

Junos[®] OS 12.3 Release Notes

Release 12.3R2
29 May 2013
Revision 8

These release notes accompany Release 12.3R2 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 Switches.

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

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

Contents

Junos OS Release Notes for ACX Series Routers	5
New Features in Junos OS Release 12.3 for ACX Series Routers	5
Firewall Filters	5
Layer 2 and Layer 3 Protocols	6
Routing Protocols	6
Time Division Multiplexing (TDM)	8
Timing and Synchronization	9
Changes in Default Behavior and Syntax in Junos OS Release 12.3 for ACX Series Routers	11
Junos OS XML API and Scripting	11
Errata and Changes in Documentation for Junos OS Release 12.3 for ACX Series Routers	11
Errata	11
Known Limitations in Junos OS Release 12.3 for ACX Series Routers	13
Class of Service	14
Firewall Filters	14
Interfaces and Chassis	15
MPLS Applications	15
Outstanding Issues in Junos OS Release 12.3 for ACX Series Routers	15
Virtual Private Network (VPN)	15

Resolved Issues in Junos OS Release 12.3 for ACX Series Routers	16
Interfaces and Chassis	16
Virtual Private Network (VPN)	16
Upgrade and Downgrade Instructions for Junos OS Release 12.3 for ACX	
Series Routers	16
Basic Procedure for Upgrading to Release 12.3	16
Upgrade and Downgrade Support Policy for Junos OS Releases	19
Downgrade from Release 12.3	19
Junos OS Release Notes for EX Series Switches	21
New Features in Junos OS Release 12.3 for EX Series Switches	21
Hardware	22
Enhanced Layer 2 Software (ELS) on EX9200 Switches	23
Access Control and Port Security	23
Class of Service (CoS)	24
Converged Networks (LAN and SAN)	24
Ethernet Switching and Spanning Trees	25
Firewall Filters and Routing Policy	26
High Availability	26
Infrastructure	27
Interfaces	28
IPv6	28
J-Web Interface	29
Layer 2 and Layer 3 Protocols	29
Management and RMON	30
MPLS	30
Virtual Chassis	30
Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX	
Series Switches	31
Infrastructure	31
Interfaces	31
Limitations in Junos OS Release 12.3 for EX Series Switches	32
Ethernet Switching and Spanning Trees	32
Firewall Filters	32
Hardware	32
High Availability	33
Infrastructure	34
Interfaces	35
J-Web Interface	35
Layer 2 and Layer 3 Protocols	36
Management and RMON	36
Virtual Chassis	37
Outstanding Issues in Junos OS Release 12.3 for EX Series Switches	38
Access Control and Port Security	38
Ethernet Switching and Spanning Trees	39
Firewall Filters	39
Hardware	39
High Availability	39
Infrastructure	39
Interfaces	40

J-Web Interface	41
Management and RMON	42
Multicast Protocols	43
Software Upgrade and Installation	43
Resolved Issues in Junos OS Release 12.3 for EX Series Switches	43
Issues Resolved in Release 12.3B1	44
Issues Resolved in Release 12.3B2	45
Issues Resolved in Release 12.3R2	54
Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches	56
Changes to Junos OS for EX Series Switches Documentation	57
Errata	57
Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches	58
Upgrade and Downgrade Support Policy for Junos OS Releases	58
Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations	58
Upgrading from Junos OS Release 10.4R3 or Later	59
Upgrading from Junos OS Release 10.4R2 or Earlier	60
Upgrading EX Series Switches Using NSSU	60
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers	64
New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers	64
Class of Service	64
Firewall Filters	69
Forwarding and Sampling	71
High Availability (HA) and Resiliency	71
Interfaces and Chassis	73
Junos OS XML API and Scripting	89
Layer 2 Features	89
MPLS Applications	91
Multicast	93
Power Management	96
Routing Protocols	96
Security	98
Subscriber Access Management	98
System Logging	109
User Interface and Configuration	110
VPLS	114
VPNs	116
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	118
Changes in Default Behavior and Syntax	118
Changes Planned for Future Releases	128
Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers	129
Outstanding Issues	129
Resolved Issues	138

Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers	148
Errata	148
Changes to the Junos OS Documentation Set	151
Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers	152
Basic Procedure for Upgrading to Release 12.3	152
Upgrade and Downgrade Support Policy for Junos OS Releases	155
Upgrading a Router with Redundant Routing Engines	156
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1	156
Upgrading the Software for a Routing Matrix	158
Upgrading Using ISSU	159
Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR	159
Downgrade from Release 12.3	160
Junos OS Release Notes for PTX Series Packet Transport Switches	162
New Features in Junos OS Release 12.3 for PTX Series Packet Transport Switches	162
Firewall Filters	163
Interfaces and Chassis	163
Network Management	164
Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Switches	165
Changes in Default Behavior and Syntax	165
Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches	165
Class of Service	166
High Availability (HA) and Resiliency	166
Interfaces and Chassis	166
Multicast	166
Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches	167
High Availability (HA) and Resiliency	167
Infrastructure	168
Interfaces and Chassis	168
MPLS Applications	168
User Interfaces and Configurations	168
Junos OS Documentation and Release Notes	169
Documentation Feedback	169
Requesting Technical Support	169
Revision History	171

Junos OS Release Notes for ACX Series Routers

- [New Features in Junos OS Release 12.3 for ACX Series Routers on page 5](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for ACX Series Routers on page 11](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for ACX Series Routers on page 11](#)
- [Known Limitations in Junos OS Release 12.3 for ACX Series Routers on page 13](#)
- [Outstanding Issues in Junos OS Release 12.3 for ACX Series Routers on page 15](#)
- [Resolved Issues in Junos OS Release 12.3 for ACX Series Routers on page 16](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for ACX Series Routers on page 16](#)

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 and the ACX2000 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:

- [Firewall Filters on page 5](#)
- [Layer 2 and Layer 3 Protocols on page 6](#)
- [Routing Protocols on page 6](#)
- [Time Division Multiplexing \(TDM\) on page 8](#)
- [Timing and Synchronization on page 9](#)

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]

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 slave 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

- [Junos OS XML API and Scripting on page 11](#)

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.

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

Errata

- [Firewall Filters on page 11](#)
- [Routing Protocols on page 13](#)

Firewall Filters

- Support for multifield classifiers is incorrectly omitted from the ACX 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 the 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 multified classifier, you can apply the multified 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]

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*]

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.]
- 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 Applications

- 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.]
- CFM MEP over L2VPN/L2 circuit on MPC no longer require the **no-control-word** command to be configured under L2VPN or l2-circuit hierarchy. [PR801746: This is a known limitation.]

Outstanding Issues in Junos OS Release 12.3 for ACX Series Routers

The following issue currently exists 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.

Virtual Private Network (VPN)

- The routing information base (RIB) groups do not work on the ACX Series routers, and this impacts the following scenarios:
 - Overlapping VPNs: When we have a common resource in one VPN that needs to be accessed by sites that are in different VPNs.
 - For leaking routes between inet.0 and VPN route table in scenarios where the VPN routers want to reach global internet routes.

[[PR736831](#)]

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.

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.]

Virtual Private Network (VPN)

- CFM MEP over L2VPN/L2 circuit on MPC no longer requires no-control-word to be configured under L2VPN or l2-circuit hierarchy. [[PR801746](#): This issue has been resolved]



NOTE: MPCs are not supported on ACX Series routers.

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 16](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 19](#)
- [Downgrade from Release 12.3 on page 19](#)

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 **bundle** 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 *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-12.3R21-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-12.3R21-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>.

Downgrade 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 21](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

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 [EX Series Switch Software Features Overview](#) and [EX Series Virtual Chassis Software Features Overview](#).

New features are described on the following pages:

- [Hardware on page 22](#)
- [Enhanced Layer 2 Software \(ELS\) on EX9200 Switches on page 23](#)
- [Access Control and Port Security on page 23](#)
- [Class of Service \(CoS\) on page 24](#)
- [Converged Networks \(LAN and SAN\) on page 24](#)
- [Ethernet Switching and Spanning Trees on page 25](#)
- [Firewall Filters and Routing Policy on page 26](#)
- [High Availability on page 26](#)
- [Infrastructure on page 27](#)
- [Interfaces on page 28](#)
- [IPv6 on page 28](#)
- [J-Web Interface on page 29](#)
- [Layer 2 and Layer 3 Protocols on page 29](#)
- [Management and RMON on page 30](#)
- [MPLS on page 30](#)
- [Virtual Chassis on page 30](#)

Hardware

Juniper Networks EX9200 Ethernet Switches—The EX9200 Programmable Switches support current and future SDN interfaces and protocols, offering the flexibility and scalability to increase business agility by simplifying the deployment and operation of cloud applications, server virtualization, and rich media collaboration tools across campuses and data centers. The EX9200 Ethernet 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 release for the EX9200 switches is Junos OS Release 12.3R2.

The EX9200 switches are modular systems that provide high availability and redundancy for all major hardware components, including Routing Engine (RE) 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 and their line cards 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 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.

- **EX9200 line cards**—Four line cards are available for EX9200 switches. 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-removable and hot-insertable.

The four EX9200 line cards 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](#).]

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 command-line interface (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.](#)]

[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.](#)]

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, lets you 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/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/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 hierarchy level [`edit protocols dot1x authenticator interface (all | interface-name)`]. By specifying the `lldp-med-bypass` statement, 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**—The EX4500 switch models that are Converged Enhanced Ethernet (CEE)–capable 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 IEEE Data Center Bridging Capability Exchange protocol (IEEE DCBX). These switches previously supported only DCBX version 1.01.

IEEE DCBX and DCBX version 1.01 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 56.]

- **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 56.]

Ethernet Switching and Spanning Trees

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

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

Support for all these switches is in addition to the previously supported EX Series switch platforms—EX2200, EX3200, and EX4200 switches. 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 cause significant 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 allows 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](#).]
- **Q-in-Q VLAN extended 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 on 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.

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, 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 allows 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](#).]

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 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

- **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. Previously, 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](#).]

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 allow 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). Previously, 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.

Related Documentation

- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)

- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

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.

Infrastructure

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

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. Previously, the frames with the Network Address Family subtype of the Chassis ID TLV having a value of 1 (IPv4) or 2 (IPv6) would be discarded, and LLDP would not detect the device as a neighbor.

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 21](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

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 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, BFD on IRB interfaces flaps if BFD is configured for subsecond timers. [PR/844951: This is a known software limitation.]
- On EX9200 switches, MVRP is not propagating the dynamically learned VLAN information that is associated with trunk interfaces. [PR/848600: 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. [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 the prestandard Cisco IP Phone 7960 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 RE**. However, no functionality is affected. [PR/477943: This is a known software limitation.]
- Routing between virtual-routing instances for local direct routes is not supported. [PR/490932: This is a known software limitation.]
- On EX4500 switches, the maintenance menu is not disabled even if you include the **lcd maintenance-menu disable** statement in the configuration. [PR/551546: This is a known software limitation.]
- When you enable the filter-id attribute on the RADIUS server for a particular client, none of the required 802.1X authentication rules are installed in the IPv6 database. Therefore, IPv6 traffic on the authenticated interface is not filtered; only IPv4 traffic is filtered on that interface. [PR/560381: This is a known software limitation.]
- On EX8200 switches, if OAM link fault management (LFM) is configured on a member of a VLAN on which Q-in-Q tunneling is also enabled, OAM PDUs cannot be transmitted to the Routing Engine. [PR/583053: This is a known software limitation.]
- When you reconfigure the 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. [PR/847772: 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.]

J-Web Interface

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

[PR/400814: This is a known software limitation.]
- In the J-Web interface, the Ethernet Switching Monitor page (Monitor > Switching > Ethernet Switching) might not display monitoring details if the switch has more than 13,000 MAC entries. [PR/425693: This is a known software limitation.]
- In the J-Web interface 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.]
- 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.]
- 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.]
- 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.]

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.]

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 21](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

Outstanding Issues in Junos OS Release 12.3 for EX Series Switches

The following are outstanding issues in Junos OS Release 12.3R2 for EX Series switches. The identifier following the description is the tracking number in our bug database.

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



NOTE: Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in [“Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers” on page 129](#).

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

- In some cases on an EX9200 switch on which the Multiple VLAN Registration Protocol (MVRP) has been configured, the switch might not propagate dynamic VLAN information about VLANs that are associated with a trunk port. [PR/840390]
- 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 EX9200 switches, BPDU blocking does not work on 40-gigabit interfaces. [PR/861293]

Firewall Filters

- On EX9200 switches, you cannot configure VLAN firewall filters for traffic leaving a VLAN. [PR/850520]

Hardware

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

High Availability

- On EX8200 Virtual Chassis, performing an NSSU upgrade from Junos OS Release 12.2R3.5 to Release 12.3R2 results in vmcores files being created on the backup XRE200 External Routing Engine, and vmcore and ksyncd core files being created on EX8200 Virtual Chassis linecards. As a workaround, use only supported NSSU upgrade paths. For information on supported NSSU upgrade paths, see [“Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches” on page 58](#). [PR/860872]

Infrastructure

- On EX8208 switches, when a line card that has no interface configurations and is not connected to any device is taken offline using the **request chassis fpc-slot slot-number offline** command, the Bidirectional Forwarding Detection process (bfd) starts and stops repeatedly. The same bfd process behavior occurs on a line card that is connected to a Layer 3 domain when another line card that is on the same switch and is connected to a Layer 2 domain is taken offline. [PR/548225]
- On EX9200 switches, the **arp** entry in the output of the **show arp** command points to a Layer 2 interface rather than to an IRB interface. [PR/781704]
- When you boot an EX2200 or EX3300 switch with Junos OS Release 12.2R1 or a later release, 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]
- 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 issue is due to the fact that, 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 peer PE router goes down, the LACP system ID will 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]

Interfaces

- If you configure a VLAN range on an access interface of an EX9200 switch, the Layer 2 Address Learning Daemon (L2ALD) might fail. As a workaround, do not configure VLAN ranges on EX9200 switch access interfaces. [PR/837608]
- 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 egresses the switch on a trunk port, the packet might be tagged twice. [PR/838597]
- 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 [enable|disable]** command. [PR/843743]
- 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 10-15 seconds to come up. [PR/847186]
- In some cases on an EX9200 switch, when the egress traffic on an interface is very high, host-originated packets (for example, OSPF hellos) that are to be sent on that interface are delayed inside the Packet Forwarding Engine. As a result, the peer device times out and the protocol flaps. [PR/848870]
- On EX9200 switches, ARP messages might not get resolved for a host added to a port that has been configured for Q-in-Q. [PR/849728]
- 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]
- On EX9200 switches, LACP flaps when GRES is done in a large configuration; in such configurations, the new master Routing Engine might take a long time to connect to a line card. As a workaround, increase the value of the redistribution timer to a value

such as 50 seconds or greater by using the **set routing-options ppm redistribution-timer seconds** command. [PR/850771]

- On EX9200 switches, Layer 3 traffic loss might occur for more than 10 seconds with MC-LAG Active-active when one of the provider-edge (PE) switches is rebooted. This LACP flap occurs because the LACP system ID changes on the other peer router. As a workaround, designate one of the PE switches as status-control-active, and configure the following on the PE switch: **set interfaces ae15 aggregated-ether-options mc-ae events iccp-peer-down prefer-status-control-active**. When this is configured on the PE switch, then when the peer goes down, the LACP system ID will not change. Note: To configure the **prefer-status-control-active** statement, you must configure the **status-control** statement as **active**. Do not configure **status-control** as **standby**. [PR/853964]

J-Web Interface

- On EX Series switches and on SRX3400, SRX3600, SRX5600, and SRX5800 devices, when you use the Microsoft Internet Explorer browser to open reports from the following pages in the J-Web interface, the reports open in the same browser session:
 - Files page (Maintain > Files)
 - History page (Maintain > Config Management > History)
 - Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
 - Static Routing page (Monitor > Routing > Route Information)
 - Support Information page (Maintain > Customer Support > Support Information)
 - View Events page (Monitor > Events and Alarms > View Events)

[PR/433883]

- In the J-Web interface, in the Port Security Configuration page, you are required to configure an action value when you configure MAC limit even though configuring an action value is not mandatory in the CLI. [PR/434836]
- In the J-Web interface on EX4200 switches; SRX100, SRX210, SRX240, and SRX650 Series Services Gateways; and all J Series devices, 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]

- When you open a J-Web session using HTTPS and then enter a username and password and click the **Login** button, the J-Web interface takes 20 seconds longer to launch and load the Dashboard page than it does if you use HTTP. [PR/549934]
- 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]

- 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]
- If you have a J-Web session open on a standalone EX Series switch, and if you then add another switch to create a Virtual Chassis, the chassis viewer might be aligned incorrectly on the dashboard. As a workaround, manually refresh the J-Web session. [PR/756711]
- 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]
- 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]
- 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]

Management and RMON

- In some cases, the hostname is not included in SNMP trap information received from EX9200 switches. [PR/840907]

- On EX9200 switches, the output of the **run monitor interface traffic** command might show more than 10Gbps for a 10-Gigabit Ethernet (xe) interface. [PR/858210]

Multicast Protocols

- On EX9200 switches, multicast traffic might be momentarily duplicated on an mrouter port 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]
- 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]
- On EX9200 switches, the mcsnoopd process creates multiple core files at `rt_nexthops_free` during GRES with Layer 3 multicast traffic. [PR/848732]
- 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]

Software Upgrade and Installation

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

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 21](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

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 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 [“Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers” on page 129](#).

- [Issues Resolved in Release 12.3B1 on page 44](#)
- [Issues Resolved in Release 12.3B2 on page 45](#)
- [Issues Resolved in Release 12.3R2 on page 54](#)

Issues Resolved in Release 12.3B1

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

- 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.]

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.]

High Availability

- 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.]

J-Web Interface

- 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.]

Management and RMON

- After a Routing Engine switchover, LACP and MIB process (mib2d) core files might be created. [PR/790966: 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.]

Virtual Chassis

- 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.]
- 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.]

Issues Resolved in Release 12.3B2

The following issues have been resolved since Junos OS Release 12.3B1. The identifier following the description is the tracking number in our bug database.

Access Control and Port Security

- For LLDP, the MAC/PHY configuration/status TLV values 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 manager process (authd) might reach a memory limit when there are approximately 250 users. As a workaround, reset the authd process when it reaches 85 percent of its RLIMIT_DATA value (that is, 85 percent of 130 MB). To check the amount of memory being used by the authd process, use the **show system processes extensive** operational mode command. [PR/783363: This issue has been resolved.]
- DHCP snooping might not 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.]
- 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.]
- 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.]

Ethernet Switching and Spanning Trees

- 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 chassis core file might be created if you remove an SFP transceiver from one of the module's interfaces. As the chassis 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.]

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 PFE internal support script and saving the output to a file. [PR/749974: This issue has been resolved.]
- You might see the following message in log files: **Kernel/ (COMPOSITE NEXT HOP) failed, err 6 (No Memory)**. [PR/751985: This issue has been resolved.]
- On EX3300 switches, if you configure more than 20 BGPv6 neighbor sessions, the command-line interface (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 daemon (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 `ppmd` core file might be created, and protocols that use `ppmd` 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 Apple Macintosh 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 daemon) 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.]
- If you issue the **set protocols rstp interface *logical-interface-name* edge** configuration command from the command-line interface (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 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 the Flash Memory 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

- After a nonstop software upgrade (NSSU) operation, OSPF might remain in the INIT state because the flooding entry is not programmed correctly. [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 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.]

Software Installation and Upgrade

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

- On EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6210 switches, a firewall filter with family set to ethernet-switching and configured for IPv4 will block 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 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.]

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 21](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

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 57](#)
- [Errata on page 57](#)

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.” 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](#).

Errata

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

- **request system software validate command**—The documentation for the **request system software validate** command incorrectly states that this command is supported on EX Series switches. This command is not supported on any EX Series switches. [This issue is being tracked by PR/803185.]
- 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 21 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 a match condition of 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.

Related Documentation

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 21](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)

- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for EX Series Switches on page 58](#)

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 58](#)
- [Upgrading to Junos OS Release 12.1R2 or Later with Existing VSTP Configurations on page 58](#)
- [Upgrading from Junos OS Release 10.4R3 or Later on page 59](#)
- [Upgrading from Junos OS Release 10.4R2 or Earlier on page 60](#)
- [Upgrading EX Series Switches Using NSSU on page 60](#)

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Upgrading to 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 60 for more information.

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

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

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

4. Install the new software package on the switch:

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

Replace *package* with one of the following paths:

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

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

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

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

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

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

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



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

5. Reboot the switch to start the new software:

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

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

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

6. After the reboot has 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 11.4 release notes. See *Upgrading from Junos OS Release 10.4R2 or Earlier* or *Upgrading from Junos OS Release 10.4R3 or Later* in the [Junos OS 11.4 Release Notes](#).

Upgrading EX Series Switches Using NSSU

You can use nonstop software upgrade (NSSU) to upgrade Junos OS releases on standalone EX6200 and EX8200 switches with dual Routing Engines and on EX3300, EX4200, EX4500, and EX8200 Virtual Chassis. For instructions on how to perform an upgrade using NSSU, see:

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

[Table 1 on page 61](#) details the switch platforms on which NSSU is supported and the required Junos OS release.

Table 1: Platform and Junos OS Upgrade Support for NSSU

Switch Platform	Upgrade from Junos OS Release x.x	Upgrade to Junos OS Release 12.3
EX3300 Virtual Chassis	Releases earlier than 12.2R1	Not supported
	12.2R1 or later	Supported
EX4200 Virtual Chassis, EX4500 Virtual Chassis, and mixed EX4200 and EX4500 Virtual Chassis	Releases earlier than 12.1R1	Not supported
	12.1R1 or later	Supported
	12.2R1 or later	Supported
EX6200 standalone switch	Releases earlier than 12.1R2	Not supported
	12.1R2 or later	Supported
	12.2R1 or later	Supported
EX8200 standalone switch	10.4R1 or 10.4R2	Not supported
	10.4R3 or later	Supported
	11.1R1 or later	Supported
	11.2R1 or later	Supported
	11.3R1 or later	Supported
	11.4R1 or later	Supported
	12.1R1 or later	Supported
	12.2R1 or later	Supported

Table 1: Platform and Junos OS Upgrade Support for NSSU (*continued*)

Switch Platform	Upgrade from Junos OS Release x.x	Upgrade to Junos OS Release 12.3
EX8200 Virtual Chassis	10.4R1 or later	Not supported
	11.1R1, 11.1R2, or 11.1R3	Not recommended
	11.1R4 or later	Supported
	11.2R1 or later	Supported
	11.3R1 or later	Supported
	11.4R1 or later	Supported
	12.1R1 or later	Supported
	12.2R1 or later	Supported



NOTE: NSSU upgrades to Junos OS Release 12.3R2—You can use the following releases for NSSU upgrades to Release 12.3R2:

- EX8200 standalone switches and EX8200 Virtual Chassis—Releases 11.4R7 or later and 12.1R5 or later
- EX3300 Virtual Chassis—Releases 12.2R3 or later and 12.3R1 or later
- EX4500 and EX4200 Virtual Chassis—Release 12.1R5 or later
- EX4550 Virtual Chassis—Releases 12.2R3 or later and 12.3R1 or later
- EX6200 switches—Releases 12.2R3 or later and 12.3R1 or later

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



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



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

**Related
Documentation**

- [New Features in Junos OS Release 12.3 for EX Series Switches on page 21](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for EX Series Switches on page 31](#)
- [Limitations in Junos OS Release 12.3 for EX Series Switches on page 32](#)
- [Outstanding Issues in Junos OS Release 12.3 for EX Series Switches on page 38](#)
- [Resolved Issues in Junos OS Release 12.3 for EX Series Switches on page 43](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.3 for EX Series Switches on page 56](#)

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 64](#)
- [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 118](#)
- [Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 129](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 148](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 152](#)

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.

- [Class of Service on page 64](#)
- [Firewall Filters on page 69](#)
- [Forwarding and Sampling on page 71](#)
- [High Availability \(HA\) and Resiliency on page 71](#)
- [Interfaces and Chassis on page 73](#)
- [Junos OS XML API and Scripting on page 89](#)
- [Layer 2 Features on page 89](#)
- [MPLS Applications on page 91](#)
- [Multicast on page 93](#)
- [Power Management on page 96](#)
- [Routing Protocols on page 96](#)
- [Security on page 98](#)
- [Subscriber Access Management on page 98](#)
- [System Logging on page 109](#)
- [User Interface and Configuration on page 110](#)
- [VPLS on page 114](#)
- [VPNs on page 116](#)

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 802.1p, 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-802.1p]** 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).

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. 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*]

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 look-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]

- **Nonstop active routing (NSR) support extended for MPLS and VPN protocols and applications using chained composite next hops (PTX Series Packet Transport Switches)**—Starting with Junos OS Release 12.1 R4, in addition to the support of chained composite next hops on PTX Series Packet Transport Switches, nonstop active routing (NSR) switchover is supported for the following MPLS and VPN protocols and applications:
 - Labeled BGP
 - Layer 2 VPNs



NOTE: Nonstop active routing is not supported for Layer 2 interworking (Layer 2 stitching).

- Layer 3 VPNs
- LDP

Nonstop active routing support for LDP includes:

- LDP unicast transit LSPs
- LDP egress LSPs for labeled internal BGP (IBGP) and external BGP (EBGP)
- LDP over RSVP transit LSPs
- LDP transit LSPs with indexed next hops
- LDP transit LSPs with unequal cost load balancing



NOTE: Nonstop active routing is not supported for LDP Point-to-Multipoint LSPs and LDP ingress LSPs.

- 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](#)..]

[[SNMP MIBs and Traps Reference](#)]

- **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 2: Maximum FRUs Supported on the MX2020 Router

FRU	Maximum Number
Routing Engines and CB	2
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**—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) fpc-slot <i>fpc-slot-number</i> mic-slot <i>mic-slot-number</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 <psm-slot-number>	request chassis spmb restart slot <i>spmb-slot-number</i>	
show chassis environment routing-engine <routing-engine-slot-number>		
show chassis environment sfb <sfb-slot-number>		

```
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
fpc-slot-number>
```

```
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>>
|<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 <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*.]

[*SNMP MIBs and Traps Reference*]

- **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 (PFEs)—**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*.

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](#).]

MPLS Applications

- **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 might 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. *[MPLS]*.

- 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 120000000 bps
- The **admitted-flow-rate** is 100000000 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 rl:rp[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*]

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 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](#).]

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 3 on page 98](#) 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 3: 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



NOTE: Although present in the code, the subscriber management features are not supported in Junos OS Release 12.3R2. Documentation for subscriber management features is included in the Junos OS Release 12.3 documentation set.

- **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 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 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 release 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]

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*]

- **Features from Junos 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](#) .]

- **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 118](#)
- [Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 129](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 148](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 152](#)

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 118](#)
- [Changes Planned for Future Releases on page 128](#)

Changes in Default Behavior and Syntax

The following are changes made to Junos OS default behavior and syntax.

- [Firewall Filters on page 119](#)
- [Interfaces and Chassis on page 119](#)
- [J-Web on page 122](#)
- [Junos OS XML API and Scripting on page 122](#)
- [Multiprotocol Label Switching \(MPLS\) on page 122](#)
- [MPLS Applications on page 122](#)
- [Multicast on page 123](#)
- [Network Address Translation \(NAT\) on page 123](#)
- [Network Management on page 124](#)
- [Routing Protocols on page 124](#)
- [Security on page 124](#)
- [Subscriber Access Management on page 125](#)
- [User Interface and Configuration on page 127](#)
- [VPNs on page 128](#)

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.

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 12.3R1, 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*]

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.

Multiprotocol Label Switching (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*]

MPLS Applications

- **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 this release, 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

Network Management

- 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]

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]

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*]

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 Router 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

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 13.3, BMP version 3 configurations are not backwards compatible with BMP version 1 configurations from earlier Junos OS releases.

[Routing Protocols]

Related Documentation

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 64](#)
- [Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 129](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 148](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 152](#)

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 152.

- [Outstanding Issues on page 129](#)
- [Resolved Issues on page 138](#)

Outstanding Issues

Class of Service (CoS)

- This cosmetic issue is specific of 3D linecards, based on Trio chipset. 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](#)
- 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](#)

Forwarding and Sampling

- Memory leak could happen to pfd, dcd, cosd, cfmd and dfcd if user frequently and repeatedly executes **show interface extensive** command from multiple telnet sessions under the following conditions:
 - Set screen-length value to small value. Screen length can be changed by the command **set cli screen-length <n>**.
 - 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](#)

General Routing

- The auto-complete feature is not working for the "show policy" command. [PR471332](#)
- If RPF and/or SCU is enabled, then any change to an ingress firewall table filter will trigger RPF/SCU reconfiguration for every prefix in the routing table. This might cause

transient high CPU utilization on the fpc which might result in SNMP stats request being time out. [PR777082](#)

- 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 get terminated. [PR817018](#)
- ADPC CORE with "in_checksum (pkt=0x7870f5aa <Address 0x7870f5aa out of bounds> [PR817318](#)
- When MX Series router running with DPC is upgraded by ISSU, some interfaces might show incorrect input packet/byte count. The incorrect count is also seen to the related interface MIB. The value will be a large number. Physical interface: xe-3/1/0, Enabled, Physical link is Up Interface index: 138, SNMP ifIndex: 5449, Generation: 141 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback: Local, Source filtering: Disabled, Flow control: Enabled Device flags : Present Running Loop-Detected Interface flags: SNMP-Traps Internal: 0x4000 Link flags : None CoS queues : 8 supported, 8 maximum usable queues Hold-times : Up 0 ms, Down 0 ms Current address: 00:24:dc:9c:7c:30, Hardware address: 00:24:dc:9c:7c:30 Last flapped : 2013-01-13 14:36:25 JST (02:07:52 ago) Statistics last cleared: Never Traffic statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps ~snip~ Logical interface xe-3/1/0.0 (Index 196614) (SNMP ifIndex 5450) (Generation 140) Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Traffic statistics: Input bytes : 3867797326912475 Output bytes : 0 Input packets: 15108583308733 Output packets: 0 Local statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Transit statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps Protocol inet, MTU: 1500, Generation: 160, Route table: 0 Flags: Sendbroadcast-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](#)

High Availability (HA) and Resiliency

- The JCS-1200 RE-JCS-1X2400-48G-S Routing Engine does not support an MTU size larger than 1500 bytes. [PR774380](#)
- The JCS-1200 RE-JCS-1X2400-48G-S Routing Engine does not support the Input and Output byte counters. When you use the "show interfaces em0 (or em1) extensive"

command, the Input and Output Bytes counter field is always displayed as zero (0). [PR783974](#)

- On TX routing systems with GRES enabled, if a mastership switch is being requested on an LCC whose backup Routing Engine's em0 interface has physically failed, this will cause all FPCs on that LCC to disconnect from the old master Routing Engine, but NOT reconnect to the new master Routing Engine. [PR799628](#)
- While doing ISSU, FEBs on router go to "offline" state with "READY Message Without Reconnect". As the result of failing to start ASICs, FEBs failed ISSU and became offline state. [PR861896](#)

Infrastructure

- Kernel fails to generate ICMP ttl expired when IP packet len is a multiple of 256. [PR829567](#)
- Delay in bringing ONLINE an FPC after it is inserted into the chassis. [PR853304](#)
- TCP is mistakenly enabling re-transmit timer for pure ACK's which is causing the FPC to reboot. [PR858489](#)

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](#)
- For Automatic Protection Switching (APS) on SONET/SDH interfaces, there are no operational mode commands that display the presence of APS mode mismatches. An APS mode mismatch occurs when one side is configured to use bidirectional mode, and the other side is configured to use unidirectional mode. [PR65800](#)
- On the Juniper Control System (JCS) platform, the control and management traffic for all Routing Engines share the same physical link on the same switch module. In rare cases, the physical link might become oversubscribed, causing the management connection to Protected System Domains (PSDs) to be dropped. [PR293126](#)
- Due to an incorrect calculation, memory heap utilization of a service PIC can go over 100% under the "show chassis pic" CLI command. There is no service impact. [PR737676](#)
- 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 is specific to the M120 hardware since there are two independent FRUs from where the PIC needs to be detached/attached. This IPC message 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](#)
- 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](#)

- With LSQ interface, the MLPPP fragments cannot use the egress queue 4 to 7 on the MLPPP member links. There is no workaround. [PR805307](#)
- Incorrect Detection timestamp in "show chassis fabric reachability" [PR811846](#)
- If we receive a MAC Move event or an L2 IFL change event, we don't immediately remove the next hops. The backup Routing Engine has to delete the NH first and then it gets deleted from the master. During this phase if pointer is stale to the NHs as in the pointer pointing to the NH is valid, but the NH has already been deleted, then you will run into this condition. [PR829093](#)
- A request (such as snmp query) for collecting the input IPv6 stats of the ae logical interface on an abc chipset is not working properly. [PR831811](#)
- When packet has to be forwarded over NH topology unilist->indirect-indexed and when the packet size is greater than egress interface MTU with DF set, then the following message might be logged and not sent the message back to source indicating "frag needed and DF set". fpc0 NH: Can not find ifl for nh 1048590 fpc0 NH(nh_get_mtu_iff) : get unilist mtu failed [PR844987](#)
- Whenever the tunnel interface -pe/-pd gets created using the MS-DPC instead of the MPC, it will not be able to process register messages. Because MPC and MS-DPC have different multicast architectures and they are incompatible, if chassis is configured in "enhanced-ip" mode this issue will be seen. Necessary changes have been made to the code so that these interfaces will not be created on MS-DPC. [PR853995](#)
- In a certain topology set up such that 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](#)
- Routing Engine reset with Fatal trap 12: page fault while in kernel mode. [PR855317](#)
- 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](#)
- Interface hold-time-down is not working properly for PIC type 10x10GE(LAN/WAN) SFPP. [PR859102](#)
- In a Junos OS IPv4/IPv6 dual-stack network, backup records are not created for services that are brought down by closing IPCP or IPV6CP. [PR790162](#)
- 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](#)
- The **set chassis fru-poweron-sequence** configuration statement is not supported for T4000 platform in 12.3R2 release. [PR868035](#)

J-Web

- For 2013 MR releases as well for all TOTs Jweb Home page copyright information should be changed from "2012" to "Copyright 2013." [PR845904](#)
- On the J-Web interface, HTTPS access might work with an invalid certificate. As a workaround, after you change the certificate, issue the "restart web-management" command to restart the J-Web interface. [PR700135](#)
- On MX Series routers, the options on the J-Web interface such as Access Concentrator, Idle Timeout, and Service Name for PPPoE logical interfaces are not supported. [PR493451](#)
- On M10i, MX240, and MX480 routers, when you use the Microsoft Internet Explorer browser to open reports from the following pages in the J-Web interface, the reports open in the same browser session:
 - Files page (Maintain > Files)
 - History page (Maintain > Config Management > History)
 - Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
 - Static Routing page (Monitor > Routing > Route Information)
 - Support Information page (Maintain > Customer Support > Support Information)
 - View Events page (Monitor > Events and Alarms > View Events)

[PR433883](#)

Layer 2 Ethernet Services

- 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 might be dropped by a dhcp server because the server-id in the client packet will not match that of the server. This will only happen once upon transition to the new method of building the server-id. Once that has happened, all future ISSU should work as before. [PR853329](#)

Multiprotocol Label Switching (MPLS)

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

Platform and Infrastructure

- Junos OS does not support dynamic ARP resolution on Ethernet interfaces that are designated for port mirroring. This causes the Packet Forwarding Engine to drop mirrored packets. As a workaround, configure the next-hop address as a static ARP entry by including the 'arp ip-address' statement at the [edit interfaces <interface-name>] hierarchy level. [PR237107](#)
- Following a large output jcs:get-input() will truncate the output after the ---(more)--- prompt and cause all following user input to be hidden. As a workaround, add the "| no-more" option when executing op scripts. [PR473816](#)
- Commit time warning is changed to trace message. [PR480082](#)
- On the process details page (Monitor > System View > Process Details) of the J-Web interface, there are multiple entries listed for a few processes that do not impact any functionality. [PR661704](#)
- On the JCS-1200 RE-JCS-1X2400-48G-S Routing Engine you can configure the MAC address on the external interfaces em0 and em1. You cannot configure the MAC address on fxp0 on the other Routing Engines supported on the JCS-1200. [PR770899](#)
- When changing configuration repeatedly, in rare conditions, some internal errors might cause the CLI process to hog 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 the user logs out from the router. [PR813673](#)
- Commit may fail, when a config object is deleted and re-added as transient change from a commit script. [PR814796](#)
- There's a problem going from 12.2 to 12.3 using ISSU. The blobs being created in 12.2 are using the newer format which is not compatible with 12.3 code. [PR818947](#)
- 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. [PR844795](#)

Routing Protocols

- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- After upgrade to Junos OS Release 10.4R9, the following messages are seen "Cancelling deferral pp0 index 131" These messages are not indicative of any problem and are only cosmetic [PR742534](#)
- The "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" to return the number of active prefixes from a BGP peer. So the new sequence of variables for the table is as follows: root@wfpro-mx4-c> show snmp mib walk jnxBgpM2PrefixCountersTable jnxBgpM2PrefixCountersAfi.0.1.1 = 1 jnxBgpM2PrefixCountersSafi.0.1.1 = 1 jnxBgpM2PrefixInPrefixes.0.1.1 = 0 jnxBgpM2PrefixInPrefixesAccepted.0.1.1 = 0 jnxBgpM2PrefixInPrefixesRejected.0.1.1 = 0 jnxBgpM2PrefixOutPrefixes.0.1.1 = 3 jnxBgpM2PrefixInPrefixesActive.0.1.1 = 0 [PR778189](#)
- This issue captures a change in behavior observed from previous releases. The adjacency hold down is taking longer than expected on passive interfaces and subsequently the issue disappears. This will not cause any functionality break since the functionality is restored eventually and seen only on passive interfaces immediately after ISSU. [PR780684](#)
- OSPF route will not be deleted from routing/forwarding table if configuration satisfies the following 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](#)
- Changing static route with qualified-next-hop and order option to next-hop option results in static route missing from route table. We need to restart routing process to see the route again [PR830634](#)
- IPv6 address syntax on Junos OS kernel and rpd logs is violation of RFC 5952 in chapter 4.2.3 [Choice in Placement of "::"]. When the length of the consecutive 16-bit 0 fields are equal, the first sequence of zero bits MUST be shortened. For example, 2002:db8:0:0:1:0:0:1 is logged in logs as 2002:db8:0:0:1::1, but 2002:db8::1:0:0:1 is correct representation. [PR840012](#)

- In Junos OS 11.4, rpd cores and snmpwalk fails at mplsL3VpnVrfRtinetCidrDestType when a route in a VRF has a datalink nexthop (such as when DHCP subscriber connects into a VRF). A warning message "Request failed: General error" will be seen. [PR840323](#)
- Upstream interface of multicast rpf not matching multicast route in Inter-AS PIM. [PR847370](#)
- 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](#)
- 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](#)

Services Applications

- When you specify a standard application at the [edit security idp idp-policy <policy-name> rulebase-ips rule <rule-name> match application] hierarchy level, IDP does not detect the attack on the nonstandard port (for example, junos:ftp on port 85). Whether it is a custom or predefined application, the application name does not matter. IDP simply looks at the protocol and port from the application definition. Only when traffic matches the protocol and port does IDP try to match or detect against the associated attack. [PR477748](#)
- Adaptive Service PIC (Service PIC II) is not supported by 12.1R4. Upon getting its configuration, the Service PIC reboots. [PR819833](#)
- Counter is not refreshing when polling the above mentioned OID via SNMP after RSP switchover. [PR829778](#)
- 11.4 Junos OS releases introduce the IKEv2 support and a stricter check on IKE/IPsec SAs proposal params. [PR843893](#)
- 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 allocate 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](#)
- MS-DPC might crash in certain scenarios when using CGNAT PBA and junos-rsh ALG. [PR862756](#)

Subscriber Access Management

- Subnet mask option is not returned to DHCP client when framed-ip-address is used with dhcp-local-server. [PR851589](#)
- Authd core experienced when multiple DHCP subscriber connection attempts require SRC for subscriber authentication. [PR862037](#)

User Interface and Configuration

- The logical router administrator can modify and delete master administrator-only configurations by performing local operations such as issuing the load override, load replace, and load update commands. [PR238991](#)
- Selecting the Monitor port for any port in the Chassis Viewer page takes the user to the common Port Monitoring page instead of the corresponding Monitoring page of the selected port. [PR446890](#)
- The J-Web interface allows the creation of duplicate term names in the Configure > Security > Filters > IPV4 Firewall Filters page. But the duplicate entry is not shown in the grid. There is no functionality impact on the J-Web interface. [PR574525](#)
- Using the IE7 browser, while deleting a user from the Configure > System Properties > User Management > Users page on the J-Web interface, the system is not showing warning message, whereas in the Firefox browser error messages are shown. [PR595932](#)
- If you access the J-Web interface using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR603669](#)
- When you configure the OSPF interface on the J-Web interface, the OSPF status changes to "disabled" from the default "enabled" status. As a workaround, instead of using the J-Web interface, use the following CLI commands to configure OSPF and the interface: - set protocols ospf area 10.1.2.5 area-range 12.26.0.0/16 - set protocols ospf area 10.1.2.5 interface pfh-0/0/0.16 [PR671052](#)
- On the J-Web interface, the next hop column in Monitor > Routing > Route Information displays only the interface address, and the corresponding IP address is missing. The title of the first column displays "static route address" instead of "Destination Address." [PR684552](#)
- Protected sections of the group hierarchy do not have their protection status displayed correctly and are not prevented from adding new elements into existing groups. [PR717527](#)
- "annotate" was not valid under firewall filter then hierarchy level and displayed "No valid completions", and led to the configuration not committing under "edit private" mode. [edit] liutao@mx480-a-re1# show | compare [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] + /* Don't process the packet here; it's IPv6, not IPv4. + * Accept it and have it be processed by the IPv6 ACL. */ accept; syntax error. liutao@mx480-a-re1# commit full [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] 'accept' outgoing comment does not match patch: [PR812111](#)
- On the J-Web interface, Configure > Routing> OSPF> Add> Interface Tab is showing only the following three interfaces by default: - pfh-0/0/0.16383 - lo0.0 - lo0.16385 To overcome this issue and to configure the desired interfaces to associated ospf

area-range, perform the following operation on the CLI: - set protocols ospf area 10.1.2.5 area-range 12.25.0.0/16 - set protocols ospf area 10.1.2.5 interface fe-0/3/1 [PR814171](#)

- On HTTPS service, the J-Web interface is not launching the chassis viewer page at IE7. [PR819717](#)

VPNs

- When you modify the frame-relay-tcc statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level of a Layer 2 VPN, the connection for the second logical interface might not come up. As a workaround, restart the chassis process (chassisd) or reboot the router. [PR32763](#)
- The deleted logical interfaces might not be freed due to references in MVPN. [PR851265](#)

Resolved Issues

Class of Service (CoS)

- 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 ifl 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 ifl <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, pfd 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 unilist 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 has multiple PFEs, 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 dump core. 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 ifl 68, is already present [15:21:23.207 LOG: Err] ASER0 SPI-4 XLR source core OOF did not go low in 20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 stop src train failed! [15:21:23.208 LOG: Err] ASER0 XLR SPI-4 sink core DPA incomplete in 20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 sink core init failed! [15:21:24.465

LOG: Err] ICHIP(0): SPI4 Stats Unexpected 2'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: [Jan1712:43:25.388 LOG: Err] SRCHIP(0): 1 Bad packets on p1 [Jan1712:43:25.389 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 (HA) 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 learnt 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 learnt 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 ifls)-->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 pfe 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's chassis-control interrupt storm may be falsely reported when a Field Replaceable Unit (FRU) is removed, inserted, or FPM button pushed. A FRU may not be recognized/booted, resulting in chassis operational failure. [PR823969](#): This issue has been resolved.
- IEEE 802.3 ah LFM stats counter "OAM current frame error event information" is not cleared correctly by cli operation. [PR827270](#): This issue has been resolved.
- If per-unit-scheduler is configured under a physical interface(ifd), and trying to delete this ifd and its sub-interfaces (ifl) 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 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 11.4R6.5 #0: 2012-11-28 21:57:12 UTC
builder@evenath.juniper.net:/volume/build/junos/11.4/release/11.4R6.5/obj-i
386/bsd/kernels/JUNIPER/kernel Build date: 2012-11-28 21:39:15 UTC Copyright (c
sysObjectID.0 = jnxProductNameMX480 sysUpTime.0 = 339594 sysContact.0 <
..... > dot3OutPauseFrames.942 = 0 dot3OutPauseFrames.943 = 0
dot3OutPauseFrames.953 = 0 dot3OutPauseFrames.954 = 0 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 unintended Packet Forwarding Engine (PFE). [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 triggers. *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 core. ?HW type of chassis/linecard/RE. "ALL" ?Suspected software feature combination. Multicast feature ?Describe if any behavior/ change to existing function - Handle the *G and S,G assert properly. [PR809338](#): This issue has been resolved.
- 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 may crash when it receives a malformed message from the master. This can occur at high scale with non-stop routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare and backup RPD will restart and the system will recover without intervention. [PR830057](#): 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 ISIS interfaces after configuration change the ISIS 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 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.
- 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 daemon (rpd) will generate flow-spec routes and install them in the routing table for these NLRIs; but these flow routes with invalid match conditions are rejected by dynamic firewall daemon (dfwd) from being added to the flowspec filters. When issue happens, the following errors could be seen:
krt_flow_trans_match_config: Failed defining match conditions
10.0.1.1,1.0.0.1,proto=6,dscp=81 krt_flow_trans_term_add: Failed adding term
10.0.1.1,1.0.0.1,proto=6,dscp=81 to filter 0x9504000 - Unknown error: 0
krt_flow_trans_filter_add: Failed sending transaction (ADD FILTER SINGLE TERM) for
filter 0x9504000 __flowspec_default_inet__ to add term 10.0.1.1,1.0.0.1,proto=6,dscp=81
- Invalid argument When the bgp peer withdraws these flow routes, they will only be
deleted but not freed, hence cause memory leak. [PR845039](#): This issue has been
resolved.
- In BGP scenario with multipath configured, if a static route which has table nexthop (such as inet.0) is configured in the same routing-instance as BGP, when an interconnect link between BGP peers is brought down or flapping, the corresponding BGP session takes 90 seconds to timeout. During this period routes received over the BGP session will stay there. For a multipath transit route received from both BGP sessions, initially both paths are resolved over the interconnect links directly. When one of the interconnect 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 daemon (rpd) crash and dump core. 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 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 may crash and core dumped when encountering unknown flow-type. Service will be impacted during the period. When issue happens, the following logs could be seen: `chassisd[1477]: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power off (jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @ 5/1/*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 8, jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 33404122)` `chassisd[1477]: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @ 5/1/*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 2, jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 192338924)` [PR834899](#): This issue has been resolved.
- In 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.

- /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 may 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 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 daemon (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 64](#)
- [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 118](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 148](#)

- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers](#) on page 152

Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers

Errata

Class of Service

- The *Example: Configuring Scheduling Modes on Aggregated Interfaces* topic fails to mention the following additional information regarding the parameters that are scaled for aggregated interface member links when the scheduler parameters are configured using scheduler maps:

Apart from transmit rate and buffer size that are scaled when the parameters are configured using scheduler maps, shaping rate is also scaled if you configure it in bits per second (bps). Shaping rate is not scaled if you configure it as a percentage of the available interface bandwidth.

[*Class of Service, Schedulers on Aggregated Ethernet and SONET/SDH Interfaces*]

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-Control Center Communications Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer ip-address** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer ip-address** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0.0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

[*Network Interfaces, Ethernet Interfaces*]

- The following additional information regarding the behavior of the `accept-data` statement for MC-LAG in an active-active bridge domain applies to the *Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview* topic:

For a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain and with VRRP configured over an integrated routing and bridging (IRB) interface, you must include the `accept-data` statement at the `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]` hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the `accept-data` option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The `accept-data` attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other layer 3 protocols and applications to work properly over multi-chassis aggregated Ethernet (mc-aeX) interfaces.

[Network Interfaces, Ethernet Interfaces]

- 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 Trio chips 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]

Routing Policy and Firewall Filters

- 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]*

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]

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]

Changes to the Junos OS Documentation Set

There are no changes to the Junos OS Documentation Set.

Related Documentation

- [New Features in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 64](#)
- [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 118](#)
- [Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 129](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 152](#)

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 152](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 155](#)
- [Upgrading a Router with Redundant Routing Engines on page 156](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 156](#)
- [Upgrading the Software for a Routing Matrix on page 158](#)
- [Upgrading Using ISSU on page 159](#)
- [Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR on page 159](#)
- [Downgrade from Release 12.3 on page 160](#)

Basic Procedure for Upgrading to Release 12.3

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

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



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



.....

NOTE: Before upgrading, back up the file system and the currently active Junos 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 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-12.3R21-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-12.3R21-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.



NOTE: Before you upgrade a router that you are using for voice traffic, you should monitor call traffic on each virtual BGF. Confirm that no emergency calls are active. When you have determined that no emergency calls are active, you can wait for nonemergency call traffic to drain as a result of graceful shutdown, or you can force a shutdown. For detailed information on how to monitor call traffic before upgrading, see the *Junos OS Multiplay Solutions Guide*.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

For more information 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 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*.

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

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

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

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

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

the standard reboot or ISSU procedures described in the other sections of these instructions.

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

To disable and reenable PIM:

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

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

[edit]
user@host# **set protocols pim nonstop-routing disable**
user@host# **activate protocols pim**
user@host# **commit**

Downgrade 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 64](#)
- [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 118](#)

- [Issues in Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 129](#)
- [Errata and Changes in Documentation for Junos OS Release 12.3 for M Series, MX Series, and T Series Routers on page 148](#)

Junos OS Release Notes for PTX Series Packet Transport Switches

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 162](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)
- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 167](#)

New Features in Junos OS Release 12.3 for PTX Series Packet Transport Switches

Powered by Junos OS, PTX Series Packet Transport Switches 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 Switches deliver several critical core functions, including industry-leading density and scalability, cost optimization, high availability, and network simplification. PTX Series Packet Transport Switches 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 Switches. 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](#).

- [Firewall Filters on page 163](#)
- [Interfaces and Chassis on page 163](#)
- [Network Management on page 164](#)

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

- **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 Switch**—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 [PTX5000 Packet Transport Switch Hardware Guide](#).]

- **Aggregated devices support increased to 64 links**—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](#).]

- **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 Packet Transport Switch PIC Guide](#).]

- **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 [the PTX Series Packet Transport Switch PIC Guide](#).]

- **WAN PHY support for PTX Series Switches**—You can configure WAN PHY mode on 10-Gigabit Ethernet interfaces on PTX Series Packet Transport Switches. Only the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number: P1-PTX-24-10G-W-SFPP) 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.

[*PTX Series Transport Switches Software Documentation*]

Network Management

- **Real-time performance monitoring (RPM) support on PTX Series Packet Transport Switches**—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](#).]

- Related Documentation**
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)
 - [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)
 - [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 167](#)

Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Switches

- [Changes in Default Behavior and Syntax on page 165](#)

Changes in Default Behavior and Syntax

- [Junos OS XML API and Scripting on page 165](#)

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.

- Related Documentation**
- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 162](#)
 - [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)
 - [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 167](#)

Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches

The following issues currently exist in Juniper Networks PTX Series Packet Transport Switches. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

Class of Service

- 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](#)]

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 RE and the back-up Routing Engine. As a workaround, disable **protocols isis prefix-export-limit**.

[[PR725478](#)]

- During a graceful Routing Engine switchover (GRES), the I2C bus on the FPC is temporarily unavailable, which might generate error messages. The FPC will recover from this condition when the GRES is complete.

[[PR743055](#)]

- 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](#)]

Interfaces and Chassis

- Having igmp enabled on the fxp0 interface can cause a discard next hop to be installed for 224/4 routes. To solve this issue, disable protocols on fxp0 and deactivate protocols (for example: IGMP, LDP). After that, activate the protocols again.

[[PR601619](#)]

Multicast

- RPD might generate a core file in some cases of wrong next-hop deletion when adding and deleting multicast next hops.

[[PR702359](#)]

Related Documentation

- [New Features in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 162](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)

- [Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 167](#)

Resolved Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches

The following issues are resolved in Juniper Networks PTX Series Packet Transport Switches. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

High Availability (HA) and Resiliency

- On PTX Series Packet Transport Switches 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 Switches when reject firewall action with nonzero reject code is present.

[[PR856473](#): This issue has been resolved.]

MPLS Applications

- On PTX Series Packet Transport Switches, only 48k longest prefix match (LPM) routes are supported. If the limit of 48,000 longest prefix match (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 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 Switches on page 162](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)
- [Outstanding Issues in Junos OS Release 12.3 for PTX Series Packet Transport Switches on page 165](#)

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

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 © 2013, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.