

System Management and Monitoring User Guide

Published
2024-12-18

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, 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.

System Management and Monitoring User Guide
Copyright © 2024 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

[About This Guide | v](#)

1

[Manage and Monitor](#)

[System Settings | 2](#)

[Specify the Physical Location of the Switch | 2](#)

[Modify the Default Time Zone on the Device | 3](#)

[Extend the Default Port Address Range | 4](#)

[Select a Fixed Source Address for Locally Generated TCP/IP Packets | 5](#)

[Rebooting and Halting a Device | 5](#)

[Hostnames | 7](#)

[Configure the Hostname of a Device in a Configuration Group | 7](#)

[Mapping the Hostname of the Switch to IP Addresses | 9](#)

[Example: Configuring the Name of the Switch, IP Address, and System ID | 9](#)

[Understanding and Configuring DNS | 10](#)

[DNS Overview | 10](#)

[Configure a DNS Name Server for Resolving Hostnames into Addresses | 11](#)

[ICMP Features | 15](#)

[Protocol Redirect Messages | 15](#)

[Pings | 17](#)

[Disable the Routing Engine Response to Multicast Ping Packets | 17](#)

[Disable Reporting IP Address and Timestamps in Ping Responses | 17](#)

[Source Quench Messages | 18](#)

[Time-to-Live \(TTL\) Expiration | 19](#)

[Rate Limit ICMP Traffic | 19](#)

[Rate Limit ICMP Error Messages | 20](#)

Alarms | 22

- System Alarms | 22

- Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types | 23

- System-Wide Alarms and Alarms for Each Interface Type | 23

System Troubleshooting | 26

- Saving Core Files Generated by Junos OS Processes | 26

- Viewing Core Files from Junos OS Processes | 27

Device Monitoring | 27

- Monitoring System Properties | 28

- Monitoring System Process Information | 31

- Monitor Interfaces | 32

- Other Tools to Configure and Monitor Devices Running Junos OS | 34

Passive Monitoring | 35

- Understanding Passive Monitoring | 35

- Example: Configuring Passive Monitoring | 36

- Requirements | 37

- Overview | 37

- Configuration | 38

- Verification | 40

- Sample Configuration for PTX10001-36MR, PTX10004, and PTX10008 Routers | 44

How to Locate a Device or Port Using the Chassis Beacon | 45

- Turning On the Chassis Beacon For the Default Interval | 46

- Turning On the Chassis Beacon For a Specified Interval | 47

Configuration Statements and Operational Commands

- Junos CLI Reference Overview | 50

About This Guide

Use this guide to manage and monitor Juniper switches with the Junos OS command line-interface.

1

CHAPTER

Manage and Monitor

System Settings | 2

Hostnames | 7

Understanding and Configuring DNS | 10

ICMP Features | 15

Alarms | 22

System Troubleshooting | 26

Device Monitoring | 27

Passive Monitoring | 35

How to Locate a Device or Port Using the Chassis Beacon | 45

System Settings

SUMMARY

Configure the system settings on your device.

IN THIS SECTION

- [Specify the Physical Location of the Switch | 2](#)
- [Modify the Default Time Zone on the Device | 3](#)
- [Extend the Default Port Address Range | 4](#)
- [Select a Fixed Source Address for Locally Generated TCP/IP Packets | 5](#)
- [Rebooting and Halting a Device | 5](#)

Specify the Physical Location of the Switch

To specify the physical location of the switch, specify the following options for the location statement at the [edit system] hierarchy level:

- altitude *feet*—Number of feet above sea level.
- building *name*—Name of the building, 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").
- country-code *code*—Two-letter country code.
- floor *number*—Floor in the building.
- hcoord *horizontal-coordinate*—Bellcore Horizontal Coordinate.
- lata *service-area*—Long-distance service area.
- latitude *degrees*—Latitude in degree format.
- longitude *degrees*—Longitude in degree format.
- npa-nxx *number*—First six digits of the phone number (area code and exchange).
- postal-code *postal-code*—Postal code.

- rack *number*—Rack number.
- vcoord *vertical-coordinate*—Bellcore Vertical Coordinate.

SEE ALSO

location (System)

[Example: Configuring the Name of the Switch, IP Address, and System ID | 9](#)

Modify the Default Time Zone on the Device

To modify the local time zone, include the `time-zone` statement at the `[edit system]` hierarchy level:

```
[edit]
user@host# set system time-zone (GMT hour-offset | time-zone)
```

You can use the `GMT hour-offset` option to set the time zone relative to UTC (GMT) time. By default, *hour-offset* is 0. You can configure this to be a value from -14 to +12.

You can also specify the *time-zone* value as a string such as PDT (Pacific Daylight Time) or WET (Western European Time), or specify the continent and major city.



NOTE: Junos OS complies with the POSIX time-zone standard, which is counter-intuitive to the way time zones are generally indicated relative to UTC. A time zone ahead of UTC (east of the Greenwich meridian) is commonly indicated as GMT +*n*; for example, the Central European Time (CET) zone is indicated as GMT +1. However, this is not true for POSIX time zone designations. POSIX indicates CET as GMT-1. If you include the `set system time-zone GMT+1` statement for a router in the CET zone, your router time will be set to one hour behind GMT, or two hours behind the actual CET time. For this reason, you might find it easier to use the POSIX time-zone strings, which you can list by entering `set system time-zone ?`.

For the time zone change to take effect for all processes running on the router or switch, you must reboot the router or switch.

The following example shows how to change the current time zone to `America/New_York`:

```
[edit]
user@host# set system time-zone America/New_York
[edit]
user@host# show
system {
    time-zone America/New_York;
}
```

SEE ALSO

[NTP Time Servers](#)

[Configure Time Zones](#)

Extend the Default Port Address Range

By default, the upper range of a port address is 5000. You can increase the range from which the port number can be selected to decrease the probability that someone can determine your port number.

- To configure Junos OS to extend the default port address range, include the `source-port` statement at the `[edit system internet-options]` hierarchy level:

```
[edit]
user@host# set system internet-options source-port upper-limit upper-limit
```

The statement `upper-limit upper-limit` is the upper limit of a source port address and can be a value from 5000 through 65,355.

SEE ALSO

[*internet-options*](#)

Select a Fixed Source Address for Locally Generated TCP/IP Packets

By default, the source address included in locally generated Transmission Control Protocol/IP (TCP/IP) packets, such as FTP traffic, and in User Datagram Protocol (UDP) and IP packets, such as Network Time Protocol (NTP) requests, is chosen as the local address for the interface on which the traffic is transmitted. This means that the local address chosen for packets to a particular destination might change from connection to connection based on the interface that the routing protocol has chosen to reach the destination when the connection is established. If multiple equal-cost next hops are present for a destination, locally generated packets use the `lo0` address as a source.

To configure the software to select a fixed address to use as the source for locally generated IP packets, include the *default-address-selection* statement at the `[edit system]` hierarchy level:

```
[edit system]
user@host# set system default-address-selection
```

If you include the *default-address-selection* statement in the configuration, the operating system chooses the system default address as the source for most locally generated IP packets. The default address is usually an address configured on the `lo0` loopback interface. For example, if you specified that SSH and telnet use a particular address, but you also have *default-address selection* configured, the system default address is used.

Rebooting and Halting a Device

To reboot the switch, issue the `request system reboot` command.

```
user@switch> request system reboot ?
Possible completions:
  <[Enter]>      Execute this command
  all-members    Reboot all virtual chassis members
  at             Time at which to perform the operation
  both-routing-engines Reboot both the Routing Engines
  fast-boot      Enable fast reboot
  hypervisor     Reboot Junos OS, host OS, and Hypervisor
  in            Number of minutes to delay before operation
  local         Reboot local virtual chassis member
  member        Reboot specific virtual chassis member (0..9)
  message       Message to display to all users
  other-routing-engine Reboot the other Routing Engine
```



```
| Pipe through a command
{master:0}
user@switch> request system reboot
Reboot the system ? [yes,no] (no) yes
Rebooting switch
```

**NOTE:**

- Not all options shown in the preceding command output are available on all devices. See the documentation for the [request system reboot](#) command for details about options.
- When you issue the `request system reboot hypervisor` command on QFX10000 switches, the reboot takes longer than a standard Junos OS reboot.

Similarly, to halt the switch, issue the `request system halt` command.



CAUTION: Before entering this command, you must have access to the switch's console port in order to bring up the Routing Engine.

```
user@switch> request system halt ?
Possible completions:
<[Enter]>      Execute this command
all-members    Halt all virtual chassis members
at             Time at which to perform the operation
backup-routing-engine  Halt backup Routing Engine
both-routing-engines  Halt both Routing Engines
in            Number of minutes to delay before operation
local         Halt local virtual chassis member
member        Halt specific virtual chassis member (0..9)
message       Message to display to all users
other-routing-engine  Halt other Routing Engine
|            Pipe through a command
```



NOTE: When you issue this command on an individual component in a QFabric system, you will receive a warning that says “Hardware-based members will halt, Virtual Junos Routing Engines will reboot.” If you want to halt only one member, use the `member` option. You cannot issue this command from the QFabric CLI.

Issuing the `request system halt` command on the switch halts the Routing Engine. To reboot a Routing Engine that has been halted, you must connect through the console.

SEE ALSO

[clear system reboot](#)[request system halt](#)[request system power-off](#)[Connecting a QFX Series Device to a Management Console](#)

RELATED DOCUMENTATION

[Disable Reporting IP Address and Timestamps in Ping Responses | 17](#)

Hostnames

IN THIS SECTION

- [Configure the Hostname of a Device in a Configuration Group | 7](#)
- [Mapping the Hostname of the Switch to IP Addresses | 9](#)
- [Example: Configuring the Name of the Switch, IP Address, and System ID | 9](#)

Configure the Hostname of a Device in a Configuration Group

The hostname of a device running Junos OS device is its identification. A network device must have its identity established to be accessible on the network. That is perhaps the most important reason to have a hostname, but a hostname also has other purposes.

Junos OS uses the configured hostname as part of the command prompt and to prepend log files and other accounting information. The hostname is also used anywhere else where knowing the device identity is important. For these reasons, we recommend that you provide hostnames that are descriptive and memorable.

You can configure the hostname at the `[edit system]` hierarchy level. Optionally, instead of configuring the hostname at the `[edit system]` hierarchy level, you can use a configuration group, as shown in this procedure. This is a recommended best practice for configuring the hostname, especially if the device has dual Routing Engines. This procedure uses groups called `re0` and `re1` as an example.



NOTE: If you configure hostnames that are longer than the CLI screen width, regardless of the terminal screen width setting, the commit operation occurs successfully. Even if the terminal screen width is less than the hostname length, the commit is successful.

To set the hostname by using a configuration group:

1. Include the `host-name` statement at the `[edit groups group-name system]` hierarchy level.

The *hostname* value must be less than 256 characters.

```
[edit groups group-name system]
host-name hostname;
```

For example:

```
[edit groups re0 system]
root@# set host-name san-jose-router0
```

```
[edit groups re1 system]
root@# set host-name san-jose-router1
```

2. If you used one or more configuration groups, apply the configuration groups, substituting the appropriate group names.

For example:

```
[edit]
user@host# set apply-groups [re0 re1]
```

3. Commit the changes.

```
[edit]
root@# commit
```


The hostname appears in the device CLI prompt.

```
san-jose-router0#
```

Mapping the Hostname of the Switch to IP Addresses

To map a hostname of a switch to one or more IP addresses, include the `inet` statement at the `[edit system static-host-mapping hostname]` hierarchy level:

```
[edit system]
static-host-mapping {
  hostname {
    inet [ addresses ];
    alias [ aliases ];
  }
}
```

hostname is the name specified by the `host-name` statement at the `[edit system]` hierarchy level.

For each host, you can specify one or more aliases.

SEE ALSO

[Configure a DNS Name Server for Resolving Hostnames into Addresses | 11](#)

Mapping the Name of the Router to IP Addresses

static-host-mapping

Example: Configuring the Name of the Switch, IP Address, and System ID

The following example shows how to configure the switch name, map the name to an IP address and alias, and configure a system identifier:

```
[edit]
user@switch# set system host-name switch1
[edit]
```



```
user@switch# set system static-host-mapping switch1 inet 192.168.1.77
[edit]
user@switch# set system static-host-mapping switch1 alias sj1
[edit]
user@switch# set system static-host-mapping switch1 sysid 1921.6800.1077
[edit]
user@switch# show
system {
    host-name switch-sj1;
    static-host-mapping {
        switch-sj1 {
            inet 192.168.1.77;
            alias sj1;
            sysid 1921.6800.1077;
        }
    }
}
```

Understanding and Configuring DNS

IN THIS SECTION

- [DNS Overview | 10](#)
- [Configure a DNS Name Server for Resolving Hostnames into Addresses | 11](#)

DNS Overview

IN THIS SECTION

- [DNS Components | 11](#)
- [DNS Server Caching | 11](#)

A Domain Name System (DNS) is a distributed hierarchical system that converts hostnames to IP addresses. The DNS is divided into sections called zones. Each zone has name servers that respond to the queries belonging to their zones.

DNS Components

DNS includes three main components:

- **DNS resolver:** Resides on the client side of the DNS. When a user sends a hostname request, the resolver sends a DNS query request to the name servers to request the hostname's IP address.
- **Name servers:** Processes the DNS query requests received from the DNS resolver and returns the IP address to the resolver.
- **Resource records:** Data elements that define the basic structure and content of the DNS.

DNS Server Caching

DNS name servers provide a hostname's IP address to users. The TTL field in the resource record defines the period for which DNS query results are cached. When the TTL value expires, the name server sends a fresh DNS query and updates the cache.

Configure a DNS Name Server for Resolving Hostnames into Addresses

You use Domain Name System (DNS) name servers to resolve hostnames to IP addresses.

Before you begin, configure your name servers with the hostname and an IP address for your Juniper Networks device. It does not matter which IP address you assign as the address of your device in the name server, as long as it is an address that reaches your device. Normally, you would use the management interface IP address, but you can choose the loopback interface IP address or a network interface IP address. You can even configure multiple addresses on the name server.

For redundancy, as a best practice, configure access to multiple name servers. You can configure a maximum of three name servers. The approach is similar to the way Web browsers resolve the names of a website to its network address.

You can use Junos OS to configure one or more domain names. The software uses these domain names to resolve hostnames that are not fully qualified (that is, hostnames for which the domain names are missing). Being able to configure domain names is convenient because you can use a hostname in configuring and operating the software without the need to reference the full domain name. After adding name server addresses and domain names to your configuration, you can use DNS resolvable hostnames in your configurations and commands instead of IP addresses.

Optionally, instead of configuring the name server at the [edit system] hierarchy level, you can use a configuration group, as shown in this procedure. This is a recommended best practice for configuring the name server.

You can route traffic between a management routing instance and a DNS name server. After you configure a routing instance at the [edit system name-server *server-ip-address*] hierarchy level, the name server becomes reachable through this routing instance.



NOTE: This management routing instance option is not supported for SRX Series Firewalls.

To enable a management routing instance for DNS, use the following configuration:

```
user@host# set system management-instance
user@host# set routing-instances mgmt_junos description description
user@host# set system name-server server-ip-address routing-instance mgmt_junos
```

If you've configured the name server using a configuration group, use the [edit groups *group-name* system name-server] hierarchy level, which is a recommended best practice for configuring the name server.

To configure the device to resolve hostnames into addresses:

1. Reference the IP addresses of your name servers.

```
[edit groups group-name system]
name-server {
    address;
}
```

The following example shows how to reference two name servers:

```
[edit groups global system]
user@host# set name-server 192.168.1.253
user@host# set name-server 192.168.1.254
user@host# show
name server {
    192.168.1.253;
    192.168.1.254;
}
```

2. (Optional) Configure the routing instance for DNS.

The following example shows how to configure the routing instance for one of the name servers:

```
[edit groups global system]
user@host# set name-server 192.168.1.253 routing-instance mgmt_junos
```

Remember to also configure the following:

- management-instance statement at the [edit system] hierarchy level
- routing-instance statement at the [edit routing-instances] hierarchy level

3. (Optional) Configure the name of the domain in which the device itself is located.

This is a good practice. The software then uses this configured domain name as the default domain name to append to hostnames that are not fully qualified.

```
[edit system]
domain-name domain-name;
```

The following example shows how to configure the domain name:

```
[edit groups global system]
user@host# set domain-name company.net
user@host# show
domain-name company.net;
```

4. (Optional) Configure a list of domains to be searched.

If your device can reach several different domains, you can configure a list of domains to be searched. Junos OS then uses this list to set an order in which it appends domain names when searching for the IP address of a host.

```
[edit groups global system]
domain-search [ domain-list ];
```

The domain list can contain up to six domain names, with a total of up to 256 characters.

The following example shows how to configure three domains to be searched. This example configures the software to search the company.net domain, next the domainone.net domain, and finally the domainonealternate.com domain when attempting to resolve unqualified hosts.

```
[edit groups global system]
domain-search [ company.net domainone.net domainonealternate.com ]
```

5. If you used a configuration group, apply the configuration group, replacing `global` with the appropriate group name.

```
[edit]
user@host# set apply-groups global
```

6. Commit the configuration.

```
user@host# commit
```

7. Verify the configuration.

If you've configured your name server with the hostname and an IP address for your device, you can issue the following commands to confirm that DNS is working and reachable. You can either use the configured hostname to confirm resolution to the IP address or use the IP address of your device to confirm resolution to the configured hostname.

```
user@host> show host host-name
user@host> show host host-ip-address
```

For example:

```
user@host> show host device.example.net
device.example.net
device.example.net has address 192.168.187.1
```

```
user@host> show host 192.168.187.1
10.187.168.192.in-addr.arpa domain name pointer device.example.net.
```


ICMP Features

SUMMARY

Use Internet Control Message Protocol (ICMP) features to diagnose network issues and check device reachability.

IN THIS SECTION

- [Protocol Redirect Messages | 15](#)
- [Pings | 17](#)
- [Source Quench Messages | 18](#)
- [Time-to-Live \(TTL\) Expiration | 19](#)
- [Rate Limit ICMP Traffic | 19](#)
- [Rate Limit ICMP Error Messages | 20](#)

Use [Feature Explorer](#) to confirm platform and release support for specific features.

Protocol Redirect Messages

IN THIS SECTION

- [Understanding Protocol Redirect Messages | 15](#)
- [Disable Protocol Redirect Messages | 16](#)

ICMP redirect, also known as protocol redirect, is a mechanism used by switches and routers to convey routing information to hosts. Devices use protocol redirect messages to notify the hosts on the same data link of the best route available for a given destination.

Understanding Protocol Redirect Messages

Protocol redirect messages inform a host to update its routing information and to send packets on an alternate route. Suppose a host tries to send a data packet through a switch S1 and S1 sends the data packet to another switch, S2. Also, suppose that a direct path from the host to S2 is available (that is, the host and S2 are on the same Ethernet segment). S1 then sends a protocol redirect message to inform the host that the best route for the destination is the direct route to S2. The host should then send packets

directly to S2 instead of sending them through S1. S2 still sends the original packet that it received from S1 to the intended destination.

Refer to [RFC-1122](#) and [RFC-4861](#) for more details on protocol redirecting.



NOTE:

- Switches do not send protocol redirect messages if the data packet contains routing information.
- All EX series switches support sending protocol redirect messages for both IPv4 and IPv6 traffic.

Disable Protocol Redirect Messages

By default, devices send protocol redirect messages for both IPv4 and IPv6 traffic. For security reasons, you may want to disable the device from sending protocol redirect messages.

To disable protocol redirect messages for the entire device, include the [no-redirects](#) or `no-redirects-ipv6` statement at the `[edit system]` hierarchy level.

- For IPv4 traffic:

```
[edit system]
user@host# set no-redirects
```

- For IPv6 traffic:

```
[edit system]
user@host# set no-redirects-ipv6
```

To re-enable the sending of redirect messages on the device, delete the `no-redirects` statement (for IPv4 traffic) or the `no-redirects-ipv6` statement (for IPv6 traffic) from the configuration.

To disable protocol redirect messages on a per-interface basis, include the [no-redirects](#) statement at the `[edit interfaces interface-name unit logical-unit-number family family]` hierarchy level.

- For IPv4 traffic:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family inet no-redirects
```


- For IPv6 traffic:

```
[edit interfaces interface-name unit logical-unit-number]  
user@host# set family inet6 no-redirects
```

Pings

IN THIS SECTION

- [Disable the Routing Engine Response to Multicast Ping Packets | 17](#)
- [Disable Reporting IP Address and Timestamps in Ping Responses | 17](#)

Pings use ICMP. A successful ping is when a device sends an ICMP echo request to a target and the target responds with an ICMP echo reply. However, there might be situations where you do not want your device to respond to ping requests.

Disable the Routing Engine Response to Multicast Ping Packets

By default, the Routing Engine responds to ICMP echo requests sent to multicast group addresses. By configuring the Routing Engine to ignore multicast ping packets, you can prevent unauthorized persons from discovering the list of provider edge (PE) devices in the network.

To disable the Routing Engine from responding to these ICMP echo requests, include the `no-multicast-echo` statement at the `[edit system]` hierarchy level:

```
[edit system]  
user@host# set no-multicast-echo
```

Disable Reporting IP Address and Timestamps in Ping Responses

When you issue the `ping` command with the `record-route` option, the Routing Engine displays the path of the ICMP echo request packets and the timestamps in the ICMP echo responses by default. By configuring the `no-ping-record-route` and `no-ping-timestamp` options, you can prevent unauthorized persons from discovering information about the provider edge (PE) device and its loopback address.

You can configure the Routing Engine to disable the setting of the `record-route` option in the IP header of the ping request packets. Disabling the `record-route` option prevents the Routing Engine from recording and displaying the path of the ICMP echo request packets in the response.

To configure the Routing Engine to disable the setting of the `record route` option, include the `no-ping-record-route` statement at the `[edit system]` hierarchy level:

```
[edit system]
user@host# set no-ping-record-route
```

To disable the reporting of timestamps in the ICMP echo responses, include the `no-ping-time-stamp` option at the `[edit system]` hierarchy level:

```
[edit system]
user@host# set no-ping-time-stamp
```

Source Quench Messages

When a device is receiving too many or undesired datagrams, it can send a source quench message to the originating device. The source quench message signals the originating device to reduce the amount of traffic it is sending.

By default, the device reacts to ICMP source quench messages. To ignore ICMP source quench messages, include the `no-source-quench` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
no-source-quench;
```

To stop ignoring ICMP source quench messages, use the `source-quench` statement:

```
[edit system internet-options]
source-quench;
```


Time-to-Live (TTL) Expiration

The time-to-live (TTL) value in a packet header determines how long the packet remains traveling through the network. The TTL value decrements with each device (or hop) the packet travels through. When a device receives a packet with a TTL value of 0, it discards the packet. The TTL expiry message is sent using ICMP.

You can configure your device to use an IPv4 address as the source address for ICMP time-to-live (TTL) expiry error messages. This means you can configure the loopback address as the source address in response to ICMP error packets. Doing this is useful when you cannot use the device address for traceroute purposes because you have duplicate IPv4 addresses in your network.

The source address must be an IPv4 address. To specify the source address, use the `ttl-expired-source-address source-address` option at the `[edit system icmp (System)]` hierarchy level:

```
[edit system icmp]
user@host# set ttl-expired-source-address source-address
```

This configuration only applies to ICMP TTL expiry messages. Other ICMP error reply messages continue to use the address of the ingress interface as the source address.

Rate Limit ICMP Traffic

To limit the rate at which ICMPv4 or ICMPv6 messages can be generated by the Routing Engine and sent to the Routing Engine, include the appropriate rate limiting statement at the `[edit system internet-options]` hierarchy level.

- For IPv4:

```
[edit system internet-options]
icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate
```

- For IPv6:

```
[edit system internet-options]
icmpv6-rate-limit bucket-size bucket-size packet-rate packet-rate
```


Rate Limit ICMP Error Messages

IN THIS SECTION

- [Why to Rate Limit ICMPv4 and ICMPv6 Error Messages | 20](#)
- [How to Rate Limit ICMPv4 and ICMPv6 Error Messages | 21](#)

By default, ICMP error messages for non-TTL-expired IPv4 and IPv6 packets are generated at the rate of 1 packet per second (pps). You can adjust this rate to a value that you decide provides sufficient information for your network without causing network congestion.



NOTE: For TTL-expired IPv4 or IPv6 packets, the rate for ICMP error messages is not configurable. It is fixed at 500 pps.

Why to Rate Limit ICMPv4 and ICMPv6 Error Messages

An example use case for adjusting the rate limit is a data center providing web services. Suppose this data center has many servers on the network that use jumbo frames with an MTU of 9100 bytes when they communicate to hosts over the Internet. These other hosts require an MTU of 1500 bytes. Unless maximum segment size (MSS) is enforced on both sides of the connection, a server might reply with a packet that is too large to be transmitted across the Internet without being fragmented when it reaches the edge router in the data center.

Because TCP/IP implementations often have Path MTU Discovery enabled by default with the `do not fragment` bit set to 1, a transit device will drop a packet that is too big rather than fragmenting it. The device will return an ICMP error message indicating the destination was unreachable because the packet was too big. The message will also provide the MTU that is required where the error occurred. The sending host should adjust the sending MSS for that connection and resend the data in smaller packet sizes to avoid the fragmentation issue.

At high core interface speeds, the default rate limit of 1 pps for the error messages may not be enough to notify all the hosts when there are many hosts in the network that require this service. The consequence is that outbound packets are silently dropped. This action can trigger additional retransmissions or back-off behaviors, depending on the volume of requests that the data center edge router is handling on each core-facing interface.

In this situation, you can increase the rate limit to enable a higher volume of oversized packets to reach the sending hosts. (Adding more core-facing interfaces can also help resolve the problem.)

How to Rate Limit ICMPv4 and ICMPv6 Error Messages

Although you configure the rate limit at the `[edit chassis]` hierarchy level, it is not a chassis-wide limit. Instead, the rate limit applies per interface family. This means, for example, that multiple physical interfaces configured with `family inet` can simultaneously generate the ICMP error messages at the configured rate.



NOTE: This rate limit takes effect only for traffic that lasts 10 seconds or longer. The rate limit is not applied to traffic with a shorter duration, such as 5 seconds or 9 seconds.

- To configure the rate limit for ICMPv4, use the `icmp` statement:

```
[edit chassis]
user@host# set icmp rate-limit rate-limit
```

Starting in Junos OS Release 19.1R1, the maximum rate increased from 50 pps to 1000 pps.

- To configure the rate limit for ICMPv6, use the `icmp6` statement:

```
[edit chassis]
user@host# set icmp6 rate-limit rate-limit
```

You must also consider that the rate limit value can interact with your DDoS protection configuration. The default bandwidth value for exceptioned packets that exceed the MTU is 250 pps. DDoS protection flags a violation when the number of packets exceeds that value. If you set the rate limit higher than the current `mtu-exceeded` bandwidth value, then you must configure the bandwidth value to match the rate limit.

For example, suppose you set the ICMP rate limit to 300 pps:

```
user@host# set chassis icmp rate-limit 300
```

You must configure the DDoS protection `mtu-exceeded` [bandwidth](#) to match that value.

```
user@host# set system ddos-protection protocols exceptions mtu-exceeded bandwidth 300
```


Alarms

IN THIS SECTION

- [System Alarms | 22](#)
- [Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types | 23](#)
- [System-Wide Alarms and Alarms for Each Interface Type | 23](#)

System Alarms

Switches provide predefined system alarms that can be triggered by a missing rescue configuration, failure to install a license for a licensed software feature, or high disk usage. You can display alarm messages by issuing the `show system alarms operational mode` command.

For example: The switch might trigger an alarm when disk usage in the `/var` partition exceeds 75 percent. A usage level between 76 and 90 percent indicates high usage and raises a minor alarm condition, whereas a usage level above 90 percent indicates that the partition is full and raises a major alarm condition.

The following sample output shows the system alarm messages that are displayed when disk usage is exceeded on the switch.

```
user@host> show system alarms


4 alarms currently active
Alarm time          Class  Description
2013-10-08 20:08:20 UTC Minor  RE 0 /var partition usage is high
2013-10-08 20:08:20 UTC Major  RE 0 /var partition is full
2013-10-08 20:08:08 UTC Minor  FPC 1 /var partition usage is high
2013-10-08 20:08:08 UTC Major  FPC 1 /var partition is full
```



BEST PRACTICE: We recommend that you regularly request a system file storage cleanup to optimize the performance of the switch and prevent generating system alarms.

Configuring Junos OS to Determine Conditions That Trigger Alarms on Different Interface Types

For the different types of PICs, you can configure which conditions trigger alarms and whether they trigger a red or yellow alarm. Red alarm conditions light the **RED ALARM** LED and trigger an audible alarm if one is connected. Yellow alarm conditions light the **YELLOW ALARM** LED and trigger an audible alarm if one is connected.

**NOTE:** By default, any failure condition on the integrated-services interface (Adaptive Services PIC) triggers a red alarm.

To configure conditions that trigger alarms and that can occur on any interface of the specified type, include the `alarm` statement at the `[edit chassis]` hierarchy level.

```
[edit chassis]
alarm {
  interface-type {
    alarm-name (red | yellow | ignore);
  }
}
```

alarm-name is the name of an alarm.

System-Wide Alarms and Alarms for Each Interface Type

Table 1 on page 23 lists the system-wide alarms and the alarms for each interface type.

Table 1: Configurable PIC Alarm Conditions

| Interface/System | Alarm Condition | Configuration Option |
|-------------------|------------------------------|----------------------|
| SONET/SDH and ATM | Link alarm indication signal | ais-l |
| | Path alarm indication signal | ais-p |
| | Signal degrade (SD) | ber-sd |

Table 1: Configurable PIC Alarm Conditions *(Continued)*

| Interface/System | Alarm Condition | Configuration Option |
|------------------|--|----------------------|
| | Signal fail (SF) | ber-sf |
| | Loss of cell delineation (ATM only) | locd |
| | Loss of framing | lof |
| | Loss of light | lol |
| | Loss of pointer | lop-p |
| | Loss of signal | los |
| | Phase-locked loop out of lock | pll |
| | Synchronous transport signal (STS) payload label (C2) mismatch | plm-p |
| | Line remote failure indication | rfl-l |
| | Path remote failure indication | rfl-p |
| | STS path (C2) unequipped | uneq-p |
| E3/T3 | Alarm indicator signal | ais |
| | Excessive numbers of zeros | exz |
| | Failure of the far end | ferf |
| | Idle alarm | idle |

Table 1: Configurable PIC Alarm Conditions *(Continued)*

| Interface/System | Alarm Condition | Configuration Option |
|----------------------------|-------------------------------|----------------------|
| | Line code violation | lcv |
| | Loss of frame | lof |
| | Loss of signal | los |
| | Phase-locked loop out of lock | pll |
| | Yellow alarm | ylw |
| Ethernet | Link has gone down | link-down |
| DS1 | Alarm indicator signal | ais |
| | Yellow alarm | ylw |
| Integrated services | Hardware or software failure | failure |
| Management Ethernet | Link has gone down | link-down |

RELATED DOCUMENTATION

[Chassis Conditions That Trigger Alarms](#)

Alarm Types and Severity Levels

[Network Management and Monitoring Guide](#)

Freeing Up System Storage Space

[show system alarms](#)

System Troubleshooting

IN THIS SECTION

- [Saving Core Files Generated by Junos OS Processes | 26](#)
- [Viewing Core Files from Junos OS Processes | 27](#)

Saving Core Files Generated by Junos OS Processes

By default, when an internal Junos OS process generates a core file, the file and associated context information are saved for debugging purposes in a compressed tar file named `/var/tmp/process-name.core.core-number.tgz`. The contextual information includes the configuration and system log message files.

- To disable the saving of core files and associated context information:

```
[edit system]
no-saved-core-context;
```

- To save the core files only:

```
[edit system]
saved-core-files number;
```

Where ***number*** is the number of core files to save and can be a value from 1 through 10.

- To save the core files along with the contextual information:

```
[edit system]
saved-core-context;
```


Viewing Core Files from Junos OS Processes

When an internal Junos OS process generates a core file, you can find the output at `/var/crash/` and `/var/tmp/`. For Junos OS Evolved, you can find the output core files at `/var/core/` for Routing Engine core files and `/var/lib/ftp/in/` for FPC core files. Using these directories provides a quick method of finding core issues across large networks.

Use the CLI command `show system core-dumps` to view core files.

```
root@host> show system core-dumps
-rw----- 1 root  wheel  268369920 Jun 18 17:59 /var/crash/vmcore.0
-rw-rw---- 1 root  field   3371008 Jun 18 17:53 /var/tmp/rpd.core.0
-rw-r--r-- 1 root  wheel   27775914 Jun 18 17:59 /var/crash/kernel.0
```

SEE ALSO

[Saving Core Files from Junos OS Processes](#)

RELATED DOCUMENTATION

[Day One: Monitoring and Troubleshooting](#)

Device Monitoring

IN THIS SECTION

- [Monitoring System Properties | 28](#)
- [Monitoring System Process Information | 31](#)
- [Monitor Interfaces | 32](#)
- [Other Tools to Configure and Monitor Devices Running Junos OS | 34](#)

Monitoring System Properties

IN THIS SECTION

- Purpose | 28
- Action | 28
- Meaning | 28

Purpose

View system properties such as the name, IP address, and resource usage.

Action

To monitor system properties in the CLI, enter the following commands:

- `show system uptime`
- `show system users`
- `show system storage`

Meaning

[Table 2 on page 28](#) summarizes key output fields in the system properties display.

Table 2: Summary of Key System Properties Output Fields

| Field | Values | Additional Information |
|---------------------|---|--|
| General Information | | |
| Serial Number | Serial number of device. | |
| Junos OS Version | Version of Junos OS active on the switch, including whether the software is for domestic or export use. | Export software is for use outside the USA and Canada. |

Table 2: Summary of Key System Properties Output Fields (*Continued*)

| Field | Values | Additional Information |
|-----------------------|--|------------------------|
| Hostname | Name of the device. | |
| IP Address | IP address of the device. | |
| Loopback Address | Loopback address. | |
| Domain Name Server | Address of the domain name server. | |
| Time Zone | Time zone on the device. | |
| Time | | |
| Current Time | Current system time, in Coordinated Universal Time (UTC). | |
| System Booted Time | Date and time when the device was last booted and how long it has been running. | |
| Protocol Started Time | Date and time when the protocols were last started and how long they have been running. | |
| Last Configured Time | Date and time when a configuration was last committed. This field also shows the name of the user who issued the last <code>commit</code> command. | |
| Load Average | CPU load average for 1, 5, and 15 minutes. | |
| Storage Media | | |

Table 2: Summary of Key System Properties Output Fields *(Continued)*

| Field | Values | Additional Information |
|--------------------------------|--|--|
| Internal Flash Memory | Usage details of internal flash memory. | |
| External Flash Memory | Usage details of external USB flash memory. | |
| Logged in Users Details | | |
| User | Username of any user logged in to the switch. | |
| Terminal | Terminal through which the user is logged in. | |
| From | System from which the user has logged in. A hyphen indicates that the user is logged in through the console. | |
| Login Time | Time when the user logged in. | This is the user@switch field in <code>show system users</code> command output. |
| Idle Time | How long the user has been idle. | |

SEE ALSO

| [show system processes](#)

Monitoring System Process Information

IN THIS SECTION

- Purpose | 31
- Action | 31
- Meaning | 31

Purpose

View the processes running on the device.

Action

To view the software processes running on the device:

```
user@switch> show system processes
```

Meaning

[Table 3 on page 31](#) summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

Table 3: Summary of System Process Information Output Fields

| Field | Values |
|----------|--|
| PID | Identifier of the process. |
| Name | Owner of the process. |
| State | Current state of the process. |
| CPU Load | Percentage of the CPU that is being used by the process. |

Table 3: Summary of System Process Information Output Fields *(Continued)*

| Field | Values |
|--------------------|---|
| Memory Utilization | Amount of memory that is being used by the process. |
| Start Time | Time of day when the process started. |

SEE ALSO

| *show system uptime*

Monitor Interfaces

IN THIS SECTION

- [Purpose | 32](#)
- [Action | 32](#)

Purpose

View general information about all physical and logical interfaces for a device.

Action

Enter the following `show` commands in the CLI to view interface status and traffic statistics.

- `show interfaces terse`



NOTE: On SRX Series Firewalls, when configuring identical IPs on a single interface, you will not see a warning message; instead, you will see a syslog message.

- `show interfaces extensive`

- `show interfaces`*interface-name*



NOTE: If you are using the J-Web user interfaces, select **Monitor>Interfaces** in the J-Web user interface. The J-Web Interfaces page displays the following details about each device interface:

- Port—Indicates the interface name.
- Admin Status—Indicates whether the interface is enabled (Up) or disabled (Down).
- Link Status—Indicates whether the interface is linked (Up) or not linked (Down).
- Address—Indicates the IP address of the interface.
- Zone—Indicates whether the zone is an untrust zone or a trust zone.
- Services—Indicates services that are enabled on the device, such as HTTP and SSH.
- Protocols—Indicates protocols that are enabled on the device, such as BGP and IGMP.
- Input Rate graph—Displays interface bandwidth utilization. Input rates are shown in bytes per second.
- Output Rate graph—Displays interface bandwidth utilization. Output rates are shown in bytes per second.
- Error Counters chart—Displays input and output error counters in the form of a bar chart.
- Packet Counters chart—Displays the number of broadcast, unicast, and multicast packet counters in the form of a pie chart. (Packet counter charts are supported only for interfaces that support MAC statistics.)

To change the interface display, use the following options:

- Port for FPC—Controls the member for which information is displayed.
- Start/Stop button—Starts or stops monitoring the selected interfaces.
- Show Graph—Displays input and output packet counters and error counters in the form of charts.
- Pop-up button—Displays the interface graphs in a separate pop-up window.
- Details—Displays extensive statistics about the selected interface, including its general status, traffic information, IP address, I/O errors, class-of-service data, and statistics.
- Refresh Interval—Indicates the duration of time after which you want the data on the page to be refreshed.

- Clear Statistics—Clears the statistics for the selected interface.

SEE ALSO

[Interfaces User Guide for Security Devices](#)

Other Tools to Configure and Monitor Devices Running Junos OS

Starting in Junos OS Release 15.1, apart from the command-line interface, Junos OS also supports the following applications, scripts, and utilities that enable you to configure and monitor devices running Junos OS:

- Junos XML Management Protocol Application Programming Interface (API)—Application programmers can use the Junos XML Management Protocol API to monitor and configure Juniper Networks devices.
- NETCONF Application Programming Interface (API)—Application programmers can also use the NETCONF API to monitor and configure Juniper Networks devices.
- Junos OS commit scripts—You can define scripts to enforce custom configuration tasks, enforce consistency, prevent common mistakes, and more. Every time you commit a new candidate configuration, the active commit scripts are called to inspect the new candidate configuration. If a configuration violates your custom rules, the script can instruct the Junos OS to perform various actions, including making changes to the configuration and generating custom, warning, and system log messages.
- Junos OS Op scripts—You can add your own commands to the operation-mode CLI. You can use these scripts to automate troubleshooting of known network problems and correct them.
- Junos OS event scripts—You can use event scripts to diagnose and fix issues, monitor the overall status of the system, and examine errors periodically. Event scripts are similar to op scripts except that certain events on the switch will trigger these scripts.
- CHEF—You can use CHEF automate the provisioning and management of compute, networking, and storage resources. Chef for Junos OS provides support for Chef on selected Junos OS devices, allowing you to automate common switching network configurations.
- Puppet—You can use PUPPET for configuration management. Puppet provides an efficient and scalable solution for managing the configurations of large numbers of devices. System administrators take advantage of Puppet to manage compute resources such as physical and virtual servers.

SEE ALSO

- [CLI User Interface Overview](#)
- [NETCONF XML Management Protocol Developer Guide](#)

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

| Release | Description |
|---------|--|
| 15.1 | Starting in Junos OS Release 15.1, apart from the command-line interface, Junos OS also supports the following applications, scripts, and utilities that enable you to configure and monitor devices running Junos OS: |

RELATED DOCUMENTATION

- [Day One: Monitoring and Troubleshooting](#)

Passive Monitoring

IN THIS SECTION

- [Understanding Passive Monitoring | 35](#)
- [Example: Configuring Passive Monitoring | 36](#)

Understanding Passive Monitoring

IN THIS SECTION

- [Passive Monitoring Benefits | 36](#)
- [Guidelines for Configuring Passive Monitoring | 36](#)

Passive monitoring is a type of network monitoring used to passively capture traffic from monitoring interfaces. When you enable passive monitoring, the device accepts and monitors traffic on the interface and forwards the traffic to monitoring tools like IDS servers and packet analyzers, or other devices such as routers or end node hosts.

Passive Monitoring Benefits

- Provides filtering capabilities for monitoring ingress and egress traffic at the Internet point of presence (PoP) where security networks are attached.

Guidelines for Configuring Passive Monitoring

- You can only configure passive monitoring at the interface level. Configuration per VLAN or logical interface is not supported.
- A passive monitoring interface cannot be an aggregated Ethernet (AE) interface.
- Monitoring tools or devices must be directly connected to the switch or router.
- Packets with more than two MPLS labels and more than two VLAN tags are dropped.
- Exception packets such as IP packet options, router alert, and TTL expiry packets are treated as regular traffic.
- Ethernet encapsulation is not supported.
- MPLS family is supported on the PTX10001-36MR, PTX10004, and PTX10008 routers.
- Link Aggregation Control Protocol (LACP) is not supported on the AE bundle connected to the monitoring tool or device.

Example: Configuring Passive Monitoring

IN THIS SECTION

- [Requirements | 37](#)
- [Overview | 37](#)
- [Configuration | 38](#)
- [Verification | 40](#)

This example shows how to configure passive monitoring on QFX10000 switches.

Requirements

This example uses the following hardware and software components:

- Two routers (R1 and R2)
- One QFX10002 switch
- Two devices, directly connected to the switch
- Junos OS Release 18.4R1 or later

Overview

IN THIS SECTION


This example describes how to configure passive monitoring on the switch.

In [Figure 1 on page 38](#), `et-0/0/2` and `et-0/0/4` are configured as passive monitoring interfaces. Packets coming into the network are exchanged between Router 1 (R1) and Router 2 (R2) in two directions (R1 to R2, R2 to R1) and are sent to the monitored interfaces. When traffic is received, a firewall filter transfers all packets to a routing instance and forwards the packets to the monitoring tools. The interfaces are then grouped into a single logical interface, known as a link aggregation group (LAG) or AE bundle. This enables the traffic to be evenly distributed across the monitoring tools effectively increasing the uplink bandwidth. If one interface fails, the bundle continues to carry traffic over the remaining interfaces.

Optionally, you can apply symmetric hashing over the passive monitor interfaces for load balancing traffic to the monitoring tools. This allows ingress and egress traffic of the same flow to be sent out through the same monitored interface. To configure symmetric hashing, include the `no-incoming-port` option under the `[edit forwarding-options enhanced-hash-key]` hierarchy. Symmetric hashing is enabled and disabled at the global level only. Per protocol hashing is not supported.

Topology

Figure 1: Passive Monitoring Topology



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 38](#)
- [Configuring Passive Monitoring | 39](#)

The following example requires you to navigate various levels in the CLI hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set interfaces et-0/0/2 passive-monitor-mode
set interfaces et-0/0/2 unit 0 family inet filter input pm
set interfaces et-0/0/4 passive-monitor-mode
set interfaces et-0/0/4 unit 0 family inet filter input pm1
set firewall family inet filter pm1 term t1 from interface et-0/0/4.0
set firewall family inet filter pm1 term t1 then count c1
set firewall family inet filter pm1 term t1 then routing-instance pm_inst
set firewall family inet filter pm term t1 from interface et-0/0/2.0
set firewall family inet filter pm term t1 then count c3
set firewall family inet filter pm term t1 then routing-instance pm_inst
```



```

set routing-instances pm_inst instance-type virtual-router
set routing-instances pm_inst interface ae0.0
set routing-instances pm_inst routing-options static route 0.0.0.0/0 next-hop 198.51.100.1
set interfaces xe-0/0/9:0 ether-options 802.3ad ae0
set interfaces xe-0/0/9:1 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 198.51.100.2/24 arp 198.51.100.1 mac
00:10:94:00:00:05
set routing-instances pm_inst interface ae0.0
set forwarding-options enhanced-hash-key inet no-incoming-port

```

Configuring Passive Monitoring

Step-by-Step Procedure

To configure passive monitoring:

1. Configure passive-monitor mode on the switch interfaces:

```

[edit]
user@switch#
set interfaces et-0/0/2 passive-monitor-mode
set interfaces et-0/0/2 unit 0 family inet filter input pm
set interfaces et-0/0/4 passive-monitor-mode
set interfaces et-0/0/4 unit 0 family inet filter input pm1

```

2. Configure a family inet firewall filter on the passive monitor interfaces to forward the traffic to a routing instance. Supported filter actions are accept, reject, count, routing-instance.

```

[edit]
user@switch#
set firewall family inet filter pm1 term t1 from interface et-0/0/4.0
set firewall family inet filter pm1 term t1 then count c1
set firewall family inet filter pm1 term t1 then routing-instance pm_inst
set firewall family inet filter pm term t1 from interface et-0/0/2.0
set firewall family inet filter pm term t1 then count c3
set firewall family inet filter pm term t1 then routing-instance pm_inst

```


3. Create a routing-instance with a static route that points to the devices.

```
[edit]
user@switch#
set routing-instances pm_inst instance-type virtual-router
set routing-instances pm_inst interface ae0.0
set routing-instances pm_inst routing-options static route 0.0.0.0/0 next-hop 198.51.100.1
```

4. Configure an AE bundle on the passive monitoring interfaces.

```
[edit]
user@switch#
set interfaces xe-0/0/9:0 ether-options 802.3ad ae0
set interfaces xe-0/0/9:1 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 198.51.100.2/24 arp 198.51.100.1 mac
00:10:94:00:00:05
set routing-instances pm_inst interface ae0.0
```

5. (Optional) Configure symmetric hashing.

```
[edit]
user@switch#
set forwarding-options enhanced-hash-key inet no-incoming-port
```

6. From configuration mode, confirm your configuration by entering the **show interfaces** command. If the command output does not display the intended configuration, repeat the instructions in this example to correct it.
7. If you are done configuring the interfaces, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- [Verify the Passive Monitoring Configuration | 41](#)
- [Verify Symmetric Hashing | 42](#)

Confirm that the configuration is working properly.

Verify the Passive Monitoring Configuration

Purpose

Verify that passive monitoring is working on the interfaces. If the interface output shows No-receive and No-transmit, this means that passive monitoring is working.

Action

From operational mode, enter the **show interfaces** command to view the passive monitoring interfaces.

```

user@host> show interfaces et-0/0/2
Physical interface: et-0/0/2, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 515
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 40Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps No-receive No-transmit Internal: 0x4000
  Link flags     : None
  CoS queues    : 8 supported, 8 maximum usable queues
  Current address: 3c:61:04:75:3c:5d, Hardware address: 3c:61:04:75:3c:5d
  Last flapped  : 2018-05-17 11:19:05 PDT (00:17:55 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects: None
  PCS statistics
    Bit errors           Seconds
    Errored blocks       0
  Ethernet FEC Mode : NONE
  Ethernet FEC statistics
    Errors
    FEC Corrected Errors      0
    FEC Uncorrected Errors    0
    FEC Corrected Errors Rate 0
    FEC Uncorrected Errors Rate 0
  PRBS Statistics : Disabled
  Interface transmit statistics: Disabled

user@host show interfaces et-0/0/4

```



```

Physical interface: et-0/0/4, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 515
  Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 40Gbps, BPDU Error: None, Loop
Detect PDU Error: None, Ethernet-Switching Error: None, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Disabled, Media type: Fiber
  Device flags   : Present Running
  Interface flags: SNMP-Traps No-receive No-transmit Internal: 0x4000
  Link flags     : None
  CoS queues    : 8 supported, 8 maximum usable queues
  Current address: 3c:61:04:75:3c:5d, Hardware address: 3c:61:04:75:3c:5d
  Last flapped  : 2018-05-17 11:19:05 PDT (00:18:17 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  Active alarms : None
  Active defects: None
  PCS statistics
    Bit errors          Seconds
    Errored blocks      0
    Ethernet FEC Mode   : NONE
    Ethernet FEC statistics
      FEC Corrected Errors      0
      FEC Uncorrected Errors    0
      FEC Corrected Errors Rate 0
      FEC Uncorrected Errors Rate 0
  PRBS Statistics : Disabled
  Interface transmit statistics: Disabled

```

Verify Symmetric Hashing

Purpose

Verify the output for symmetric hashing. The incoming port fields for `inet`, `inet6` and `L2` should all be set to No.

Action

From configuration mode, enter the **show forwarding-options enhanced-hash-key** command.

Slot 0


```

Seed value for Hash function      0: 3626023417
Seed value for Hash function      1: 3626023417
Seed value for Hash function      2: 3626023417
Seed value for Hash function      3: 3626023417

```

Inet settings:

```

    IPV4 dest address:    Yes
    IPV4 source address:  Yes
    L4 Dest Port:        Yes
    L4 Source Port:       Yes
Incoming port:           No

```

Inet6 settings:

```

    IPV6 dest address:    Yes
    IPV6 source address:  Yes
    L4 Dest Port:        Yes
    L4 Source Port:       Yes
Incoming port:           No

```

L2 settings:

```

    Dest Mac address:     No
    Source Mac address:   No
    Vlan Id:              Yes
    Inner-vlan Id:        No
    Incoming port:        No

```

GRE settings:

```

    Key:                  No
    Protocol:             No

```

MPLS settings:

```

MPLS Enabled:           Yes

```

VXLAN settings:

```

    VXLAN VNID:           No

```


Sample Configuration for PTX10001-36MR, PTX10004, and PTX10008 Routers

The following is a sample configuration for the PTX10001-36MR, PTX10004, and PTX10008 routers with family mpls support.

```
set interfaces et-0/0/13 passive-monitor-mode
set interfaces et-0/0/13 passive-monitor-mode
set interfaces et-0/0/13 unit 0 family inet filter input ipv4pmFilter
set interfaces et-0/0/13 unit 0 family inet6 filter input ipv6pmFilter
set interfaces et-0/0/13 unit 0 family mpls filter input mplspmFilter
set interfaces et-0/0/5 ether-options 802.3ad ae0
set interfaces et-0/0/7 ether-options 802.3ad ae0
set interfaces ae0 unit 0 family inet address 192.168.1.1/24 arp 192.168.1.10 mac
00:00:00:11:11:11
set interfaces ae0 unit 0 family inet6 address 2001:db8:1::1/64 ndp 2001:db8:1::10 mac
00:00:00:11:11:11
set routing-instances pm_inst routing-options rib pm_inst.inet6.0 static route 0::0/0 next-hop
2001:db8:1::10
set routing-instances pm_inst routing-options static route 0.0.0.0/0 next-hop 192.168.1.10
set routing-instances pm_inst instance-type virtual-router
set routing-instances pm_inst interface ae0.0
set firewall family inet filter ipv4pmFilter term t1 then count C1
set firewall family inet filter ipv4pmFilter term t1 then routing-instance pm_inst
set firewall family inet6 filter ipv6pmFilter term t2 then count C2
set firewall family inet6 filter ipv6pmFilter term t2 then routing-instance pm_inst
set firewall family mpls filter ipv4pmfilter term t1 then count C1
set firewall family mpls filter ipv4pmfilter term t1 then routing-instance pm_inst
set firewall family mpls filter ipv4pmfilter term t1 from ip-version ipv4 ip-protocol-except 255
set firewall family mpls filter ipv6pmfilter term t2 then count C2
set firewall family mpls filter ipv6pmfilter term t2 then routing-instance pm_inst
set firewall family mpls filter ipv6pmfilter term t2 from ip-version ipv6 next-header-except 255
```

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

| Release | Description |
|---------|--|
| 18.4R1 | Starting in Junos OS Release 18.4R1, passive monitoring is supported on QFX10000 switches. |
| 18.4R1 | Starting in Junos OS Evolved 19.4R1, passive monitoring is supported on PTX10003 routers. |

How to Locate a Device or Port Using the Chassis Beacon

IN THIS SECTION

- [Turning On the Chassis Beacon For the Default Interval | 46](#)
- [Turning On the Chassis Beacon For a Specified Interval | 47](#)

By default, when a network port and its associated link are active, the status LED for that port blinks green at a rate of 8 blinks per second. With the chassis beacon feature, you can use the `request chassis beacon` command to slow the current rate at which the status LED blinks green to 2 blinks per second. The slower and steadier green light acts as a beacon that you, as a network administrator in a remote office, can enable to guide a network installer in a busy data center or lab to a Juniper Networks device or port on the device.

You can use the following options with the chassis beacon feature:

- Turn on the beacon for:
 - 5 minutes (default)
 - A specified number of minutes (1 through 120)
- Turn off the beacon:
 - Immediately
 - After a specified number of minutes (1 through 120) elapses

You can use these options on all network ports on an FPC or just one network port on an FPC.

To turn the beacon on or off on a Virtual Chassis, you must:

- Issue the `request chassis beacon` command on the primary switch in the Virtual Chassis.
- When specifying the FPC slot number, use the target Virtual Chassis member number.

You can slow the rate at which the status LED blinks green to 2 blinks per second. The slower and steadier green light acts as a beacon that guides a network installer in a busy data center or lab to a Juniper Networks device or port on the device.

This topic covers the available options in the following use cases:

Turning On the Chassis Beacon For the Default Interval

You can turn on the chassis beacon for the default interval, which is 5 minutes.

1. Turn on the chassis beacon using one of the following commands:

- a. For all network ports on a specified FPC:

```
user@switch> request chassis beacon fpc slot-number on
```

- b. For a specified network port on an FPC:

```
user@switch> request chassis beacon fpc slot-number pic-slot slot-number port port-number on
```

After you turn on the chassis beacon, you can expect the following behavior:

- The chassis beacon overrides the current state of the status LED for all or the specified network port on the FPC.
 - If you turn on the beacon for only one network port, the status LEDs for the remaining network ports on the FPC are turned off.
 - Unless you issue a command to explicitly turn off the chassis beacon before the default interval is over, it turns off after 5 minutes. The state of the status LED for all ports or the specified port returns to the state it was in before you turned on the chassis beacon.
2. If you want to turn the chassis beacon off before the 5-minute interval is over, use one of the following commands:
- a. For all network ports on a specified FPC:

```
user@switch> request chassis beacon fpc slot-number off
```

- b. For a specified network port on an FPC:

```
user@switch> request chassis beacon fpc slot-number pic-slot slot-number port port-number off
```


Turning On the Chassis Beacon For a Specified Interval

You can turn on the chassis beacon for 1 through 120 minutes.

1. Turn on the chassis beacon using one of the following commands:

- a. For all network ports on a specified FPC:

```
user@switch> request chassis beacon fpc slot-number on timer number-of-minutes
```

- b. For a specified network port on an FPC:

```
user@switch> request chassis beacon fpc slot-number pic-slot slot-number port port-number
on timer number-of-minutes
```

After you turn on the chassis beacon, you can expect the following behavior:

- The chassis beacon overrides the current state of the status LEDs for all or one network port on the FPC.
 - If you turn on the chassis beacon for only one network port, the status LEDs for the remaining network ports on the FPC are turned off.
 - The chassis beacon stays on until you explicitly issue a command to turn it off.
2. You can turn off the chassis beacon immediately or after a specified time interval (1 through 120 minutes) is over.
 - a. To turn off the chassis beacon immediately, use one of the following commands:

For all network ports on a specified FPC:

```
user@switch> request chassis beacon fpc slot-number off
```

OR

For a specified network port on an FPC:

```
user@switch> request chassis beacon fpc slot-number pic-slot slot-number port port-number
off
```

- b. To turn off the chassis beacon after a specified time interval of 1 through 120 minutes is over, use one of the following commands:

For all network ports on a specified FPC:

```
user@switch> request chassis beacon fpc slot-number off timer number-of-minutes
```

OR

For a specified network port on an FPC:

```
user@switch> request chassis beacon fpc slot-number pic-slot slot-number port port-number  
off timer number-of-minutes
```

After you turn off the chassis beacon, the state of the status LED for all or one port on the FPC returns to the state it was in before you turned on the chassis beacon.

2

CHAPTER

Configuration Statements and Operational Commands

[Junos CLI Reference Overview](#) | 50

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- [Junos CLI Reference](#)

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- [Configuration Statements](#)
- [Operational Commands](#)