configuration. [PR828005](#): This issue has been resolved.
- This applies to all Juniper M, MX, and T Series routers. In certain GRES scenarios, the backup Routing Engine may not have the complete state of the NH database from the active Routing Engine and may send duplicate NH add messages to Packet Forwarding Engine with same NH Ids when it becomes active. This could potentially cause undesirable behavior in forwarding resulting in broken forwarding state and/or FPC cores. To limit the affect of these duplicate NH add messages, only certain duplicate NH adds messages which can be handled gracefully are allowed and all other duplicate add messages are rejected. There is no work-around for this problem. [PR843907](#): This issue has been resolved.
- On MX Series and T4000, when output Filter-Based Forwarding (FBF) destined to a routing-instance is configured, the packets matched by the FBF filter may be discarded or sent to the unintended Packet Forwarding Engine. [PR845700](#): This issue has been resolved.

Routing Protocols

- If maximum-paths or maximum-prefixes is configured for a route table, these limits are displayed in the output of "show route summary". In affected releases, these limits were omitted from the output of "show route summary". [PR753013](#): This issue has been resolved.
- Due to duplication of the traffic, assert will be triggered. *G and S,G assert is not handled properly hence few assert entries will not be deleted due to Routing Engine switchover which result in a core file. ?HW type of chassis/linecard/RE. "ALL" ?Suspected software feature combination. Multicast feature ?Describe if any behavior/ change to existing function - Handle the *G and S,G assert properly. [PR809338](#): This issue has been resolved.
- This PR fixes a bug by which the receiving EBGP speaker mistakenly accepts a session establishment attempt from an EBGP peer address that is not directly connected because it did not check to see if the address to which peer wants to establish a session belongs to the receiving interface or not. [PR816531](#): This issue has been resolved.
- Changes to add-path prefix-policy do not get absorbed automatically, and require a manual soft-clearing of the BGP session. [PR818789](#): This issue has been resolved.

- RPD on the backup Routing Engine might crash when it receives a malformed message from the master. This can occur at high scale with nonstop active routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare and backup RPD will restart and the system will recover without intervention. [PR830057](#): This issue has been resolved.
- Multiple route nexthops will not be returned via SNMP for ipCidrRouteTable object. [PR831553](#): This issue has been resolved.
- If LDP-SYNC *<hold-down>* timer is configured under IS-IS interfaces after configuration change the IS-IS interfaces can go to *<hold-down>* state. [PR831871](#): This issue has been resolved.
- 1) What triggers the bug to be happened? =>Enabling PIM - Bidir feature (possibly pim rp with 224.0.0.0/4 group) and rpd restart. This issue is hit during regression test for PIM bidir. 2) HW type of chassis/linecard/Routing Engine. If it affects all, just say ?all?. =>all. 3) Suspected software feature combination. (If customer turns on feature X along with Y, they might hit, etc) =>PIM - Bidir feature (rp configured) and rpd restart is causing the issue. 4) Describe if any behavior/ change to existing function =>None. [PR836629](#): This issue has been resolved.
- On EX8200 switches, multiple rpd process core files might be created on the backup Routing Engine after a nonstop software upgrade (NSSU) has been performed while multicast traffic is on the switch. [PR841848](#): This issue has been resolved.
- IS-IS reports prefix-export-limit exceeded even though the number of exported routes is smaller than the configured value of prefix-export-limit. [PR844224](#): This issue has been resolved.
- In scenarios that use BGP to distribute traffic flow specifications, if the received flow-spec Network Layer Reachability Information (NLRI) contains invalid argument (such as dscp is larger than 63), routing protocol process (rpd) will generate flow-spec routes and install them in the routing table for these NLRIs; but these flow routes with invalid match conditions are rejected by dynamic firewall daemon (dfwd) from being added to the flowspec filters. When issue happens, the following errors could be seen:
krt_flow_trans_match_config: Failed defining match conditions
10.0.1.1,1.0.0.1,proto=6,dscp=81 krt_flow_trans_term_add: Failed adding term
10.0.1.1,1.0.0.1,proto=6,dscp=81 to filter 0x9504000 - Unknown error: 0
krt_flow_trans_filter_add: Failed sending transaction (ADD FILTER SINGLE TERM) for
filter 0x9504000 __flowspec_default_inet__ to add term 10.0.1.1,1.0.0.1,proto=6,dscp=81
- Invalid argument When the bgp peer withdraws these flow routes, they will only be
deleted but not freed, hence cause memory leak. [PR845039](#): This issue has been
resolved.
- In BGP scenario with multipath configured, if a static route which has table nexthop (such as inet.0) is configured in the same routing-instance as BGP, when an interconnect link between BGP peers is brought down or flapping, the corresponding BGP session takes 90 seconds to timeout. During this period routes received over the BGP session will stay there. For a multipath transit route received from both BGP sessions, initially both paths are resolved over the interconnect links directly. When one of the interconnect link is brought down or flapping, that path will be resolved over the static default route which has table nexthop (such as inet.0). So now, one path is resolved

over a router nexthop and the other path is resolved over a table nexthop. This will cause routing protocol process (rpd) crash and generate a core file. This issue usually occurs in BGP/L3VPN environment. The core files could be seen by executing CLI command **show system core-dumps**. [PR851807](#): This issue has been resolved.

Services Applications

- SIP ALG was not allowing SIP 603 decline message. [PR822679](#): This issue has been resolved.
- jlt2pd crash_thr_send_sig (thread=0x8a5e000, sig=6) at `../../../../src/bsd/lib/libthr/thread/thr_kern.c:91` jlt2pd crash exhibited in environment where MX480 was configured as LAC and terminating 500 l2tp subscribers. [PR824760](#): This issue has been resolved.
- When MX Series uses MS-DPC to provide the tunnelling service for flow-tap traffic, if there is SCU/DCU configured on the same slot of the flow-tap traffic ingress interface, all the flow-taped sampled packets will be dropped. It is caused by the wrong nexthop linking when DCU is configured. [PR825958](#): This issue has been resolved.
- This issue is seen when two l2tp users get connected to same routing-instance and they get same framed routes. When last connected user disconnects this issue can be seen. [PR832034](#): This issue has been resolved.
- In the case of a stateful proxy, two SIP users behind the NAT device (so-called SIP hairpinning) will be unable to signal the call. [PR832364](#): This issue has been resolved.
- With RTSP ALG enabled, RTSP keep-alive packets might be dropped if it's already Ack'ed by the receiver. [PR834198](#): This issue has been resolved.
- In Carrier Grade NAT (CGNAT) scenario, without any configuration change, under some conditions, MS-DPC PIC might crash and create a core file when encountering unknown flow-type. Service will be impacted during the period. When issue happens, the following logs could be seen: `chassisd[1477]: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power off (jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @ 5/1/*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 8, jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 33404122)` `chassisd[1477]: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @ 5/1/*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 2, jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 192338924)` [PR834899](#): This issue has been resolved.
- In scenarios which use sp interface, such as IPSec VPN, multiservice process (mspd) will memory leak during sp interface flapping. The memory usage of mspd process can be checked by following CLI command: `user@router> show system processes extensive | match "PID | mspd"` (Note: The "RES" field means "Current amount of resident memory, in kilobytes") `PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND` `2048 root 1 96 0 36216K 34820K select 0:10 0.00% mspd` When the memory usage of mspd process increases to system limit (about 131072KB), the following logs could be seen: `/kernel: %KERN-5: Process (2048,mspd) attempted to exceed RLIMIT_DATA: attempted 131076 KB Max 131072 KB` [PR836735](#): This issue has been resolved.

- When DHCP subscribers login and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: GENCFG: op 24 (Lawful Intercept) failed; err 5 (Invalid)."
[PR837877](#): This issue has been resolved.
- The "hot-standby" CLI knob under [edit interfaces <RSP-interface-name> redundancy-options] is made hidden for the Redundant Service PIC (RSP).
[PR838762](#): This issue has been resolved.
- If flow-tap or radius-flow-tap is configured and logging, dfcd might be leaking file descriptors. RPD may crash and write a core with a signature like "kern.maxfiles limit exceeded by uid 0" due to this issue.
[PR842124](#): This issue has been resolved.
- Service PIC might crash under certain race conditions when receiving sip invite packets.
[PR843047](#): This issue has been resolved.
- Service PIC might crash in corner cases when receiving specific SIP REGISTER.
[PR843479](#): This issue has been resolved.
- Service PIC might crash in corner cases when EIM is enabled for SIP ALG.
[PR847124](#): This issue has been resolved.
- spd core generated during switchover when CGAT config is there. Issue is well understood now and has been fixed in later releases.
[PR854206](#): This issue has been resolved.

Subscriber Access Management

- Snmpwalk requests sent to MX Series returns multiple duplicate records for jnxUserAAAAccessPool.
[PR840640](#): This issue has been resolved.

VPNs

- In a scaled Multicast VPN setup, where many selective provider tunnels are used, and the MVPN instance is deleted, RPD can sometimes crash.
[PR801667](#): This issue has been resolved.
- In BGP-MVPN, when the number of multicast routes falls below the threshold, the earlier suppressed MVPN multicast routes because of limit are not added back again. For MVPN, there was no mechanism to trigger the processing of cmcast entries that were not added earlier. The fix is to queue the cmcast entries that are suppressed for multicast route addition in a new list. When the reuse limit is reached, this list is walked and used to add back the entries.
[PR841105](#): This issue has been resolved.
- In I2circuit (Martini I2vpn) scenarios where a backup neighbor is being defined along the 'standby' knob, after deleting this backup neighbor from configuration, its associated vc-route is not being eliminated. Later if user deletes the I2circuit neighbor or restarts routing protocol process (rpd), rpd process will crash and core dumped.
[PR841522](#): This issue has been resolved.

Related Documentation

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)

- [Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 181](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 357](#)

Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

This section discusses the following topics:

- [Basic Procedure for Upgrading to Release 12.3 on page 357](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 359](#)
- [Upgrading a Router with Redundant Routing Engines on page 359](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 360](#)
- [Upgrading the Software for a Routing Matrix on page 361](#)
- [Upgrading Using Unified ISSU on page 362](#)
- [Downgrading from Release 12.3 on page 362](#)

Basic Procedure for Upgrading to Release 12.3

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



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

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

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

The download and installation process for Junos OS Release 12.3 is as follows:

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



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

Customers in the United States and Canada use the following command:

```
user@host> request system software add validate reboot source/jinstall-12.3R11  
1-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot source/jinstall-12.3R11  
1-export-signed.tgz
```

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

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

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots

successfully. This is the default behavior when the software package being added is a different release.

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 12.3 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 [Junos OS Installation and Upgrade Guide](#).

Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1

In releases prior to Junos OS Release 10.1, the draft-rosen multicast VPN feature implements the unicast **lo0.x** address configured within that instance as the source address used to establish PIM neighbors and create the multicast tunnel. In this mode, the multicast VPN loopback address is used for reverse path forwarding (RPF) route resolution to create the reverse path tree (RPT), or multicast tunnel. The multicast VPN loopback address is also used as the source address in outgoing PIM control messages.

In Junos OS Release 10.1 and later, you can use the router's main instance loopback (**lo0.0**) address (rather than the multicast VPN loopback address) to establish the PIM state for the multicast VPN. We strongly recommend that you perform the following procedure when upgrading to Junos OS Release 10.1 if your draft-rosen multicast VPN network includes both Juniper Network routers and other vendors' routers functioning as provider edge (PE) routers. Doing so preserves multicast VPN connectivity throughout the upgrade process.

Because Junos OS Release 10.1 supports using the router's main instance loopback (**lo0.0**) address, it is no longer necessary for the multicast VPN loopback address to match the main instance loopback address **lo0.0** to maintain interoperability.



NOTE: You might want to maintain a multicast VPN instance **lo0.x** address to use for protocol peering (such as IBGP sessions), or as a stable router identifier, or to support the PIM bootstrap server function within the VPN instance.

Complete the following steps when upgrading routers in your draft-rosen multicast VPN network to Junos OS Release 10.1 if you want to configure the routers's main instance loopback address for draft-rosen multicast VPN:

1. Upgrade all M7i and M10i routers to Junos OS Release 10.1 before you configure the loopback address for draft-rosen Multicast VPN.



NOTE: Do not configure the new feature until all the M7i and M10i routers in the network have been upgraded to Junos OS Release 10.1.

2. After you have upgraded all routers, configure each router's main instance loopback address as the source address for multicast interfaces. Include the **default-vpn-source interface-name loopback-interface-name** statement at the **[edit protocols pim]** hierarchy level.

3. After you have configured the router's main loopback address on each PE router, delete the multicast VPN loopback address (**lo0.x**) from all routers.

We 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 *Junos OS Multicast Configuration Guide*.

Upgrading the Software for a Routing Matrix

A routing matrix can use either a TX Matrix router as the switch-card chassis (SCC) or a TX Matrix Plus router as the switch-fabric chassis (SFC). By default, when you upgrade software for a TX Matrix router or a TX Matrix Plus router, the new image is loaded onto the TX Matrix or TX Matrix Plus router (specified in the Junos OS CLI by using the **scc** or **sfc** option) and distributed to all T640 routers or T1600 routers in the routing matrix (specified in the Junos OS CLI by using the **lcc** option). To avoid network disruption during the upgrade, ensure 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 T640 routers or T1600 routers (LCC) are all **re0** or are all **re1**.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and T640 routers or T1600 routers (LCC) are all **re1** or are all **re0**.
- All master Routing Engines in all routers run the same version of software. This is necessary for the routing matrix to operate.
- All master and backup Routing Engines run the same version of software before beginning the upgrade procedure. Different versions of 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 Routing Engines.



NOTE: It is considered best practice to make sure that all master Routing Engines are re0 and all backup Routing Engines are re1 (or vice versa). For the purposes of this document, the master Routing Engine is re0 and the backup Routing Engine is re1.

To upgrade the software for a routing matrix, 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 Feature Guide](#) or the [Routing Matrix with a TX Matrix Plus Feature 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 [Junos High Availability Configuration Guide](#).

Downgrading from Release 12.3

To downgrade from Release 12.3 to another supported release, follow the procedure for upgrading, but replace the 12.3 **jinstall** package with one that corresponds to the appropriate release.



NOTE: You cannot downgrade more than three releases. For example, if your routing platform is running Junos OS Release 11.4, you can downgrade the software to Release 10.4 directly, but not to Release 10.3 or earlier; as a workaround, you can first downgrade to Release 10.4 and then downgrade to Release 10.3.

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

**Related
Documentation**

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 102](#)
- [Changes in Default Behavior and Syntax, and for Future Releases in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 162](#)
- [Known Behavior in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 181](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 184](#)
- [Outstanding Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 216](#)
- [Resolved Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 248](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 364](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 368](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 370](#)
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 371](#)
- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 374](#)

New Features in Junos OS Release 12.3 for PTX Series Packet Transport Routers

Powered by Junos OS, PTX Series Packet Transport Routers are a portfolio of high-performance platforms designed for the service provider supercore. These systems deliver powerful capabilities based on innovative silicon and a forwarding architecture focused on MPLS and Ethernet. PTX Series Packet Transport Routers deliver several critical core functions, including industry-leading density and scalability, cost optimization, high availability, and network simplification. PTX Series Packet Transport Routers supported in this release include the PTX5000 system.

The following features have been added to Junos OS Release 12.3 for the PTX Series Packet Transport Routers. Following the description is the title of the manual or manuals to consult for further information:



.....

NOTE: Features described in the Junos OS 12.1X48R4 Release Notes are supported in Junos OS 12.3 except for real-time performance monitoring (RPM) support. See [Junos OS 12.1X48R4 Release Notes for Juniper Networks PTX Series Packet Transport Switches](#).

.....

- [Hardware on page 364](#)
- [Firewall Filters on page 366](#)
- [Interfaces and Chassis on page 366](#)
- [Network Management and Monitoring on page 367](#)
- [User Interface and Configuration on page 368](#)

Hardware

- **CFP-GEN2-CGE-ER4 and CFP-GEN2-100GBASE-LR4 (PTX5000)**—The CFP-GEN2-CGE-ER4 transceiver (part number: 740-049763) provides a duplex LC connector and supports the 100GBASE-ER4 optical interface specification and monitoring. The CFP-GEN2-100GBASE-LR4 transceiver (part number: 740-047682) provides a duplex LC connector and supports the 100GBASE-LR4 optical interface specification and monitoring. Starting in Junos OS Release 13.3, the “GEN2” optics have

been redesigned with newer versions of internal components for reduced power consumption. The following interface module supports the CFP-GEN2-CGE-ER4 and CFP-GEN2-100GBASE-LR4 transceivers. For more information about interface modules, see the *Interface Module Reference* for your router.

- 100-Gigabit Ethernet PIC with CFP (model number: P1-PTX-2-100GE-CFP)—Supported in Junos OS Release 12.3R5, 13.2R3, 13.3R1, and later

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#).]

- **Support for three-phase delta AC power distribution unit (PDU), three-phase wye AC PDU, and the AC power supply module (PSM) on the PTX5000**—The PTX5000 now supports the three-phase delta AC PDU, three-phase wye AC PDU, and the AC PSM. You can use the **show chassis hardware**, **show chassis hardware-models**, and **show environment pdu** commands to view these hardware components.



NOTE: Mixing the DC and AC power components on the same chassis is not supported.

[See the [PTX5000 Packet Transport Switch Hardware Guide](#).]

- **Support for new 60 A DC power distribution unit (PDU) and 60 A DC power supply module (PSM) on the PTX5000 Packet Transport Router**—Each 60 A PDU has four dual-input input power trays. Each DC power cable, which you must provide, requires a 4-AWG cable lug minimum.



NOTE: Mixing the 60 A DC PDU and 120 A DC PDU is not supported except during upgrade. The 60 A DC PSM is supported only in the 60 A DC PDU.

[See the [PTX5000 Packet Transport Switch Hardware Guide](#).]

- **SFPP-10GE-ZR transceiver**—The PTX5000 supports the SFPP-10GE-ZR transceiver on the 10-Gigabit Ethernet PIC with SFP+ (model number: P1-PTX-24-10GE-SFPP). The SFPP-10GE-ZR transceiver supports the 10GBASE-Z optical interface standard. For more information, see the “Cables and Connectors” section in the PIC guide.

[See [10-Gigabit Ethernet 10GBASE Optical Interface Specifications](#) and the [PTX Series Interface Module Reference](#).]

- **CFP-100GBASE-ER4 and CFP-100GBASE-SR10 transceivers**—The PTX5000 supports the CFP-100GBASE-ER4 and CFP-100GBASE-SR10 transceivers on the 100-Gigabit Ethernet PIC with CFP (model number: P1-PTX-2-100GE-CFP). The CFP-100GBASE-ER4 transceiver supports the 100GBASE-ER4 optical interface standard. The CFP-100GBASE-SR10 transceiver supports the 100GBASE-SR10 optical interface standard. For more information, see the “Cables and connectors” section in the PIC guide.

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#) and [PTX Series Interface Module Reference](#).]

- **SFP-10G-ZR-OTN-XT (PTX Series)**—The SFP-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 SFP-10G-ZR-OTN-XT transceiver:

PTX Series:

- 10-Gigabit Ethernet PIC with SFP+ (model number: P1-PTX-24-10GE-SFP)—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-SFP)—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](#).]

Firewall Filters

- **DSCP and Traffic Class firewall filter match conditions on the loopback interface**—You can set the DSCP value for IPv4 traffic and Traffic Class value for IPv6 traffic in firewall filters that you apply to the loopback (**lo.0**) interface. To configure these forwarding class and DSCP values, apply an output filter to the **lo.0** interface.

[See [Standard Firewall Filter Match Conditions for IPv4 Traffic](#) and [Standard Firewall Filter Match Conditions for IPv6 Traffic](#).]

Interfaces and Chassis

- **Aggregated devices support increased to 64 links (PTX Series)**—This feature adds support for specifying up to 64 links for aggregated Ethernet devices. You set the number of links in the **maximum-links** statement at the **[edit chassis aggregated-devices]** hierarchy level.

[See [Configuring Junos OS for Supporting Aggregated Devices](#).]

- **WAN PHY support for the PTX Series**—You can configure WAN PHY mode on 10-Gigabit Ethernet interfaces on PTX Series Packet Transport Routers. Only the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number: P1-PTX-24-10G-W-SFP) supports WAN PHY mode. To configure WAN PHY mode, include the **framing wan-phy** statement at the **[edit interfaces xe-fpc/pic/port]** hierarchy level. The default framing mode for 10-Gigabit Ethernet interfaces is LAN PHY.

[See [PTX Series Packet Transport Switches Software Documentation](#).]

- **SFP-10G-CT50-ZR (PTX Series)**—The SFP-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:

PTX Series:

- 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](#).]

Network Management and Monitoring

- **Real-time performance monitoring (RPM) support on PTX Series Packet Transport Routers**—Real-time performance monitoring (RPM) allows the user to perform service-level monitoring. When RPM is configured on a device, the device calculates network performance based on packet response time, jitter, and packet loss. You can configure these values to be gathered by the following types of requests: Hypertext Transfer Protocol (HTTP) GET, Internet Control Message Protocol (ICMP), TCP, or UDP. The device gathers RPM statistics by sending out probes to a specified probe target, identified by an IP address or URL. When the target receives a probe, it generates responses that are received by the device. You set the probe options in the **probe** statement at the **[edit services rpm]** hierarchy level. You set the server to receive the probes in the **probe-server** statement at the **[edit services rpm]** hierarchy level. You use the **show services rpm probe-results** and **show services rpm history-results** commands to view the results of the most recent RPM probes.



NOTE: The PTX5000 supports up to 200 concurrent RPM sessions.

The following statements at the **[edit services rpm]** hierarchy are not supported:

- **twamp**
- **bgp logical-system**
- **probe owner test name destination-interface**
- **probe owner test name hardware-timestamp**
- **probe owner test name one-way-hardware-timestamp**
- **probe-server icmp**
- **probe-server tcp destination-interface**
- **probe-server udp destination-interface**

[See [Configuring Real-Time Performance Monitoring](#).]

User Interface and Configuration

- **Features from Junos OS 12.1X48 are now integrated in Junos OS Release 12.3 (PTX Series)**— All features supported in the Junos OS Release 12.1X48 release are supported in Junos OS Release 12.3.

[See [PTX Series Packet Transport Switch Software Documentation](#).]

Related Documentation

- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Router](#) ongoing on page 368
- [Errata and Changes in Documentation for Junos OS Release 12.3 for PTX Series Packet Transport Routers](#) on page 370
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router](#) ongoing on page 371
- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router](#) ongoing on page 374

Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Routers

- [Changes in Default Behavior and Syntax](#) on page 368

Changes in Default Behavior and Syntax

- [IPv6](#) on page 368
- [Junos OS XML API and Scripting](#) on page 368
- [Network Management and Monitoring](#) on page 369
- [Security](#) on page 369

IPv6

- **Change in automatically generated virtual-link-local-address for VRRP over IPv6**— The seventh byte in the automatically generated virtual-link-local-address for VRRP over IPv6 is 0x02. This change makes the VRRP over IPv6 feature in Junos OS 12.2R5, 12.3R3, 13.1R3, and later releases. Inoperable with Junos OS 12.2R1, 12.2 R2, 12.2 R3, 12.2R4, 12.3R1, 12.3R2, 13.1R1, and 13.3R2 releases if automatically generated virtual-link-local-address ID used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

Junos OS XML API and Scripting

- **IPv6 address text representation is stored internally and displayed in command output using lowercase**—Starting with Junos OS Release 11.1R1, IPv6 addresses are stored internally and displayed in the command output using lowercase. Scripts that match on an uppercase text representation of IPv6 addresses should be adjusted to either match on lowercase or perform case-insensitive matches.

- **<get-configuration> RPC with inherit="inherit" attribute returns correct time attributes for committed configuration**—In Junos OS Release 12.3R1, when you configured some interfaces using the interface-range configuration statement, if you later requested the committed configuration using the <get-configuration> RPC with the inherit="inherit" and database="committed" attributes, the device returned `junos:changed-localtime` and `junos:changed-seconds` in the RPC reply instead of `junos:commit-localtime` and `junos:commit-seconds`. This issue is fixed in Junos OS Release 12.3R2 and later releases so that the device returns the expected attributes in the RPC reply.

Network Management and Monitoring

- **New system log message indicating the difference in the Packet Forwarding Engine counter value (PTX Series)**—Effective in Junos OS Release 12.3R9, if the counter value of a Packet Forwarding Engine is reported to be less than its previous value, then the residual counter value is added to the newly reported value only for that specific counter. In that case, the CLI shows the `MIB2D_COUNTER DECREASING` system log message for that specific counter.

Security

- In all supported Junos OS releases, regular expressions can no longer be configured if they require more than 64MB of memory or more than 256 recursions for parsing.

This change in the behavior of Junos OS is in line with the Free BSD 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).

Related Documentation

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 364](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 370](#)
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 371](#)
- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 374](#)

Errata and Changes in Documentation for Junos OS Release 12.3 for PTX Series Packet Transport Routers

Errata

- The *OSPF Configuration Guide* incorrectly includes the **transmit-interval** statement at the `[edit protocols ospf area area interface interface-name]` hierarchy level. The **transmit-interval** statement at this hierarchy level is deprecated in the Junos OS command-line interface.

[*OSPF Configuration Guide*]

Network Management and Monitoring

- The documentation fails to clearly describe the characters that can be used for SNMPv3 authentication passwords. Besides numbers, uppercase letters, and lowercase letters, the following special characters are supported:

`.,/\<>;:'[]{}~!@#$%^*_+=-``

In addition, the following special characters are also supported, but you must enclose them within quotation marks ("") if you enter them on the CLI; if you use a Network Management System to enter the password, the quotation marks are not required:

`| & () ?`

The documentation also fails to clearly state that characters entered by simultaneously pressing the Ctrl key and additional keys are not supported. [PR/883083: This issue has been resolved]

- The syntax of the **filter-interfaces** statement in the *SNMP Configuration Statement* section is incorrect. The correct syntax is as follows:

```
filter-interfaces {
  all-internal-interfaces;
  interfaces interface-names{
    interface 1;
    interface 2;
  }
}
```

Related Documentation

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 364](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 368](#)
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 371](#)
- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 374](#)

Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Routers

The following issues currently exist in Juniper Networks PTX Series Packet Transport Routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

Class of Service (CoS)

- If a subset of the available queues is configured with transmit rates that add to 100 percent and the offered load to those queues exceeds 100 percent, the remaining queues can become starved for bandwidth. This situation leads to fatal egress TQ ASIC failures. This requires the FPC to be restarted to resume normal operation. [PR849914](#): This issue has been resolved.
- Without an explicit EXP rewrite configuration, the EXP bits of the top label is rewritten on the egress MPLS frame following the default EXP rewrite profile. [PR976481](#)

General Routing

- The PTX Series router is not supposed to generate pause frames even if it gets congestion. PTX might aggressively drop all ingress packets when queuing memory runs out. [PR873028](#)
- If a subset of the available queues is configured with transmit rates that add to 100 percent and the offered load to those queues exceeds 100 percent, the remaining queues can become starved for bandwidth. This situation leads to fatal egress TQ ASIC failures. This requires the FPC to be restarted to resume normal operation. Remote End connected to an interface of PTX 4x100GE PIC might observe interface flap with active MAC RF, even if the link switching is being done well within the interface hold-down timer configured at Remote End. [PR909635](#)
- Cflowd flow output shows 0 for Output ifIndex. [PR964745](#)
- In the P2MP environment with AE interface. When flapping interface continually, in rare condition, if unilist nexthop is added failure, the system does not clean it up properly, which leads to an FPC crash. [PR980622](#)
- When there is link/node protection/ECMP for RSVP/LDP transit or egress LSPs with huge scaling and continuous flapping of LSPs like auto-bandwidth case, traffic might get black-holed upon LSP re-optimizations. The issue would get triggered if the same unilist list-id (unilist list-id is a unique id for unilist nexthop) is allocated for two different unilist forwarding topologies. This situation arises when the unilist list-id wraps around after max value of 65535. After the wraparound, if there is long living list-id (which can be due to a node/link protected LSP that has not been re-optimized for long time), the Packet Forwarding Engine assigns the same list-id during allocation (upon other LSP re-optimizations) and this will trigger the issue as the new unilist will be directed to incorrect interface. [PR1043747](#)
- In LDP tunneling over single hop RSVP based LSP environment, after enabling "chained-composite-next-hop", the router may fail to create the chained composite next hops if the label value of VPN is equal with the label value of LDP. [PR1058146](#)

General Routing

- The PTX Series router is not supposed to generate pause frames even if it gets congestion. The behavior is to drop aggressively if it ever runs out of queuing memory. [PR873028](#)
- Remote End connected to an interface of PTX 4x100GE PIC might observe interface flap with active MAC RF, even if the link switching is being done well within the interface and hold-down timer configured at Remote End. [PR909635](#)
- This PR fixes the issue where cflowd flow output ifIndex being exported as 0. Unless there is a critical business need, we do not plan to backport the fix to releases earlier than 14.1. [PR964745](#)
- In the P2MP environment with AE interface. When flapping interface continually, in rare condition, if unilist nexthop is added failure, the system does not clean up it properly, this leads to FPC crash. [PR980622](#)
- On PTX5000, the packet drop is observed along with the parity error read from l3bnd_ht entry corresponding to certain addresses. With this SRAM parity error, ASIC will unconditionally drop the packet even PTX Series does not use l3bnd_ht during lookup. The parity check for l3bnd_ht lookup for PTX5000 will be disabled to avoid the SRAM parity error and packet drop as a workaround. We also add new log message to report the counter value change for slu.hw_err trap count - TL[<num>]: SLU hw error count <xxx> (prev count <yyy>). [PR1012513](#)
- When there is link/node protection/ECMP for RSVP/LDP transit or egress LSPs with huge scaling and continuous flapping of LSPs like auto-bandwidth case, traffic might get black-holed upon LSP re-optimizations. The issue would get triggered if the same unilist list-id (unilist list-id is a unique id for unilist nexthop) is allocated for 2 different unilist forwarding topology. This situation arises when the unilist list-id wraps around after max value of 65535. After the wraparound, if there is long living list-id (which can be due to some node/link protected LSP that has not been re-optimized for long time), the Packet Forwarding Engine assigns the same list-id during allocation (upon other LSP re-optimizations) and this will trigger the issue as the new unilist will be directed to incorrect interface. [PR1043747](#)
- The MIB "jnxDomMib" is not returning proper result from SNMP query on PTX3000/PTX5000. [PR1045804](#)
- In LDP tunneling over single hop RSVP based LSP environment, after enabling "chained-composite-next-hop", the router may fail to create the chained composite next hops if the label value of VPN is equal with the label value of LDP. [PR1058146](#)
- When a switchover is done from one Routing Engine to the other, in graceful-switchover redundancy mode, there is a brief period early in the transition of the SIB to online state, during which unsolicited (not corresponding to an attempt by the CPU to access the SIB via PCIe) errors are received at the downstream PCIe port on the CB to the SIB. The fix is to mute the generation of such errors during this brief period of the switchover. [PR1068237](#)
- On PTX Series platform, SIB temperature-threshold configuration cannot be changed for off-lined SIBs. [PR1097355](#)

Interfaces and Chassis

- During subscriber login/logout the following error log might occur on the device configured with GRES/NSR. /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV (pp0.1073751222). [PR1058958](#)

MPLS

- On PTX Series platforms, if link-protection and fast-reroute (FRR) are configured together, the routing protocol daemon (rpd) crashes and generates a core file. The root cause here is PTX Series platforms do NOT support both these configurations and assume to have only 2 gateways (one being primary and other being either bypass OR detour). [PR851047](#)
- When issue "traceroute mpls rsvp lsp-name" from the MPLS LSP ingress node, if there are PTX Series routers on the LSP path, PTX Series would not list correct downstream router's IP in the TLV of the response packet. [PR966986](#)
- When a PTX Series router is at the merge-point (MP) of a bypass LSP, if MPLS explicit-null has been enabled on the router, and the loopback interface has not been configured under protocol RSVP, the bypass LSP might not work correctly. [PR1012221](#)
- In MPLS traffic engineering with link or node protection enabled, after adding Shared Risk Link Group (SRLG) configuration, the bypass LSP might ignore the constraint and use an unexpected path. [PR1034636](#)
- When fast-reroute, node-link-protection or link-protection is configured, if a Shared Risk Link Group (SRLG) is associated with a link used by a LSP ingressing at a router, then on deleting the SRLG configuration from the router, the SRLG entry still stays in the SRLG table even after the re-optimization of this LSP. [PR1061988](#)

Network Management and Monitoring

- At some instances, the master Routing Engine fails to export some syslog facility messages to the external syslog server after Routing Engine switchover. [PR898030](#)
- While the router is rebooting and SNMP polling is not stopped, SNMP requests might land on mib2d process before Routing Engine protocol mastership is resolved, causing the mib2d process crash. [PR1114001](#)

Routing Protocols

- PTX Series routers accept traffic from remote sources to enable the remote source to be learned and advertised by MSDP so that receivers in other MSDP areas can join the source. To configure this feature, use the accept-remote-source configuration statement at the [edit protocols pim interface interface-name] hierarchy level. Note: On PTX Series routers requiring tunnel services, the PIM accept-remote-source configuration statement is not supported. [See accept-remote-source.] [PR891500](#)

Related Documentation

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 364](#)

- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Router](#) ongoing on page 368
- [Errata and Changes in Documentation for Junos OS Release 12.3 for PTX Series Packet Transport Routers](#) on page 370
- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router](#) ongoing on page 374

Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Routers

- [Current Release](#) on page 374
- [Previous Releases](#) on page 374

Current Release

The following issues are resolved in Juniper Networks PTX Series Packet Transport Routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

- **Interfaces and Chassis**

Interfaces and Chassis

- After removing a child link from AE bundle, in the output of "show interface <AE> detail", the packets count on the remaining child link spikes, then if you add back the previous child link, the count recovers to normal. [PR1091425](#): This issue has been resolved.

Previous Releases

- [Release 12.3R10](#) on page 374
- [Release 12.3R9](#) on page 376
- [Release 12.3R8](#) on page 376
- [Release 12.3R7](#) on page 376
- [Release 12.3R6](#) on page 377
- [Release 12.3R5](#) on page 378
- [Release 12.3R4](#) on page 380
- [Release 12.3R3](#) on page 384
- [Release 12.3R2](#) on page 384

Release 12.3R10

General Routing

- On PTX5000, the packet drop is observed along with the parity error read from l3bnd_ht entry corresponding to certain addresses. With this SRAM parity error, ASIC will unconditionally drop the packet even if the PTX5000 does not use l3bnd_ht during lookup. The parity check for l3bnd_ht lookup for the PTX5000 will be disabled to avoid the SRAM parity error and packet drop as a workaround. We also add new log message

to report the counter value change for slt.hw_err trap count - TL[<num>]: SLU hw error count <xxx> (prev count <yyy>). [PR1012513](#): This issue has been resolved.

- When the port on 24x 10GE(LWO) SFP+ (which never went link up since the PIC is onlined) is configured as CLI loopback, the ports will receive framing error until the interface gets physically linked up. (i.e. with real fiber instead of CLI loop). There would be no problem in normal use. This is only seen in self-loopback testing with CLI loopback. [PR1057364](#): This issue has been resolved.

Platform and Infrastructure

- Distributed protocol adjacencies (LFM/BFD/and so on) may experience a delay in keepalives transmission and/or processing due to a prolonged CPU usage on the FPC microkernel on PTX5000 Type 5-3D FPCs. The delay in keepalive transmission/processing may result in a mis-diagnosis of a link fault by the peer devices. The issue is seen several seconds after a Routing Engine mastership switch with NSR enabled, and the fault condition will clear after a couple of minutes. [PR849148](#): This issue has been resolved.

Routing Protocols

- When running Simple Network Management Protocol (SNMP) polling to specific IS-IS Management Information Base (MIB) with invalid variable, it will cause routing protocol process (rpd) crash. [PR1060485](#): This issue has been resolved.

Release 12.3R9

General Routing

- In the P2MP environment with AE interface. When flapping interface continually, in rare condition, if unilist nexthop is added failure, the system does not clean up properly and this leads to FPC crash. [PR980622](#): This issue has been resolved.

Release 12.3R8

General Routing

- SFP+-10G-ZR (part number = 740-052562) is not fully supported on P1-PTX-24-10G-W-SFPP pic. Inserting the optic on P1-PTX-24-10G-W-SFPP pic can cause FPC core file on the pic. [PR974783](#): This issue has been resolved.

Interfaces and Chassis

- On PTX Series platform, performing Routing Engine switchover might cause flabel (fabric token) to be out of sync between master Routing Engine and the backup Routing Engine, and results in an FPC crash. [PR981202](#): This issue has been resolved.

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](#): This issue has been resolved.

Release 12.3R7

Interfaces and Chassis

- Kernel crash may happen when a router running a junos install with the fix to PR 937774 is rebooted. This problem will not be observed during the upgrade to this junos install. It occurs late enough in the shutdown procedure that it shouldn't interfere with normal operation. [PR956691](#): This issue has been resolved.

- Sometimes the COSD generates a core file when add/delete child interface on the LAG bundle. [PR961119](#): This issue has been resolved.

Routing Policy and Firewall Filters

- On PTX Series platform, when a firewall filter has many terms, all the terms might not work correctly due to incorrect order of terms due to mis-programming. [PR973545](#): This issue has been resolved.

Release 12.3R6

General Routing

- If rpd ACK feature is enabled through command "indirect-next-hop-change-acknowledgements", when a route being added and a quick route change happens on the same route, high routing protocol process (rpd) CPU utilization might be seen and stays high (above 90%) until rpd is restarted. [PR953712](#): This issue has been resolved.

High Availability (HA) and Resiliency

- RPD on backup Routing Engine might hit out of memory condition and crash if BGP protocol experiences many flaps. [PR904721](#): This issue has been resolved.

Interfaces and Chassis

- On PTX Series Packet Transport Routers, a change to the 'oam lfm pdu holdtime' on an interface is not updated correctly. This results in an incorrect LFM state, which should be reported as Adjacency Lost. As a workaround, issue the **clear oam ethernet link-fault-management state** command from the CLI to correctly update the 'pdu holdtimer.' [PR792763](#): This issue has been resolved.
- Ethernet-CCC encapsulation allow both untagged and tagged packets to flow through. [PR807808](#): This issue has been resolved
- After an FPC comes online, CCL ERR counter registers need to be cleared for the link between the FPC and SIB. These counters get cleared after the link training and are 8B long. Also a 30-day history is maintained as well, that makes it 30 8B contiguous locations. Due to this bug only the first 30B out of the 8x30B get initialized leaving others with transient counter values from the SIB/FPC online event. This can lead to unintended CRC errors reported leading to FPC/SIB link alarms (link failure) if FPC/SIB uptime is between 4 - 30 days. Also since the AggrCRCErrCnt maintains a sum of a 30-day history which is a roll over counter, it can lead to a negative value indicated as a high error count. This has been fixed in the recent Junos OS releases to initialize all the 8x30B of the counter to correctly represent the current state of the link. [PR948185](#): This issue has been resolved.

IPv6

- PTX TLCHIP drops transit and host-bound packets containing the same source and destination IPs due to a protection mechanism built into TLCHIP. Such packets are counted as "Data error". This forces a change in loopback mode configuration on Ethernet interfaces. [PR934364](#): This issue has been resolved.

MPLS

- When Packet Forwarding Engine fast reroute (FRR) applications are in use (such as mpls facility backup, fast-reroute or loop free alternates), a primary path interface flap could be triggered due to Operation, Administration, and Maintenance (OAM) link failure detection or by Bidirectional Forwarding Detection (BFD). However, this interface flap might lead to a permanent use of the backup path, which means the original primary path could not be active again. [PR955231](#): This issue has been resolved.

Routing Protocols

- On PTX Series platform, after short protocol adjacencies flaps, rpd and kernel next-hops might not be in sync, resulting in equal-cost multi-path (ECMP) not working correctly. [PR911307](#): This issue has been resolved.

Software Installation and Upgrade

- In this case, since the high level package (i.e. jinstall) is signed, the underlying component packages are not required to be signed explicitly. However, the infrastructure was written such a way to display a warning message if the component package is not signed (that is, jpfe). [PR932974](#): This issue has been resolved.

User Interface and Configuration

- Configuration of an extended community such as: rt-import:*:* src-as:*:* fails because the wildcard is not allowed during the configuration validation process. [PR944400](#): This issue has been resolved.

Release 12.3R5

Class of Service (CoS)

- In an RSVP point-to-multipoint crossover/pass-through scenario, more than one sub-LSP can use the same PHOP and NHOP. If link protection is enabled in the above mentioned scenario, when a 'primary link up' event is immediately followed by a Path Tear message, disassociation of the routes/next hops are sequential in nature. When the routes/next hops disassociation is in progress, if a sub-LSP receives a path tear/PSB delete will lead to this core. [PR739375](#): This issue has been resolved.
- The ability to configure buffer size for SH queues was added. [PR770583](#): This issue has been resolved.

General Routing

- Processing of a neighbor advertisement can get into an infinite loop in the kernel, given a special set of events with regard to the Neighbor cache entry state and the incoming neighbor advertisement. [PR756656](#): This issue has been resolved.

High Availability (HA) and Resiliency

- FPC might randomly crash during unified ISSU. It will be kept offline after the unified ISSU period. [PR773960](#): This issue has been resolved.
- Distributed protocol adjacencies (LFM/BFD/ and so on) might experience a delay in keepalives transmission and/or processing due to prolonged CPU usage on the FPC microkernel on T4000 Type 5-3D FPCs. The delay in keepalive transmission/processing might result in a misdiagnosis of a link fault by the peer devices. The issue is seen several seconds after a Routing Engine mastership switch with NSR enabled. The fault condition will clear after a couple of minutes. [PR849148](#): This issue has been resolved.
- VPLS connections in MI state. In rare scenarios, the routing protocol process can fail to read the mesh-group information from kernel, which might result in the VPLS connections for that routing-instance to stay in MI (Mesh-Group ID not available) state. The workaround is to deactivate/activate the routing-instance. [PR892593](#): This issue has been resolved.

Interfaces and Chassis

- On PTX5000, when we issue the **debug cos halp ifd <ifd_index>** commands in a remote FPC, the FPC crashes. The core files could be seen by executing the CLI command **show system core-dumps**. [PR814935](#): This issue has been resolved.
- Kernel message 'Only parameters changed ...Sending to Slave side' is seen continuously on both master and backup Routing Engines. [PR820414](#): This issue has been resolved.
- When an FPC goes bad due to hardware failure and is stuck in a boot mode, it might affect Routing Engine-Packet Forwarding Engine communication for other FPCs since all private next-hop index space got depleted.

The following syslog entries are reported. /kernel: %KERN-4: Nexthop index allocation failed: private index space exhausted [PR831233](#): This issue has been resolved.

- This issue is triggered with affected Junos OS versions, under some special conditions, when only one end of an AE link sees LACP timeouts or there is intermittent LACP loss on the AE link. This trigger causes an issue only with these specific Junos OS versions (that do not have this PR fix) because of a change in default behavior where the AE member link was considered to be UP in any state other than DETACHED. The PR fix affects the following two changes 1) The default behavior has been restored to what it was before - which is, in any LACP state other than COLLECTING_DISTRIBUTING, the AE member is considered to be DOWN. 2) A fast-failover knob has been introduced that, if configured, causes the behavior to change such that in any LACP state other than DETACHED, the AE member is considered to be UP. Note that this issue is platform independent. [PR908059](#): This issue has been resolved.
- PTX Series and T4000's FPC crash can be triggered by a "Single Bit Error" (SBE) event after accessing a protected memory region, as indicated in the following log: "System

Exception: Illegal data access to protected memory." [PR919681](#): This issue has been resolved.

IPv6

- Setting OSPF overload via the configuration sets both the metric field in router LSAs as well as the te-metric field in opaque LSAs to 65535 or $2^{16}-1$. Since te-metric is a 32-bit field, it should be set to $2^{32}-1$. [PR797293](#): This issue has been resolved.
- Changing the domain-name doesn't reflect in DNS query unless a Commit full is done. This bug in management daemon (mgd) has been resolved by ensuring mgd propagates the new domain-name to file /var/etc/resolv.conf, so that this can be used for future DNS queries. [PR918552](#): This issue has been resolved.

Network Management and Monitoring

- Multiple SNMP queries for large volumes of information might cause Mib2d to grow in size and eventually create a core file. Mib2d will restart, possibly multiple times, but should recover by itself. [PR742186](#): This issue has been resolved.

Software Installation and Upgrade

- In this case, since the high level package (i.e. jinstall) is signed, the underlying component packages are not required to be signed explicitly. However the infra was written such a way to display warning message if the component package is not signed (i.e. jpfe). [PR932974](#): This issue has been resolved.

User Interface and Configuration

- Configuration mode access is locked after connection to router dropped. [PR745280](#): This issue has been resolved.

VPNs

- On the PTX Series, the routing protocol process (rpd) and the kernel might be out of sync regarding the forwarding nexthops after short protocol adjacencies flaps. [PR911307](#): This issue has been resolved.

Release 12.3R4

Class of Service (CoS)

- At Junos OS Release 12.1, excess-rate was an unsupported statement under [edit class-of-service] schedulers on PTX Series Packet Transport Routers. The excess-rate statement is now supported for scheduler configurations on PTX Series Packet Transport Routers. Subsequent versions of the *Junos OS Class of Service Configuration Guide* and other related documentation will be updated to reflect this change on the PTX Series. [PR738552](#): This issue has been resolved.
- Changing the preference on an LSP was considered a catastrophic event, tearing down the current path and then re-establishing a new one. This PR makes the preference change minor and only needs a new path to be re-signalled in a make-before-break manner. [PR897182](#): This issue has been resolved.

General Routing

- Processing of a neighbor advertisement can get into an infinite loop in the kernel, given a special set of events with regard to the Neighbor cache entry state and the incoming neighbor advertisement. [PR756656](#): This issue has been resolved.

High Availability (HA) and Resiliency

- In a situation where **prefix-export-limit** and NSR are configured together, when there are Routing Engine mastership switches, the IS-IS overload bit might be set after the NSR switchover. This issue is triggered due to inconsistent state between the master Routing Engine and the backup Routing Engine. As a workaround, disable **protocols isis prefix-export-limit**. [PR725478](#): This issue has been resolved.
- LACP status disagreement after Routing Engine switchover. [PR751745](#): This issue has been resolved.
- RPD on the backup Routing Engine might crash when it receives a malformed message from the master. This can occur at high scale with nonstop active routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare. The backup RPD will restart and the system will recover without intervention. [PR830057](#): This issue has been resolved.

Infrastructure

- The Junos OS kernel might crash because of a timing issue in the `ttymodem()` internal I/O processing routine. The crash can be triggered by simple remote access (such as Telnet or SSH) to the device. [PR755448](#): This issue has been resolved.

Interfaces and Chassis

- NSR switchover does not work with aggressive hello and hold-timers. This is a system limitation. Even the default 3 sec timer interval (for LAN interfaces) will not work. Cannot use such aggressive timers in scale scenario. Similar issue is seen in PR 719301 (though the scale numbers are different). To ensure zero traffic loss during GRES or NSR switchover ensure the following. 1. For non-"point-to-point" interfaces increase the hello-timer to something big around 30 seconds and holdtime to 90 seconds on all interfaces, on all routers. OR 2. Configure the interfaces as "point-to-point" under IS-IS on all routers. [PR772136](#): This issue has been resolved.
- When an FPC goes bad due to hardware failure and is stuck in a boot mode, it might affect Routing Engine-Packet Forwarding Engine communication for other FPCs since all private next-hop index space got depleted. The following syslog entries are reported. /kernel: %KERN-4: Nexthop index allocation failed: private index space exhausted. [PR831233](#): This issue has been resolved.
- Interrupt storm happened when press craft button with "craft-lockout". [PR870410](#): This issue has been resolved.

IPv6

- The core file is due to a null pointer dereference in ND6 code in the kernel and this bug was introduced when new feature HFRR was added. An IPv6 route that points to a discard next-hop will not require ND6 cache entry, and this check has been coded in to fix this issue. [PR755066](#): This issue has been resolved.

MPLS

- Some MPLS LSPs might run into the stuck status on the following triggers: 1. Aggressive link flapping on the LSP path 2. RSVP session got cleared from both ingress and egress within 1 second interval. The LSP remains down on the ingress router indefinitely because RSVP RESV messages were stuck on one of the transit LSRs, and were never sent to its upstream. [PR751729](#): This issue has been resolved.
- On PTX Series, in l2circuit/l2vpn or VPLS scenario, with chained composite-next-hop used, while performing certain l2circuit/l2vpn/vpls pings, routing protocol process (rpd) might crash and create a core file. When the issue happens, the following behavior could be observed: user@router> ping mpls l2circuit interface et-1/0/10.2 Info request to rpd timed out, exiting. [PR755489](#): This issue has been resolved.
- When a PTX Series Packet Transport Switch is a penultimate hop of one P2MP LSP branch and acts as a transit LSR on another branch for the same P2MP LSP, the MPLS packets going out from the penultimate hop branch might be tagged with incorrect Ethertype field. There is no workaround. [PR867246](#): This issue has been resolved.

Multicast

- RPD might generate a core file in some cases of incorrect next-hop deletion when adding and deleting multicast next hops. [PR702359](#): This issue has been resolved.

Platform and Infrastructure

- Following error messages in logs when performing commit synchronize mgd[1951]: UI_COMMIT: User 'regress' requested 'commit synchronize' operation (comment: none) rpd[2787]: junos_dfw_trans_purge:1445 Error "session" is not allocated. rpd[2787]: junos_dfw_session_close:1120 Error "session" is NULL or its socket has an invalid value before session close. The error is purely cosmetic and happens whenever anything being touched that results in RPD being notified, including just activating and deactivating interfaces. [PR737438](#): This issue has been resolved.

Routing Policy and Firewall Filters

- If you issue the **show krt next-hop** or **show krt iflist-next-hop** command, and if you later delete a route or the route is removed, an rpd core file might be created. [PR727014](#): This issue has been resolved.
- In scale LDP scenario (about 250 LDP neighbors), routing protocol process (rpd) crashed and created a core file while issuing CLI command **show ldp neighbor** because system tries to remove an active LDP adjacency incorrectly. The core files could be seen by executing CLI command **show system core-dumps**. If issue happens, the following logs could be seen: init: routing (PID 4305) terminated by signal number 6. Core dumped! init: routing (PID 32773) started. [PR747109](#): This issue has been resolved.

- When packet with TTL expiry is dropped on PTX Series as penultimate hop router, the following message can be seen: `/kernel: rnh_comp_output(): rnh 1578: no af 2 iff context for chaining; discarding packet` The issue is resolved in 12.1X48-D30 and later releases. [PR785366](#): This issue has been resolved.
- RPD on the backup Routing Engine might crash when it receives a malformed message from the master. This can occur at high scale with nonstop active routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare and backup RPD will restart and the system will recover without intervention. [PR830057](#): This issue has been resolved.
- Routing Engine might cause kernel panic. [PR851086](#): This issue has been resolved.
- Fixed continuous FPC crash on PTX Series when reject firewall action with non-zero reject code is present. [PR856473](#): This issue has been resolved.
- On PTX Series, while deactivating or activating a firewall filter that has tcp-flags in the match condition on a loopback interface (e.g. lo0.0), memory corruption could occur when the filter configuration is pushed to the Packet Forwarding Engine, or is removed from the Packet Forwarding Engine, causing all FPCs to crash and generate core files. The following is logged by the FPCs a few seconds prior to the crash: `fpc1 dfw_match_branch_db_destroy:77filter index 1, dfw 0x20bb2a90, match_branch_db not empty on filter delete fpc2 dfw_match_branch_db_destroy:77filter index 1, dfw 0x205a6340, match_branch_db not empty on filter delete fpc0 dfw_match_branch_db_destroy:77filter index 1, dfw 0x20471c38, match_branch_db not empty on filter delete` [PR874512](#): This issue has been resolved.

Routing Protocols

- RPD (routing protocol process) cored on receipt of RESV message with unexpected NHOP address. To avoid the crash, the solution is to drop RESV message with different NHOP IP address. Then the LSP will time out due to lack of refresh by RESV message and session reset. [PR887734](#): This issue has been resolved.
- When running the command `"monitor label-switched-path <lsp-name>"` on the PTX Series platform to display the real-time status of the specified RSVP label-switched path (LSP), the routing protocol process (rpd) might generate core files. [PR773439](#): This issue has been resolved.

Subscriber Access Management

- "Power Supply failure"/"Power Supply Removed" messages and SNMP trap occur hourly. [PR860223](#): This issue has been resolved.

User Interface and Configuration

- In scenario where telnet session is disconnected ungracefully while accessing "load merge terminal" prompt, problem can be exhibited with other CLI users unable to access configuration mode. [PR745280](#): This issue has been resolved.

Release 12.3R3***Interfaces and Chassis***

- Distributed protocol adjacencies (LFM/BFD/etc) might experience a delay in keepalives transmission and/or processing due to a prolonged CPU usage on the FPC microkernel on PTX5000 type 5-3D FPCs. The delay in keepalive transmission/processing can result in a mis-diagnosis of a link fault by the peer devices. The issue is seen several seconds after a Routing Engine mastership switch with nonstop active routing is enabled, and the fault condition will clear after a couple of minutes. [PR849148](#): This issue has been resolved.

Release 12.3R2***High Availability and Resiliency***

- On PTX Series Packet Transport Routers with nonstop active routing configured, if LDP is deleted or deactivated from the master Routing Engine, the Layer 2 circuit connections enter an incorrect encapsulation information state. The Layer 2 circuit connection transitions to the correct state when LDP is reactivated on the master Routing Engine. [PR799258](#): This issue has been resolved.
- Deletion of IPv6 addresses with a prefix of /128 from an interface can cause the Routing Engine to crash. [PR799755](#): This issue has been resolved.
- When the Flexible PIC Concentrator (FPC) restarted after performing a master Routing Engine switchover, the aggregate interface flag was set to "down". Any traffic that entered this FPC and traversed the equal-cost multipath (ECMP) to the aggregate interface was dropped. [PR809383](#): This issue has been resolved.

Infrastructure

- In an IPv6 scenario, when `ipv6-duplicate-addr-detection-transmits` is configured with a value of zero, IPv6 Neighbor Discovery might not function properly. [PR805837](#): This issue has been resolved.

Interfaces and Chassis

- FPC crashes on PTX Series Packet Transport Routers when reject firewall action with nonzero reject code is present. [PR856473](#): This issue has been resolved.

IPv6

- On PTX Series Packet Transport Routers, only 48k longest prefix match (LPM) routes are supported. If the limit of 48,000 LPM routes is exceeded, the kernel routing table (KRT) queue can be stuck with the error "Longest Prefix Match(LPM) route limit is exceeded." As a workaround, reduce the number of LPM routes because only 48000 LPM routes are supported. [PR801271](#): This issue has been resolved.

User Interfaces and Configurations

- PTX Series does not allow configuring buffer sizes on SH queues. [PR770583](#): This issue has been resolved

Related Documentation

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 364](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 368](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for PTX Series Packet Transport Routers on page 370](#)
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Router ongoing on page 371](#)

Junos OS Documentation and Release Notes

For a list of related Junos OS documentation, see <http://www.juniper.net/techpubs/software/junos/>.

If the information in the latest release notes differs from the information in the documentation, follow the *Junos OS Release Notes*.

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <http://www.juniper.net/books>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/customers/support/downloads/710059.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.

- JTAC Hours of Operation —The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool located at <https://tools.juniper.net/SerialNumberEntitlementSearch/>.

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/> .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at <http://www.juniper.net/support/requesting-support.html>.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to **support@juniper.net**. For documentation issues, fill out the bug report form located at <https://www.juniper.net/cgi-bin/docbugreport/>.

Revision History

15 October 2015—Revision 3, Junos OS 12.3 R11— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

9 October 2015—Revision 2, Junos OS 12.3 R11— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

8 October 2015—Revision 2, Junos OS 12.3 R11— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

1 October 2015—Revision 1, Junos OS 12.3 R11— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

26 August 2015—Revision 5, Junos OS 12.3 R10— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

12 August 2015—Revision 4, Junos OS 12.3 R10— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

16 July 2015—Revision 3, Junos OS 12.3 R10— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

9 July 2015—Revision 2, Junos OS 12.3 R10— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

2 July 2015—Revision 1, Junos OS 12.3 R10— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

5 March 2015—Revision 3, Junos OS 12.3 R9— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

26 February 2015—Revision 2, Junos OS 12.3 R9— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

19 February 2015—Revision 1, Junos OS 12.3 R9— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

9 October 2014—Revision 3, Junos OS 12.3 R8— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

1 October 2014—Revision 2, Junos OS 12.3 R8— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

25 September 2014—Revision 1, Junos OS 12.3 R8— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

1 July 2014—Revision 3, Junos OS 12.3 R7— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

24 June 2014—Revision 2, Junos OS 12.3 R7— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

17 June 2014—Revision 1, Junos OS 12.3 R7— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

2 April 2014—Revision 3, Junos OS 12.3 R6— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

25 March 2014—Revision 2, Junos OS 12.3 R6— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

18 March 2014—Revision 1, Junos OS 12.3 R6— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

17 January 2014—Revision 5, Junos OS 12.3 R5— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

14 January 2014—Revision 4, Junos OS 12.3 R5— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

10 January 2014—Revision 3, Junos OS 12.3 R5— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

7 January 2014—Revision 2, Junos OS 12.3 R5— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

23 December 2013—Revision 1, Junos OS 12.3 R5— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

21 November 2013—Revision 4, Junos OS 12.3 R4— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

03 October 2013—Revision 3, Junos OS 12.3 R4— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

25 September 2013—Revision 2, Junos OS 12.3 R4— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

18 September 2013—Revision 1, Junos OS 12.3 R4— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

23 August 2013—Revision 4, Junos OS 12.3 R3— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

12 July 2013—Revision 3, Junos OS 12.3 R3— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

27 June 2013—Revision 2, Junos OS 12.3 R3— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

19 June 2013—Revision 1, Junos OS 12.3 R3— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

29 May 2013—Revision 8, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

15 May 2013—Revision 7, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

07 May 2013—Revision 6, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

29 April 2013—Revision 5, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

09 April 2013—Revision 4, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

03 April 2013—Revision 3, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

29 March 2013—Revision 2, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

26 March 2013—Revision 1, Junos OS 12.3 R2— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

6 March 2013—Revision 4, Junos OS 12.3 R1— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

21 February 2013—Revision 3, Junos OS 12.3 R1— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

08 February 2013—Revision 2, Junos OS 12.3 R1— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

31 January 2013—Revision 1, Junos OS 12.3 R1— ACX Series, EX Series, PTX Series and the M Series, MX Series, and T Series.

Copyright © 2015, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.