

Junos® OS

Transport and Internet Protocols User Guide

Published
2022-12-15

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.

Junos® OS Transport and Internet Protocols User Guide
Copyright © 2022 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 | viii

1

Understanding IP Support on Junos OS

Junos OS Support for IPv4, IPv6, and MPLS Routing Protocols | 2

2

Configure Transport and Internet Protocol Features

ARP Learning and Aging Options | 6

Configuring Passive ARP Learning for Backup VRRP Devices | 6

Configuring a Delay in Gratuitous ARP Requests | 7

Sending a Gratuitous ARP Request When an Interface is Online | 8

Purging ARP Entries | 8

Adjusting the ARP Aging Timer | 8

Disabling Neighbor Discovery | 9

Example: Configuring ARP Cache Protection | 10

Requirements | 10

Overview | 11

Configuration | 14

Verification | 16

Troubleshooting | 19

ICMP Features | 20

Protocol Redirect Messages | 20

Pings | 22

Disable the Routing Engine Response to Multicast Ping Packets | 23

Disable Reporting IP Address and Timestamps in Ping Responses | 23

Source Quench Messages | 24

Time-to-Live (TTL) Expiration | 24

Rate Limit ICMP Traffic | 25

Rate Limit ICMP Error Messages | 25

ICMP Extension Option for Selective Error Messages | 27

IPv6 Features | 29

Configure IPv6 Duplicate Address Detection Attempts | 30

Accept IPv6 Packets with a Zero Hop Limit | 30

Process IPv4-mapped IPv6 Addresses | 30

Process 6PE Traceroutes | 31

Path MTU Discovery | 32

Configuring Path MTU Discovery on Outgoing TCP Connections | 32

Configuring IP-IP Path MTU Discovery on IP-IP Tunnel Connections | 32

Configuring Path MTU Discovery on Outgoing GRE Tunnel Connections | 33

TCP | 34

Security for TCP Headers with SYN and FIN Flags Set | 34

Disable TCP RFC 1323 Extensions | 35

Configure TCP MSS for Session Negotiation | 36

Configuring TCP MSS on T Series and M Series Routers, and MX Series Routers Using a Service Card | 37

Configuring TCP MSS Inline on MX Series Routers Using MPC Line Cards | 37

Select a Fixed Source Address for Locally Generated TCP/IP Packets | 38

TCP Authentication | 39

IP Subnet Support | 40

VRF Support | 41

TCP Authentication Option (TCP-AO) | 45

TCP-AO for BGP and LDP Sessions | 45

Example: Configure a Keychain (TCP-AO) | 48

Example: Use TCP-AO to Authenticate a BGP Session | 52

Requirements | 52

Overview | 53
Configuration | 53

Example: Use TCP-AO to Authenticate an LDP Session | 59

Requirements | 59
Overview | 60
Configuration | 60
Verification | 64

Example: Use TCP-AO to Authenticate RPKI Validation Sessions | 65

Overview | 66
Requirements | 66
Topology | 66
Configuration | 67

3

Configure Port Security

System Settings | 76

Specifying the Physical Location of the Switch | 76
Modifying the Default Time Zone for a Router or Switch Running Junos OS | 77
Configuring Junos OS to Extend the Default Port Address Range | 78
Configuring Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets | 79
Rebooting and Halting a Device | 80

Password Authentication for Console Access to PICs | 82

4

Configuration Statements

allow-6pe-traceroute | 85
allow-v4mapped-packets | 86
arp | 88
arp-max-cache | 93
arp-new-hold-limit | 95
arp-system-cache-limit | 96
auxiliary | 99

authentication-key-chains (TCP-AO) | 101

console (System Ports) | 103

default-address-selection | 106

diag-port-authentication | 108

extended-statistics | 110

icmp (System) | 112

icmp (Error Message Rate Limit) | 114

icmp6 (Error Message Rate Limit) | 116

internet-options | 118

nd-maxucast-retry | 122

no-multicast-echo | 124

non-subscriber-no-reply | 125

no-ping-record-route | 127

no-ping-time-stamp | 128

path-mtu-discovery (Tunnel) | 130

proactive-arp-detection | 132

tcpao-auth-mismatch | 133

tcp-mss | 135

5

Operational Commands

clear arp | 139

clear multicast snooping statistics | 141

show arp | 143

show system statistics arp | 150

show system statistics icmp | 160

show system statistics icmp6 | 167

`show system statistics igmp` | 177

`show system statistics ip` | 183

`show system statistics ip6` | 195

`show system statistics tcp` | 205

About This Guide

Use this guide to configure the common transport and Internet protocol options.

1

CHAPTER

Understanding IP Support on Junos OS

Junos OS Support for IPv4, IPv6, and MPLS Routing Protocols | 2

Junos OS Support for IPv4, IPv6, and MPLS Routing Protocols

Junos OS implements full IP routing functionality, providing support for IP version 4 and IP version 6 (IPv4 and IPv6, respectively). The routing protocols are fully interoperable with existing IP routing protocols, and they have been developed to provide the scale and control necessary for the Internet core.

Junos OS supports the following unicast routing protocols:

- BGP—Border Gateway Protocol version 4 is an *EGP* that guarantees loop-free exchange of routing information between routing domains (also called autonomous systems). BGP, in conjunction with Junos OS routing policies, provides a system of administrative checks and balances that can be used to implement peering and transit agreements.
- ICMP—Internet Control Message Protocol router discovery enables hosts to discover the addresses of operational routers on the subnet.
- IS-IS—Intermediate System to Intermediate System is a link-state *IGP* for IP networks that uses the *SPF* algorithm, which also is referred to as the *Dijkstra* algorithm, to determine routes. The Junos OS supports a new and complete implementation of the protocol, addressing issues of scale, convergence, and resilience.
- OSPF—Open Shortest Path First is an IGP that was developed for IP networks by the Internet Engineering Task Force (*IETF*). OSPF is a link-state protocol that makes routing decisions based on the *SPF* algorithm.

OSPF Version 2 supports IPv4. OSPF Version 3 supports IPv6. The fundamental mechanisms of OSPF such as flooding, designated router (*DR*) election, area-based topologies, and the *SPF* calculations remain unchanged in OSPF Version 3. Some differences exist either because of changes in protocol semantics between IPv4 and IPv6, or because of the need to handle the increased address size of IPv6.

- RIP—Routing Information Protocol version 2 is a distance-vector IGP for IP networks based on the *Bellman-Ford* algorithm. RIP dynamically routes packets between a subscriber and a service provider without the subscriber having to configure BGP or to participate in the service provider's *IGP* discovery process.

Junos OS also provides the following routing and Multiprotocol Label Switching (*MPLS*) applications protocols:

- *Unicast* routing protocols:
 - BGP

- ICMP
- IS-IS
- OSPF Version 2
- RIP Version 2
- Multicast routing protocols:
 - DVMRP—Distance Vector Multicast Routing Protocol is a *dense-mode (flood-and-prune)* multicast routing protocol.
 - IGMP—Internet Group Management Protocol versions 1 and 2 are used to manage membership in multicast groups.
 - MSDP—Multicast Source Discovery Protocol enables multiple Protocol Independent Multicast (*PIM*) *sparse mode* domains to be joined. A rendezvous point (*RP*) in a PIM sparse mode domain has a peer relationship with an RP in another domain, enabling it to discover multicast sources from other domains.
 - PIM sparse mode and dense mode—Protocol-Independent Multicast is a multicast routing protocol. PIM sparse mode routes to multicast groups that might span wide-area and interdomain internets. PIM dense mode is a flood-and-prune protocol.
 - SAP/SDP—Session Announcement Protocol and Session Description Protocol handle conference session announcements.
- MPLS applications protocols:
 - LDP—The Label Distribution Protocol provides a mechanism for distributing labels in non-traffic-engineered applications. LDP enables routers to establish label-switched paths (LSPs) through a network by mapping network layer routing information directly to data-link layer switched paths. LSPs created by LDP can also traverse LSPs created by the Resource Reservation Protocol (*RSVP*).
 - MPLS—Multiprotocol Label Switching, formerly known as tag switching, enables you to manually or dynamically configure LSPs through a network. It lets you direct traffic through particular paths rather than rely on the IGP least-cost algorithm to choose a path.
 - RSVP—The Resource Reservation Protocol version 1 provides a mechanism for engineering network traffic patterns that is independent of the shortest path decided upon by a routing protocol. RSVP itself is not a routing protocol; it operates with current and future unicast and multicast routing protocols. The primary purpose of RSVP is to support dynamic signaling for MPLS LSPs.

RELATED DOCUMENTATION

[Junos OS Overview](#)

2

CHAPTER

Configure Transport and Internet Protocol Features

ARP Learning and Aging Options | 6

Example: Configuring ARP Cache Protection | 10

ICMP Features | 20

IPv6 Features | 29

Path MTU Discovery | 32

TCP | 34

TCP Authentication Option (TCP-AO) | 45

ARP Learning and Aging Options

IN THIS SECTION

- [Configuring Passive ARP Learning for Backup VRRP Devices | 6](#)
- [Configuring a Delay in Gratuitous ARP Requests | 7](#)
- [Sending a Gratuitous ARP Request When an Interface is Online | 8](#)
- [Purging ARP Entries | 8](#)
- [Adjusting the ARP Aging Timer | 8](#)
- [Disabling Neighbor Discovery | 9](#)

Address Resolution Protocol (ARP) is a protocol used by IPv4 and IPv6 to map IP network addresses to MAC addresses. Use this topic to set passive ARP learning and ARP aging options for network devices. In these situations, a switch operates as a virtual router.

Configuring Passive ARP Learning for Backup VRRP Devices

By default, the backup Virtual Router Redundancy Protocol (VRRP) device drops ARP requests for the VRRP-IP to VRRP-MAC address translation. The backup device does not learn the ARP (IP-to-MAC address) mappings for the hosts sending the requests. When it detects a failure of the primary device and becomes the new primary, the backup device must learn all the entries that were present in the ARP cache of the primary device. In environments with many directly attached hosts, such as metro Ethernet environments for a router, the backup device may have to learn a large number of ARP entries. This can cause a significant transition delay, during which traffic transmitted to some of the hosts might be dropped.

Passive ARP learning enables the ARP cache in the backup device to hold approximately the same contents as the ARP cache in the primary device. When a backup device becomes the primary device, the new primary device will already know the entries in the ARP cache of what used to be the primary device, reducing the transition delay.

To enable passive ARP learning, include the `passive-learning` statement at the `[edit system arp]` hierarchy level:

```
[edit system arp]
passive-learning;
```

While a device is operating as the primary, the passive learning configuration has no operational impact. The primary (or a standalone) device always learns ARP entries from incoming requests. The configuration takes effect only when the device is operating as a backup device.

We recommend setting passive learning on both the backup and primary VRRP device. Otherwise, you will need to remember to configure passive learning on a primary device after it becomes a backup device.

Configuring a Delay in Gratuitous ARP Requests

By default, the Junos OS sends gratuitous ARP requests immediately after you make network-related configuration changes on an interface, like a VLAN ID, MAC address, or IP address change. It also sends gratuitous ARP requests if a failover occurs and the device becomes the new primary device.

The Packet Forwarding Engine may drop some initial request packets if the IP address configuration updates have not been fully processed by the time a gratuitous ARP request is sent. To avoid dropping request packets, you can configure a delay in gratuitous ARP requests.

To configure a delay in gratuitous ARP requests, include the `gratuitous-arp-delay seconds` statement at the `[edit system arp]` hierarchy level:

```
[edit system arp]
gratuitous-arp-delay seconds;
```

We recommend that you configure a value in the range of 3 through 6 seconds.

Sending a Gratuitous ARP Request When an Interface is Online

To configure the device to automatically send a gratuitous ARP request when an interface is online, include the `gratuitous-arp-on-ifup` statement at the `[edit system arp]` hierarchy level:

```
[edit system arp]
gratuitous-arp-on-ifup;
```

Purging ARP Entries

To configure a device to purge obsolete ARP entries in the cache when an interface goes offline, include the `purging` statement at the `[edit system arp]` hierarchy level:

```
[edit system arp]
purging;
```

Purging is configured to delete ARP entries immediately after an interface that has gone offline is detected. If purging is not configured, ARP entries in the ARP table are retried after they have expired and are deleted if there is no ARP response within the default timeout value of 20 minutes. The default timeout value can be changed to other values using the `aging-timer` statement, as explained below.

Adjusting the ARP Aging Timer

By default, the ARP aging timer is set at 20 minutes. In environments with many directly attached hosts, such as metro Ethernet environments, increasing the amount of time between ARP updates by configuring the ARP aging timer can improve performance. Thousands of clients timing out at the same time might impact packet forwarding performance. In environments where there are devices connected with lower ARP aging timers (less than 20 minutes), decreasing the ARP aging timer can improve performance by preventing the flooding of traffic toward next hops with expired ARP entries. In most environments, the default ARP aging timer value does not need to be adjusted.

The range of the ARP aging timer is 1 through 240 minutes. To configure a system-wide ARP aging timer, include the `aging-timer` statement at the `[edit system arp]` hierarchy level:

```
[edit system arp]
aging-timer minutes;
```

You can also configure the ARP aging timer for each logical interface of family type `inet`. To configure the ARP aging timer at the logical interface level, specify the `aging-timer` statement and the timer value in minutes at the `[edit system arp interfaces interface-name]` hierarchy level:

```
[edit system arp interfaces interface-name]
aging-timer minutes;
```

To configure the ARP aging timer for a specific interface in a logical system, include the `aging-timer` statement and the timer value in minutes at the `[edit logical-systems logical-system-name system arp interfaces interface-name]` hierarchy level:

```
[edit logical-systems logical-system-name system arp interfaces interface-name]
aging-timer minutes;
```

NOTE: If the aging timer value is configured both at the system and the logical interface levels, the value configured at the logical interface level takes precedence for the specific logical interface.

The timer value you configure takes effect as ARP entries expire. Each refreshed ARP entry receives the new timer value. The new timer value does not apply to ARP entries that exist at the time you commit the configuration.

Disabling Neighbor Discovery

You can prevent the device from learning the MAC addresses of its neighbors through ARP or neighbor discovery for IPv4 and IPv6 neighbors. To disable ARP address learning by not sending ARP requests and not learning from ARP replies, use the `no-neighbor-learn` configuration statement.

To disable neighbor discovery for IPv4 neighbors:

```
[edit interfaces interface-name unit interface-unit-number family inet]  
no-neighbor-learn;
```

To disable neighbor discovery for IPv6 neighbors:

```
[edit interfaces interface-name unit interface-unit-number family inet6]  
no-neighbor-learn;
```

Example: Configuring ARP Cache Protection

IN THIS SECTION

- [Requirements | 10](#)
- [Overview | 11](#)
- [Configuration | 14](#)
- [Verification | 16](#)
- [Troubleshooting | 19](#)

You can configure an ARP cache limit for resolved and unresolved next-hop entries in the cache. This example shows how to configure ARP cache protection by specifying a maximum count and hold limit for resolved and unresolved next-hop entries in the ARP cache. This limit can be specified globally for all interfaces, or locally on a particular interface of the device. The benefit of configuring such a limit on the ARP cache is to protect the device from denial-of-service (DoS) attacks.

Requirements

This example uses the following hardware and software components:

- Two routers that can be a combination of M, MX, and T Series routers.

- Two host devices connected to the routers.
- Junos OS Release 16.1 or later running on the routers.

Overview

IN THIS SECTION

- [Topology | 13](#)

Sending IP packets on a multiaccess network requires mapping from an IP address to a media access control (MAC) address (the physical or hardware address). In an Ethernet environment, ARP is used to map a MAC address to an IP address. Hosts that use ARP maintain a cache of discovered Internet-to-Ethernet address mappings to minimize the number of ARP broadcast messages.

To keep the cache from growing too large, by default, an entry is removed from the cache if it is not used within a certain period of time. In addition to this, starting in Junos OS Release 16.1, you can manage the number of ARP cache entries by configuring a limit on the resolved and unresolved next-hop entries.

The ARP cache feature supports two types of limits:

- Count—Count limit is the maximum number of next hops that can be created in the ARP cache.
- Hold—Hold limit is the maximum number of hold routes pointing to a particular interface that can be retained before getting added to the ARP cache.

The ARP cache limits are executed at two levels:

- Local—Local limits are configured per interface and are defined for resolved and unresolved entries in the ARP cache.
- Global—Global limits apply system-wide. A global limit is further defined separately for the public interfaces and management interfaces, for example, fxp0. The management interface has a single global limit and no local limit. The global limit enforces a system-wide cap on entries for the ARP cache, including private Internal routing interfaces (IRIs) for internal routing instances, for example, em0 and em1.

Small-sized platforms: ACX, EX22XX, EX3200, EX33XX, and SRX; default is 20,000. Medium-sized platforms: EX4200, EX45XX, EX4300, EX62XX, and MX; default is 75,000. All other platforms, default is 100,000. You can modify this limit by configuring the ARP next-hop cache protection feature.

- To configure the ARP cache count limit for resolved and unresolved next-hop entries globally, include the `arp-system-cache-limit` statement at the `[edit system]` hierarchy level.
- To configure the ARP cache count limit for resolved and unresolved next-hop entries locally, include the `arp-system-cache-limit` statement at the `[edit interfaces interface-name unit interface-unit-number family inet]` hierarchy level.
- To configure the ARP cache hold limit for unresolved next-hop entries locally, include the `arp-new-hold-limit` statement at the `[edit interfaces interface-name unit interface-unit-number family inet]` hierarchy level.

NOTE: The ARP cache hold limit is configured on a per-interface basis only, and cannot be configured at the system level.

The ARP cache next-hop entries get allotted to different types of interfaces differently, irrespective of the ARP cache protection feature configuration.

1. By default, 200 entries get allotted to IRIs.
2. 80 percent of the remaining entries get allotted to public interfaces.
3. 20 percent of the remaining entries get allotted to management interfaces.


When the ARP next-hop entries exceed the configured count limit, new entries are either discarded, or kept under the hold counter, if a hold limit is configured for that interface. The ARP next-hop hold limit specifies the maximum number of hold entries or hold routes that point to a particular interface. When the number of hold entries exceeds the configured hold limit, the drop counter for that interface is affected drastically, as the new hold entries create a loop and continue to increment until there is bandwidth to accommodate them.

NOTE: After modifying the default ARP next-hop cache limit on an interface, the interface must be deactivated and reactivated for the newly configured values to take effect.

Topology

Figure 1 on page 13 illustrates a simple two-router topology with ARP cache protection enabled. Routers R1 and R2 are each connected to hosts, Host1 and Host2, respectively.

Figure 1: ARP Cache Protection



For example, if Router R1 is configured with an `arp-system-cache-limit` of 220 globally, and it receives 230 ARP entries, on the first interface receiving the entries (say, `ge-0/0/0`), the following actions are performed:

1. When 230 entries are received, the global limit of 220 entries is applied to the system, where the configured limit is divided among the different types of interfaces, and the remaining entries received on a particular interface get discarded.
2. Out of the 220 cached entries, by default, 200 entries are allocated for IRI interfaces.
3. Out of the remaining 20 entries, 80 percent of the entries (16 entries) are sent to public interfaces and 20 percent of the entries (4 entries) are sent to the management interface. If the 230 ARP entries are received on the public interface, only the cache limit of 16 entries is retained, and the remaining 214 entries get discarded.

In addition, if `ge-0/0/0` on Router R1 is configured with an `arp-new-hold-limit` value of 8, the following actions are performed:

1. Out of the 230 received entries, only 220 entries are cached in the ARP table. However, instead of discarding the remaining entries, the hold entries are sent to the hold counter of `ge-0/0/0`, and then the remaining entries are sent to the drop counter of `ge-0/0/0`.
2. Depending on availability of bandwidth, the eight hold entries are cached in the ARP table of `ge-0/0/0` before taking any newly received entries into account.
3. The drop counter of `ge-0/0/0`, however, does not increment by single entries. The discarded hold entries in the drop counter form a loop and add to the entries count until there is bandwidth on the

interface to accommodate all the entries. Therefore, additions to the drop counter have a drastic effect on the interface performance.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 14](#)
- [Procedure | 15](#)
- [Results | 15](#)

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.

R1

```
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/0/0 unit 0 family inet arp-new-hold-limit 8
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/24
set interfaces lo0 unit 0 family inet address 10.10.10.1/32
set system arp-system-cache-limit 220
```

R2

```
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.2/24
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/24
set interfaces lo0 unit 0 family inet address 10.20.20.1/32
```

Procedure

Step-by-Step Procedure

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

To configure Router R1 with ARP cache protection:

1. Configure the interfaces of Router R1.

```
[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet address 192.0.2.1/24
user@R1# set ge-0/0/1 unit 0 family inet address 192.0.2.1/24
user@R1# set lo0 unit 0 family inet address 10.10.10.1/32
```

2. Configure ARP cache protection globally for all the interfaces of Router R1.

```
[edit system]
user@R1# set arp-system-cache-limit 220
```

3. Configure a hold limit on the ARP cache entries of interface ge-0/0/0 of Router R1.

```
[edit interfaces]
user@R1# set ge-0/0/0 unit 0 family inet arp-new-hold-limit 8
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces` and `show system` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R1# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
  }
}
```

```
}
ge-0/0/1 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 10.10.10.1/32;
    }
  }
}
```

```
user@R1# show system
arp-system-cache-limit 220 ;
```

Verification

IN THIS SECTION

- [Verifying Global ARP Next-Hop Cache Limit | 16](#)
- [Verifying Local ARP Next-Hop Cache Limit | 17](#)

Confirm that the configuration is working properly.

Verifying Global ARP Next-Hop Cache Limit

Purpose

Verify the system-wide ARP next-hop cache limits and the allocation of next-hop entries for different interfaces.

Action

From operational mode, run the **show system statistics arp** command.

```
user@R1> show system statistics arp
arp:
    717253 datagrams received
    47 ARP requests received
    31 ARP replies received
    285 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 unrestricted proxy requests not proxied
*****
    220 Max System ARP nh cache limit
    16 Max Public ARP nh cache limit
    200 Max IRI ARP nh cache limit
    4 Max Management intf ARP nh cache limit
    16 Current Public ARP next-hops present
    1 Current IRI ARP next-hops present
    2 Current Management ARP next-hops present
    2457 Total ARP next-hops creation failed as limit reached
    2454 Public ARP next-hops creation failed as public limit reached
    3 IRI ARP next-hops creation failed as iri limit reached
    0 Management ARP next-hops creation failed as mgt limit reached
```

Meaning

The global ARP next-hop cache limits are displayed in the output, along with the allocation of next-hop entries for IRI, public, and management interfaces.

Verifying Local ARP Next-Hop Cache Limit

Purpose

Verify the interface ARP next-hop cache limit.

Action

From operational mode, run the **show interfaces** *interface-name* command.

```
user@R1> show interface fxp0
fxp0
Physical interface: fxp0, Enabled, Physical link is Up
  Interface index: 1, SNMP ifIndex: 1
  Type: Ethernet, Link-level type: Ethernet, MTU: 1514, Speed: 100mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Link type      : Full-Duplex
  Current address: 00:a0:a5:62:8e:39, Hardware address: 00:a0:a5:62:8e:39
  Last flapped   : 2014-10-16 10:23:29 PDT (16:27:21 ago)
    Input packets : 0
    Output packets: 0

Logical interface fxp0.0 (Index 3) (SNMP ifIndex 13)
  Flags: Up SNMP-Traps Encapsulation: ENET2
  Bandwidth: 0
  Input packets : 23
  Output packets: 4
  Protocol inet, MTU: 1500
  Max nh cache: 220 New hold nh limit: 8, Curr nh cnt: 2, Curr new hold cnt: 0, NH drop cnt: 0
    Flags: Sendbroadcast-pkt-to-re, Is-Primary
    Addresses, Flags: Is-Default Is-Preferred Is-Primary
      Destination: 10.209.0/18, Local: 10.209.3.69, Broadcast: 10.209.63.255
```

Meaning

The local ARP next-hop cache count and hold limits for the management interface is displayed in the output.

Troubleshooting

IN THIS SECTION

- [Troubleshooting System Log Messages | 19](#)

To troubleshoot the ARP cache protection configuration, see:

Troubleshooting System Log Messages

Problem

System log messages are generated to record events when the ARP cache limits are exceeded.

Solution

To interpret the system log messages, refer to the following:

- **Feb 08 17:12:39 [TRACE] [R1]: Public intf soft (80%) arp nh cache limit reached**—Router R1 has reached 80 percent of the allowed ARP next-hop cache limit for public interfaces.
- **Feb 08 17:07:43 [TRACE] [R1]: Public intf hard arp nh cache limit reached**—Router R1 has reached the maximum allowed limit for ARP next-hop cache entries on the public interface.
- **Feb 08 17:15:14 [TRACE] [R1]: Max cache soft (80%) arp nh cache limit for intf idx 325 reached**—Router R1 has reached 80 percent of the configured global ARP next-hop cache limit for all its interfaces.
- **Feb 08 17:19:41 [TRACE] [R1]: Max cache hard arp nh cache limit for intf idx 325 reached**—Router R1 has reached the maximum configured global ARP next-hop cache limit for all its interfaces.

Release History Table

Release	Description
16.1	Starting in Junos OS Release 16.1, you can configure an ARP cache limit for resolved and unresolved next-hop entries in the cache.

RELATED DOCUMENTATION

[arp-system-cache-limit](#) | 96

[arp-new-hold-limit](#) | 95

ICMP Features

IN THIS SECTION

- [Protocol Redirect Messages](#) | 20
- [Pings](#) | 22
- [Source Quench Messages](#) | 24
- [Time-to-Live \(TTL\) Expiration](#) | 24
- [Rate Limit ICMP Traffic](#) | 25
- [Rate Limit ICMP Error Messages](#) | 25
- [ICMP Extension Option for Selective Error Messages](#) | 27

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

Protocol Redirect Messages

IN THIS SECTION

- [Understanding Protocol Redirect Messages](#) | 21
- [Disable Protocol Redirect Messages](#) | 21

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 | 23](#)
- [Disable Reporting IP Address and Timestamps in Ping Responses | 23](#)

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 | 26](#)
- [How to Rate Limit ICMPv4 and ICMPv6 Error Messages | 26](#)

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 `don't-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
```

ICMP Extension Option for Selective Error Messages

IN THIS SECTION

- [Benefits of ICMP Extension | 28](#)
- [How to Enable ICMP Extension | 28](#)

An IP device uses the ICMP protocol to diagnose network communications problems, particularly to determine whether a datagram is arriving at its intended destination in a timely manner. If a datagram does not arrive at the intended destination, ICMP reports an appropriate error message to the originating IP device.

When network problems prevent IP packet delivery, network devices use ICMP to generate error messages to the source IP address. ICMPv4 and ICMPv6 provide an extension option for selective error messages.

Benefits of ICMP Extension

ICMP extension helps to identify the interface and other information as follows:

- ICMPv4 and ICMPv6 messages couldn't identify the interface of a datagram that cannot be processed on an unnumbered interface.
- ICMP messages are created by determining the source address of an incoming interface and sending packets to the origination device; however, the origin device has no way of knowing where the ICMP message originated.

The ICMP extension enables you to identify the network device responding to the ICMP message that includes the following information:

- A datagram received through an IP interface.
- A datagram arrived at the sub-IP component of an IP interface.
- The IP interface in which the datagram would be forwarded.
- Next-hop IP address to which it would have been forwarded.

We've implemented RFC5837 to enable us to append additional fields to select ICMP (IPv4 and IPv6) messages for both numbered and unnumbered aggregated Ethernet interfaces:

- ICMPv4 Time Exceeded
- ICMPv4 Destination Unreachable
- ICMPv6 Time Exceeded
- ICMPv6 Destination Unreachable

NOTE: The ICMPv6 extension is only supported for numbered interfaces.

How to Enable ICMP Extension

To enable the ICMPv4 extension:

```
[edit chassis]
user@host# set system allow-icmp4-extension
```

To disable the ICMPv4 extension, delete the configuration:

```
[edit chassis]
user@host# delete system allow-icmp4-extension
```

To enable the ICMPv6 extension:

```
[edit chassis]
user@host# set system allow-icmp6-extension
```

To disable the ICMPv6 extension, delete the configuration:

```
[edit chassis]
user@host# delete system allow-icmp6-extension
```

RELATED DOCUMENTATION

| [TCP](#) | [34](#)

IPv6 Features

IN THIS SECTION

- [Configure IPv6 Duplicate Address Detection Attempts](#) | [30](#)
- [Accept IPv6 Packets with a Zero Hop Limit](#) | [30](#)
- [Process IPv4-mapped IPv6 Addresses](#) | [30](#)
- [Process 6PE Traceroutes](#) | [31](#)

Configure IPv6 Duplicate Address Detection Attempts

To set the number of attempts the device makes to detect IPv6 duplicate addresses, use the `ipv6-duplicate-addr-detection-transmits` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
ipv6-duplicate-addr-detection-transmits;
```

Accept IPv6 Packets with a Zero Hop Limit

By default, incoming IPv6 packets that have a zero hop limit value in their header are rejected both when they are addressed to the local host and when they are transiting the device. To accept zero hop-limit packets addressed to the local host, include the `no-ipv6-reject-zero-hop-limit` statement at the `[edit system internet-options]` hierarchy level. Transit packets are still dropped.

```
[edit system internet-options]
no-ipv6-reject-zero-hop-limit;
```

To re-enable rejection of these packets, use the following statement:

```
[edit system internet-options]
ipv6-reject-zero-hop-limit;
```

Process IPv4-mapped IPv6 Addresses

By default, the Junos OS disables the processing of IPv4-mapped IPv6 packets to protect against malicious packets from entering the network. You may want to enable IPv4-mapped IPv6 packets:

- To ensure smooth packet flow in a mixed routing environment of IPv4 and IPv6 networks.
- So that IPv6 packets aren't dropped in a pure IPv4 routing environment.
- When you are transitioning your routing environment from IPv4 to IPv6 networks.

To enable the processing of IPv4-mapped IPv6 packets, use the [allow-v4mapped-packets](#) statement:

```
[edit system]
allow-v4mapped-packets;
```

NOTE: We recommend that you configure this statement only after fully understanding the security implications of allowing IPv4-mapped IPv6 packets in your network.

Process 6PE Traceroutes

In a dual-stack IPv6 network connected over an IPv4 MPLS network, the P routers in the IPv4 MPLS backbone do not have an IPv6 family. Consequently, the transit P routers are not shown in the output when you do an IPv6 traceroute. To generate an ICMPv6 echo request and a TTL expired response packet to and from the intermediate transit routers in the 6PE network, use the [allow-6pe-traceroute](#) statement:

```
[edit system]
allow-6pe-traceroute;
```

RELATED DOCUMENTATION

[Understanding IPv6](#)

[IPv6 Neighbor Discovery Overview](#)

[Path MTU Discovery](#) | 32

Junos OS Support for IPv4, IPv6, and MPLS Routing Protocols

Path MTU Discovery

IN THIS SECTION

- [Configuring Path MTU Discovery on Outgoing TCP Connections | 32](#)
- [Configuring IP-IP Path MTU Discovery on IP-IP Tunnel Connections | 32](#)
- [Configuring Path MTU Discovery on Outgoing GRE Tunnel Connections | 33](#)

Configuring Path MTU Discovery on Outgoing TCP Connections

By default, path MTU discovery on outgoing TCP connections is enabled. To disable path MTU discovery, include the `no-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]  
no-path-mtu-discovery;
```

To reenable path MTU discovery on outgoing TCP connections, include the `path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]  
path-mtu-discovery;
```

Configuring IP-IP Path MTU Discovery on IP-IP Tunnel Connections

By default, path maximum transmission unit (MTU) discovery on outgoing IP-IP tunnel connections is enabled.

To disable IP-IP path MTU discovery, include the `no-ipip-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
no-ipip-path-mtu-discovery;
```

To reenale IP-IP path MTU discovery, include the `ipip-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
ipip-path-mtu-discovery;
```

Configuring Path MTU Discovery on Outgoing GRE Tunnel Connections

By default, path MTU discovery on outgoing GRE tunnel connections is enabled. To disable GRE path MTU discovery, include the `no-gre-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
no-gre-path-mtu-discovery;
```

To re-enable GRE path MTU discovery, include the `gre-path-mtu-discovery` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
gre-path-mtu-discovery;
```

NOTE: To verify details of the path MTU on outgoing GRE tunnels, use the command `show interfaces (GRE)`.

RELATED DOCUMENTATION

TCP

IN THIS SECTION

- [Security for TCP Headers with SYN and FIN Flags Set | 34](#)
- [Disable TCP RFC 1323 Extensions | 35](#)
- [Configure TCP MSS for Session Negotiation | 36](#)
- [Select a Fixed Source Address for Locally Generated TCP/IP Packets | 38](#)
- [TCP Authentication | 39](#)

Many applications and services use TCP to communicate. Configure TCP options to improve link quality and security.


Security for TCP Headers with SYN and FIN Flags Set

By default, your device accepts packets that have both the SYN and FIN bits set in the TCP flag. Configure your device to drop packets with both the SYN and FIN bits set to reduce security vulnerabilities.

Both the SYN and FIN control flags are not normally set in the same TCP segment header. The SYN flag synchronizes sequence numbers to initiate a TCP connection. The FIN flag indicates the end of data transmission to finish a TCP connection. Their purposes are mutually exclusive. A TCP header with the

SYN and FIN flags set is anomalous TCP behavior, causing various responses from the recipient, depending on the OS. See [Figure 2 on page 35](#).

Figure 2: TCP Header with SYN and FIN Flags Set



An attacker can send a segment with both flags set to see what kind of system reply is returned and thereby determine what kind of OS is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks. When you enable the `tcp-drop-synfin-set` statement, Junos OS checks if the SYN and FIN flags are set in TCP headers. If it discovers such a header, it drops the packet.

```
[edit system internet-options]
tcp-drop-synfin-set;
```

Disable TCP RFC 1323 Extensions

To disable RFC 1323 TCP extensions, include the `no-tcp-rfc1323` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
no-tcp-rfc1323;
```

To disable the Protection Against Wrapped Sequence (PAWS) number extension (described in RFC 1323, *TCP Extensions for High Performance*), include the `no-tcp-rfc1323-paws` statement at the `[edit system internet-options]` hierarchy level:

```
[edit system internet-options]
no-tcp-rfc1323-paws;
```

Configure TCP MSS for Session Negotiation

IN THIS SECTION

- [Configuring TCP MSS on T Series and M Series Routers, and MX Series Routers Using a Service Card | 37](#)
- [Configuring TCP MSS Inline on MX Series Routers Using MPC Line Cards | 37](#)

During session connection establishment, two peers agree in negotiations to determine the IP segment size of packets that they will exchange during their communication. The TCP MSS (maximum segment size) value in TCP SYN packets specifies the maximum number of bytes that a TCP packet's data field, or segment, can contain. An MSS value that is set too high can result in an IP datagram that is too large to send and that must be fragmented. Fragmentation can incur additional overhead cost and packet loss.

To diminish the likelihood of fragmentation and to protect against packet loss, you can use the `tcp-mss` statement to specify a lower TCP MSS value. The `tcp-mss` statement applies to all IPv4 TCP SYN packets traversing all the router's ingress interfaces whose MSS value is higher than the one you specify. You cannot exempt particular ports from its effects.

The following section describes how to configure TCP MSS on T Series, M Series, and MX Series routers.

Configuring TCP MSS on T Series and M Series Routers, and MX Series Routers Using a Service Card

To specify a TCP MSS value on T Series and M Series routers as well as MX Series routers using a service card, include the `tcp-mss mss-value` statement at the `[edit services service-set service-set-name]` hierarchy level:

```
[edit services service-set service-set-name]
tcp-mss mss-value;
```

The range of the `tcp-mss mss-value` parameter is from 536 through 65535 bytes.

Add the service set to any interface for which you want to adjust the TCP-MSS value:

```
[edit interfaces interface-name unit 0 family family service]
input service-set service-set-name;
output service-set service-set-name;
```

To view statistics of SYN packets received and SYN packets whose MSS value is modified, issue the `show services service-sets statistics tcp-mss operational mode` command.

For further information about configuring TCP MSS on T Series and M Series routers, see the [Junos OS Services Interfaces Library for Routing Devices](#).

Configuring TCP MSS Inline on MX Series Routers Using MPC Line Cards

To specify a TCP MSS value on MX Series routers that use MPC line cards, include the `tcp-mss` statement at the `[edit interfaces interface-name unit logical-unit-number family family]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family family]
tcp-mss mss-value;
```

The range of the `mss-value` parameter is from 64 through 65,535 bytes. The TCP MSS value must be lower than the MTU of the interface.

This statement is supported on the following interfaces: gr- (GRE), ge- (Gigabit Ethernet), xe- (10-Gigabit Ethernet), and et- (40-Gigabit and 100-Gigabit Ethernet). Families supported are `inet` and `inet6`.

NOTE: Configuring TCP MSS inline on MX Series routers using MPC line cards works only for traffic exiting/egressing the interface, not traffic entering/ingressing the interface.

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

Locally generated IP packets are the packets that are produced by applications running on the Routing Engine. Junos OS chooses a source address for these packets so that the application peers can respond. It also enables you to specify the source address on a per application basis. To serve this purpose, the Telnet CLI command contains the `source-address` argument.

This section introduces the `default-address-selection` statement:

```
[edit system]
default-address-selection;
```

If you specifically choose the source address, as in the case of Telnet, `default-address-selection` does not influence the source address selection. The source address becomes the one that is specified with the `source-address` argument (provided the address is a valid address specified on the interface of a router). If the source address is not specified or if the specified address is invalid, `default-address-selection` influences the default source address selection.

If the source address is not explicitly specified as in the case of Telnet, then by default (when `default-address-selection` is not specified) the source address chosen for locally generated IP packets is the IP address of the outgoing interface. This indicates that depending on the chosen outgoing interface, the source address might be different for different invocations of a given application.

If the interface is unnumbered (no IP address is specified on an interface), Junos OS uses a predictable algorithm to determine the default source address. If `default-address-selection` is specified, Junos OS uses the algorithm to choose the source address irrespective of whether the outgoing interface is numbered. This indicates that with `default-address-selection`, you can influence Junos OS to provide the same source address in locally generated IP packets regardless of the outgoing interface.

The behavior of source address selection by Junos OS can be summed up as shown in the following table:

Table 1: Source Address Selection

Outgoing Interface	When default-address-selection Is Specified	When default-address-selection Is Not Specified
Unnumbered	Use default-address-selection	Use default-address-selection

Table 1: Source Address Selection *(Continued)*

Outgoing Interface	When default-address-selection Is Specified	When default-address-selection Is Not Specified
Numbered	Use default-address-selection	Use IP address of outgoing interface

See [Configuring Default, Primary, and Preferred Addresses and Interfaces](#) for more information about the default address source selection algorithm.

NOTE: For IP packets sent by IP routing protocols (including OSPF, RIP, RSVP, and the multicast protocols, but not including IS-IS), the local address selection is often constrained by the protocol specification so that the protocol operates correctly. When this constraint exists in the routing protocol, the packet's source address is unaffected by the presence of the default-address-selection statement in the configuration. For protocols in which the local address is unconstrained by the protocol specification like IBGP and multihop EBGP, if you do not configure a specific local address when configuring the protocol, the local address is chosen using the same method as other locally generated IP packets.

TCP Authentication

IN THIS SECTION

- [IP Subnet Support | 40](#)
- [VRF Support | 41](#)

Enabling a TCP authentication method enhances the security and ensures the authenticity of TCP segments exchanged during BGP and LDP sessions. Junos devices support three main types of TCP authentication: TCP MD5, TCP Authentication Option (TCP-AO), and TCP keychain-based authentication. For more information about TCP-AO, see ["TCP Authentication Option \(TCP-AO\)" on page 45](#).

NOTE: Although Junos devices support both the TCP-AO and TCP MD5 authentication methods, you cannot use both at the same time for a given connection.

IP Subnet Support

IN THIS SECTION

- BGP | 40
- LDP | 41

Prior to Junos OS Evolved Release 22.4R1, Junos devices only permit you to use TCP authentication with a specific address. This means you can only authenticate TCP connections to remote peers with known IP addresses.

Starting in Junos OS Evolved Release 22.4R1, TCP-AO and TCP MD5 authentication support IP subnets for LDP and BGP sessions. When you configure TCP authentication with a network address and a prefix length, your chosen TCP authentication method authenticates TCP connections to the entire range of addresses under that subnet. This means you can authenticate TCP connections without needing to know the exact IP addresses of the destination devices.

When IP subnets overlap, the authentication method uses the longest prefix match (LPM) to determine the exact authentication key for a specific TCP session.

BGP

To configure prefix-based authentication for BGP sessions, include the `allow (all | prefix-list)` statement at either of the following hierarchies:

- `[edit protocols bgp group group-name]`
- `[edit protocols bgp group group-name dynamic-neighbor dyn-name]`

You can use IPV4 or IPV6 addresses for the subnet.

In this example, TCP MD5 authenticates TCP connections to devices in the 10.0.3.0/24 subnet for all BGP sessions:

```
[edit protocols]
bgp {
  group one {
    authentication-key "$ABC123";
    allow 10.0.3.0/24;
    dynamic-neighbor dyn_one {
      allow 10.0.3.0/24;
      authentication-key "$ABC123";
    }
  }
}
```

LDP

To configure prefix-based authentication for LDP, configure TCP authentication under the *session-group ip-prefix* hierarchy. You must use an IPv4 address.

In this example, LDP uses TCP-AO to authenticate any TCP connection with a device that has an address in the 10.0.0.0/24 subnet:

```
[edit protocols ldp]
session-group 10.0.0.0/24 {
  authentication-algorithm ao;
  authentication-key-chain tcpao;
}
```

For how to configure your TCP-AO keychain, see ["TCP Authentication Option \(TCP-AO\)" on page 45](#).

VRF Support

IN THIS SECTION

- [BGP | 42](#)
- [LDP | 44](#)

In releases prior to Junos OS Evolved Release 22.4R1, TCP MD5 and TCP-AO ignore virtual routing and forwarding (VRF) instances. The device ignores TCP MD5 and TCP-AO configurations under non-default routing instances. When you configure TCP MD5 or TCP-AO under the default VRF instance, the device applies that authentication method to all TCP sessions that have destinations inside the IP address range for that VRF instance. If a TCP session belonged to non-default VRF instance but had the same destination IP address as the default VRF instance, TCP MD5 and TCP-AO would apply the same authentication key to two TCP connections with the same destination IP address.

Starting in Junos OS Evolved Release 22.4R1, TCP-AO and TCP MD5 authentication are VRF aware in BGP and LDP sessions. You can configure TCP-AO and TCP MD5 under non-default routing instances. The TCP authentication method you configure under a routing instance is only applied to the TCP sessions inside that VRF instance. If a TCP connection in a different VRF instance has the same destination IP address, the TCP authentication method does not get applied to that TCP connection if the VRF instance does not have TCP authentication configured for the peer.

Configure VRF-based TCP authentication as you normally would, but under a `routing-instances` hierarchy level. To use TCP MD5 authentication, include the `authentication-key authentication-key` statement. To use TCP-AO, include the following statements:

```
user@device# set authentication-algorithm ao
user@device# set authentication-key-chain keychain
```

For how to configure your TCP-AO keychain, see ["TCP Authentication Option \(TCP-AO\)" on page 45](#).

You can combine VRF-aware configurations with IP subnets. This enables you to authenticate connections to a range of addresses inside the VRF instance.

BGP

Configure VRF-based TCP authentication for BGP sessions at any of the following hierarchy levels:

- [edit routing-instances *vrf-instance* protocols bgp]
- [edit routing-instances *vrf-instance* protocols bgp group *group-name*]
- [edit routing-instances *vrf-instance* protocols bgp group *group-name* neighbor *neighbor-ip*]
- [edit routing-instances *vrf-instance* protocols bgp group *group-name* dynamic-neighbor *dyn-name*]

If you configure VRF-based authentication at the `dynamic-neighbor` level, include the `allow` statement along with your chosen authentication method configuration. For example, to use TCP-AO with a dynamic neighbor:

```
[edit routing-instances vrf-instance protocols bgp group group-name dynamic-neighbor dyn-name]
user@device# set allow (all | prefix-list)
user@device# set authentication-algorithm ao
user@device# set authentication-key-chain keychain
```

In the following example, BGP uses TCP authentication to ensure the security of TCP connections in a VRF instance called `vrf-one`. In group one, BGP uses TCP MD5 to authenticate connections to the neighbor with the IP address 10.0.1.1. It uses TCP-AO to authenticate connections to the neighbor with the IP address 10.0.1.2.

In group two, BGP uses TCP-AO to authenticate connections to any device in the 10.0.0.0/24 subnet.

```
[edit routing-instances]
vrf-one {
  protocols {
    bgp {
      group one {
        peer-as 22;
        neighbor 10.0.1.1 {
          authentication-key "$ABC123"; ## SECRET-DATA
        }
        neighbor 10.0.1.2 {
          authentication-algorithm ao;
          authentication-key-chain tcpao;
        }
      }
      group two {
        peer-as 22;
        dynamic-neighbor dyn_two {
          allow 10.0.0.0/24;
          authentication-algorithm ao;
          authentication-key-chain tcpao;
        }
      }
    }
  }
}
```

LDP

Configure VRF-based authentication for LDP sessions at any of the following hierarchy levels:

- [edit routing-instances *vrf-instance* protocols ldp]
- [edit routing-instances *vrf-instance* protocols ldp session *session-ip*]
- [edit routing-instances *vrf-instance* protocols ldp session-group *ip-prefix*]

In this example, TCP-AO authenticates TCP connections in a VRF instance called *vrf-two*. It authenticates TCP connections to the address 10.0.1.1 as well as any address in the 10.0.0.0/24 subnet.

```
[edit routing-instances]
vrf-two {
  protocols {
    ldp {
      session 10.0.1.1 {
        authentication-algorithm ao;
        authentication-key-chain tcpao;
      }
      session-group 10.0.0.0/24 {
        authentication-algorithm ao;
        authentication-key-chain tcpao;
      }
    }
  }
}
```

Release History Table

Release	Description
22.4R1	Starting in Junos OS Evolved Release 22.4R1, you can configure TCP-AO or TCP MD5 authentication with an IP subnet to include the entire range of addresses under that subnet.
22.4R1	Starting in Junos OS Evolved Release 22.4R1, TCP authentication is VRF aware.

RELATED DOCUMENTATION

Configuring Junos OS to Extend the Default Port Address Range 78
Protocol Redirect Messages 20

TCP Authentication Option (TCP-AO)

SUMMARY

Learn about TCP Authentication Option (TCP-AO) for BGP and LDP sessions.

IN THIS SECTION

- [TCP-AO for BGP and LDP Sessions | 45](#)
- [Example: Configure a Keychain \(TCP-AO\) | 48](#)
- [Example: Use TCP-AO to Authenticate a BGP Session | 52](#)
- [Example: Use TCP-AO to Authenticate an LDP Session | 59](#)
- [Example: Use TCP-AO to Authenticate RPKI Validation Sessions | 65](#)

TCP-AO for BGP and LDP Sessions

IN THIS SECTION

- [Benefits of TCP-AO | 45](#)
- [What is TCP-AO? | 46](#)
- [Configuration | 48](#)

The BGP and LDP protocols use TCP for transport. TCP-AO is a new authentication method proposed through *RFC5925, The TCP Authentication Option* to enhance the security and authenticity of TCP segments exchanged during BGP and LDP sessions. It also supports both IPv4 and IPv6 traffic.

Benefits of TCP-AO

TCP-AO provides the following benefits over TCP MD5:

- **Stronger algorithms**—Supports multiple stronger authentication algorithms such as HMAC-SHA-1-96 and AES-128-CMAC-96 (mandated by *RFC5925, The TCP Authentication Option*). HMAC-

SHA-1-96 is a hash-based MAC and AES-128-CMAC-96 is a cipher-based MAC, thus making the message digest more complex and secure than the digest created by using the MD5 algorithm.

- **Two-Fold security**—In the TCP-AO method, the configured Authentication algorithm is used in two stages: Once to generate an internal traffic key from a user-configured key and then to generate a message digest using the generated traffic key, whereas in the TCP MD5 method, the MD5 algorithm generates a message digest using its user-configured key.
- **Better Key Management and Agility**—You can configure up to 64 keys for a session and you can add them at any time during the lifetime of a session. It provides a simple key coordination mechanism by giving the ability to change keys (move from one key to another) within the same connection without causing any TCP connection closure. Changing TCP MD5 keys during an established connection might cause a flap or restart in the connection.
- **Suitable for long-lived connections**—More suitable for long-lived connections for routing protocols such as BGP and LDP and across repeated instances of a single connection.

What is TCP-AO?

TCP-AO provides a framework to:

- Support multiple stronger algorithms, such as HMAC-SHA1 and AES-128 to create an internal traffic key and message digest.
- Add a new user-configured key to re-generate internal traffic keys for an established connection and a mechanism to synchronize key change between BGP or LDP peers.

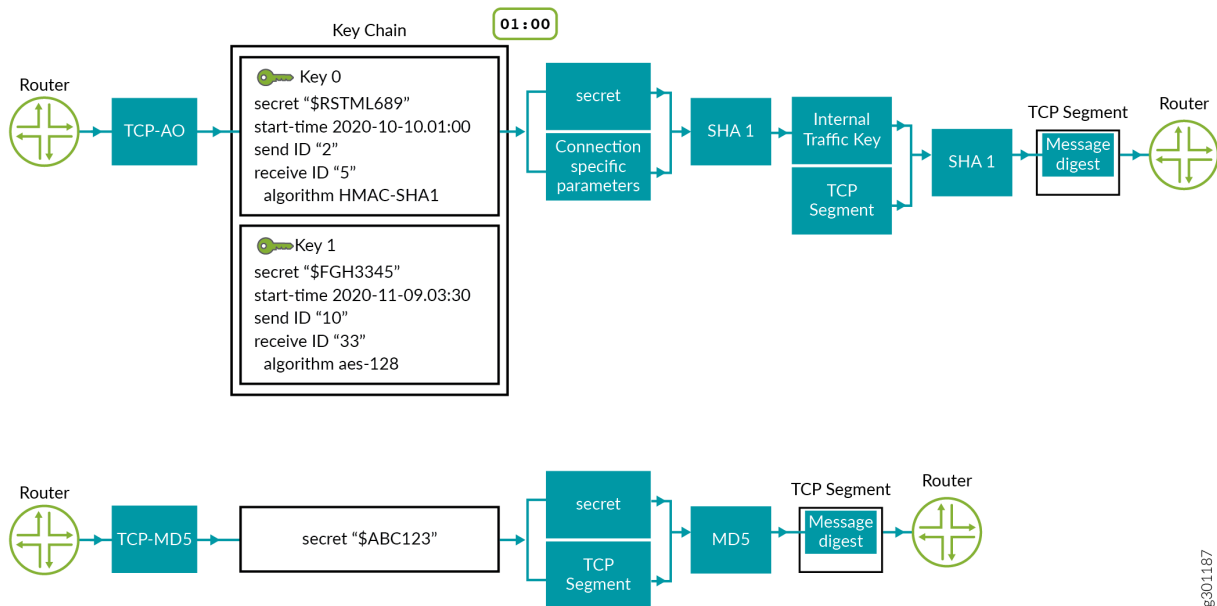
In earlier releases, Junos devices only supported the TCP MD5 authentication method for BGP and LDP sessions. The MD5 method supports only the MD5 algorithm, which is less secure than TCP-AO. In addition, changing a MD5 key normally disrupts the TCP session, unlike TCP-AO. TCP MD5 is defined in *RFC2385, Protection of BGP Sessions via the TCP MD5 Signature Option*. For more information about TCP MD5, see ["TCP" on page 34](#).

NOTE:

- While Junos devices support both the TCP-AO and TCP MD5 authentication methods, you cannot use both at the same time for a given connection.
- TCP-AO supports [Nonstop Active Routing](#).

The following diagram explains the difference between TCP-AO and TCP MD5 authentication. The first flow shows the configuration and processing flow for TCP-AO and the second flow shows the configuration and processing flow for TCP-MD5.

Figure 3: TCP-AO in comparison with TCP MD5



Below is an explanation of the processing flows shown in Figure 1:

- **TCP-AO**—The user has configured two keys in the keychain (key 0 and key 1) with all required parameters. The keychain supports two algorithms: HMAC SHA1 and AES-128 (mandated per RFC5925). TCP fetches key 0, which is the key that is currently active, as shown by the timestamp in the figure. In the example, key 0 is configured with HMAC-SHA1.

SHA1 takes the "secret" (from the key 0 configuration) and connection specific parameters for encryption and generates an internal traffic key.

SHA1 again encrypts the internal traffic key and the TCP segment to generate the message digest. The digest is copied to the TCP-AO MAC field of the TCP-AO option in the TCP segment. The segment is then sent to the receiving device.

- **TCP-MD5**—The user has configured a single key because TCP MD5 option supports only one key for a connection. Further, it only supports the MD5 algorithm. The MD5 algorithm takes the "secret" from the key and the TCP segment for encryption and generates a message digest. This message digest is then copied to MD5 digest field in the TCP segment and is sent to the receiving device.

Configuration

First, configure a keychain. Then apply TCP-AO to the BGP or LDP session.

To configure a keychain for TCP-AO (with one key), configure the following statement at the [edit security] hierarchy level.

```
[edit security]
user@router# set authentication-key-chains key-chain key-chain key id secret secretpassword
start-time YYYY-MM-DD.HH:MM algorithm ao ao-attribute send-id send-id recv-id recv-id
cryptographic-algorithm cryptographic-algorithm tcp-ao-option enabled
```

To apply TCP-AO to a BGP session (with the configured keychain), configure the following statement at the [edit protocols] hierarchy level.

```
[edit protocols]
user@router# set bgp group group neighbor neighbor authentication-algorithm ao
user@router# set bgp group group neighbor neighbor authentication-key-chain key-chain
```

To apply TCP-AO to an LDP session (with the configured keychain), configure the following statement at the [edit protocols] hierarchy level.

```
[edit protocols]
user@router# set ldp session session authentication-algorithm ao
user@router# set ldp session session authentication-key-chain key-chain
```

Example: Configure a Keychain (TCP-AO)

SUMMARY

This example shows you how to create a TCP-AO keychain to authenticate a BGP or LDP session.

This example uses the following hardware and software components:

- MX Series or PTX Series routers.

- Junos OS Release 20.3R1 or later version.

This example shows you how to create a TCP-AO keychain to authenticate a BGP or LDP session.

In this example, you can create a keychain `new_auth_key` with two keys, key 0 and key 1 on devices R1 and R2.

1. To create a keychain `new_auth_key` with the first key, (key 0):

NOTE: Copy the following commands, paste them into a text file, remove any line breaks and change any details necessary to match your network configuration, copy and paste the commands into the CLI.

R1

```
[edit security]
user@R1# set authentication-key-chains key-chain new_auth_key key 0 secret secretpassword
start-time 2020-10-10.03:00 algorithm ao ao-attribute send-id 3 recv-id 8 cryptographic-
algorithm hmac-sha-1-96 tcp-ao-option enabled
```

R2 (with send-id and recv-id values reversed)

```
[edit security]
user@R2# set authentication-key-chains key-chain new_auth_key key 0 secret secretpassword
start-time 2020-10-10.03:00 algorithm ao ao-attribute send-id 8 recv-id 3 cryptographic-
algorithm hmac-sha-1-96 tcp-ao-option enabled
```

Consider the following parameters while configuring a keychain:

Table 2: Keychain Parameters

Parameter	Description
key-chain	Enter a unique name.
key	Enter a unique key ID.
secret	Enter a unique password.

Table 2: Keychain Parameters (Continued)

Parameter	Description
start-time	Enter a unique time in <i>YYYY-MM-DD.HH:MM</i> format to specify the start time of the key.
algorithm	Enter algorithm <i>ao</i> .
send-id and recv-id	Enter any two numbers between 0 and 255. You must not use these numbers for any other key within that keychain.
cryptographic-algorithm	Choose either <i>hmac-sha-1-96</i> or <i>aes-128-cmac-96</i> .
tcp-ao-option	Choose <i>enabled</i> to enable the TCP-AO option.

- To add another key (key 1), after creating key 0:

R1

```
[edit security authentication-key-chains key-chain new_auth_key]
user@R1# set key 1 secret password start-time 2020-11-11.04:00 algorithm ao ao-attribute send-
id 1 recv-id 2 cryptographic-algorithm aes-128-cmac-96 tcp-ao-option enabled
```

R2 (with send-id and recv-id values reversed)

```
[edit security authentication-key-chains key-chain new_auth_key]
user@R2# set key 1 secret password start-time 2020-11-11.04:00 algorithm ao ao-attribute send-
id 2 recv-id 1 cryptographic-algorithm aes-128-cmac-96 tcp-ao-option enabled
```

- Enter `commit` from configuration mode on both devices to activate your changes.
- To verify the keychain `new_auth_key` with the 2 keys configured, use the `show security authentication-key-chains` command from configuration mode.

The following is sample output based on this example:

```
user@R1# show security authentication-key-chains key-chain new_auth_key {
  key 0 {
    secret "$RSTML689"; ## SECRET-DATA
```

```

start-time "2020-10-10.03:00:00 -0700";
algorithm ao;
ao-attribute {
    send-id 3;
    recv-id 8;
    tcp-ao-option enabled;
    cryptographic-algorithm hmac-sha-1-96;
}
}
key 1 {
    secret "$FFGH3345"; ## SECRET-DATA
    start-time "2020-11-11.04:00:00 -0800";
    algorithm ao;
    ao-attribute {
        send-id 1;
        recv-id 2;
        tcp-ao-option enabled;
        cryptographic-algorithm aes-128-cmac-96;
    }
}
}

```

You have successfully created a keychain!

To delete a keychain, use the `delete security authentication-key-chains key-chain key-chain-name` command from configuration mode.

NOTE:

- You can associate only one TCP-AO keychain with a BGP or LDP session during its life-time. You cannot point another keychain to the session in its life-time.
- We recommend a minimum interval of 30 minutes between the start-time of any two subsequent keys within a keychain.
- Once a keychain is configured and in use by a TCP connection, you cannot change the `send-id` or `recv-id` values of its active key. However, you can change the other parameters in the key, and any new connection associated with the updated keychain will take the updated parameters for its connection establishment.

- Starting in Junos OS Release 21.2R1, you can use the `tcpao-auth-mismatch allow-without-tcpao` to allow the connection establishment without TCP-AO if any one TCP endpoint does not have TCP-AO configured on it.

To display information about existing keychains (if any) from the operational mode, use the `show security keychain` command. Here is sample output:

```
user@R1> show security keychain
```

Keychain	Active-ID		Next-ID		Transition	Tolerance
	Send	Receive	Send	Receive		
new_auth_key	1	1	None	None	None	3600 (secs)

Example: Use TCP-AO to Authenticate a BGP Session

SUMMARY

This example shows you how to authenticate a BGP session using a TCP Authentication Option (TCP-AO) keychain.

IN THIS SECTION

- [Requirements | 52](#)
- [Overview | 53](#)
- [Configuration | 53](#)

Requirements

This example uses the following hardware and software components:

- MX Series or PTX Series routers.
- Junos OS Release 20.3R1 or later version.
- Configure a keychain `new_auth_key`. See ["Configure a Keychain \(TCP-AO\)" on page 48](#).

Overview

IN THIS SECTION

- [Topology | 53](#)

BGP uses TCP as its transport protocol. TCP-AO is a method you can use to authenticate BGP sessions. You can apply a TCP-AO keychain at the BGP neighbor or at BGP group levels of the configuration hierarchy.

Topology

Figure 4: Topology for BGP Authentication



Configuration

IN THIS SECTION

- [Verification | 57](#)

In this example, you associate the TCP-AO authentication keychain `new_auth_key` and authentication algorithm `ao` on both devices to authenticate a BGP session.

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, and then copy and paste the commands into the CLI.

R1

```
[edit]
set interfaces ge-0/0/1 description R1-to-R2-Link
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.11/32
set routing-options router-id 192.168.0.11
set routing-options autonomous-system 65500
set protocols bgp group ebgp_grp type external
set protocols bgp group ebgp_grp peer-as 65501
set protocols bgp group ebgp_grp neighbor 192.0.2.2
set protocols bgp group ebgp_grp neighbor 192.0.2.2 authentication-key-chain new_auth_key
set protocols bgp group ebgp_grp neighbor 192.0.2.2 authentication-algorithm ao
```

R2

```
[edit]
set interfaces ge-0/0/1 description R2-to-R1-Link
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.2/30
set interfaces lo0 unit 0 family inet address 192.168.0.12/32
set routing-options router-id 192.168.0.12
set routing-options autonomous-system 65501
set protocols bgp group ebgp_grp type external
set protocols bgp group ebgp_grp peer-as 65500
set protocols bgp group ebgp_grp neighbor 192.0.2.1
set protocols bgp group ebgp_grp neighbor 192.0.2.1 authentication-key-chain new_auth_key
set protocols bgp group ebgp_grp neighbor 192.0.2.1 authentication-algorithm ao
```

Step-By-Step Procedure

1. Enter configuration mode.
2. Configure basic settings such as the interface IP address, interface description, a loopback address, router-ID, AS number on both devices.

R1

```
[edit]
user@R1# set interfaces ge-0/0/1 description R1-to-R2-Link
user@R1# set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/30
user@R1# set interfaces lo0 unit 0 family inet address 192.168.0.11/32
```

```
user@R1# set routing-options router-id 192.168.0.11
user@R1# set routing-options autonomous-system 65500
```

R2

```
[edit]
user@R2# set interfaces ge-0/0/1 description R2-to-R1-Link
user@R2# set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.2/30
user@R2# set interfaces lo0 unit 0 family inet address 192.168.0.12/32
user@R2# set routing-options router-id 192.168.0.12
user@R2# set routing-options autonomous-system 65501
```

3. Configure an EBGp between R1 and R2.

R1

```
[edit]
user@R1# set protocols bgp group ebgp_grp type external
user@R1# set protocols bgp group ebgp_grp peer-as 65501
user@R1# set protocols bgp group ebgp_grp neighbor 192.0.2.2
```

R2

```
[edit]
user@R2# set protocols bgp group ebgp_grp type external
user@R2# set protocols bgp group ebgp_grp peer-as 65500
user@R2# set protocols bgp group ebgp_grp neighbor 192.0.2.1
```

4. Associate the authentication keychain `new_auth_key` and the authentication algorithm `ao` to the BGP session on both devices.

R1

```
[edit]
user@R1# set protocols bgp group ebgp_grp neighbor 192.0.2.2 authentication-key-chain
new_auth_key
user@R1# set protocols bgp group ebgp_grp neighbor 192.0.2.2 authentication-algorithm ao
```

R2

```
[edit]
user@R2# set protocols bgp group ebgp_grp neighbor 192.0.2.1 authentication-key-chain
new_auth_key
user@R2# set protocols bgp group ebgp_grp neighbor 192.0.2.1 authentication-algorithm ao
```

5. Enter `commit` from configuration mode on both devices.

Once you commit the configurations statements on both devices the BGP session should establish using the TCP-AO authentication method.

Results

Confirm your configurations by using the `show interfaces`, `show routing-options`, and `show protocols` commands from configuration mode.

user@R1# **show interfaces**

```
ge-0/0/1 {
  description R1-to-R2-Link;
  unit 0 {
    family inet {
      address 192.0.2.1/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.11/32;
    }
  }
}
```

user@R1# **show routing-options**

```
autonomous-system 65500;
```

user@R1# **show protocols**

```
bgp {  
  group ebgp_grp {  
    type external;  
    peer-as 65500;  
    neighbor 192.0.2.1 {  
      authentication-algorithm ao;  
      authentication-key-chain new_auth_key;  
    }  
  }  
}
```

```
bgp {  
  group ebgp_grp {  
    type external;  
    peer-as 65551;  
    neighbor 192.0.2.2 {  
      authentication-algorithm ao;  
      authentication-key-chain new_auth_key;  
    }  
  }  
}
```

Verification

IN THIS SECTION

- [Verify BGP Session Establishment | 58](#)
- [Verify BGP Session is Using TCP-AO | 58](#)

Verify BGP Session Establishment

Purpose

Confirm BGP session establishment output after enabling TCP-AO.

Action

View a BGP summary of BGP session state with the `show bgp summary` operational mode command.

```
user@R1> show bgp summary
Threading mode: BGP I/O
Default eBGP mode: advertise - accept, receive - accept
Groups: 1 Peers: 1 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State   Pending
inet.0
              0          0          0          0          0          0
Peer          AS      InPkt    OutPkt    OutQ   Flaps Last Up/Dwn State|#Active/
Received/Accepted/Damped...
192.0.2.2      65501         6         4         0         0      1:19 Establ
inet.0: 0/0/0/0
```

Meaning

The highlighted output values indicate that BGP has successfully established a session with the TCP-AO authentication method 1:19 minute ago.

Verify BGP Session is Using TCP-AO

Purpose

Verify a BGP neighbor is authenticated with the TCP-AO keychain.

Action

Use the `show bgp neighbor neighbor` command to view configuration details for BGP peers. To filter only authentication-specific details in the output, use the pipe (`|`) function and match on authentication, as shown:

```
user@R1> show bgp neighbor 192.0.2.2 | match authentication
Authentication key chain: new_auth_key
Authentication algorithm: ao
```

Meaning

The output indicates that authentication keychain `new_auth_key` and Authentication algorithm `ao` is applied to the BGP neighbor `192.0.2.2`.

Example: Use TCP-AO to Authenticate an LDP Session

SUMMARY

This example shows you how to authenticate an LDP session using a TCP Authentication Option (TCP-AO) keychain.

IN THIS SECTION

- [Requirements | 59](#)
- [Overview | 60](#)
- [Configuration | 60](#)
- [Verification | 64](#)

Requirements


This example uses the following hardware and software components:

- MX Series or PTX Series routers.
- Junos OS Release 20.3R1 or later version.
- Configure a keychain `new_auth_key`. See "[Configure a Keychain \(TCP-AO\)](#)" on page 48.

Overview

Label Distribution Protocol (LDP) is an MPLS signaling protocol. It allows routers to establish label-switched paths (LSPs) through a network. TCP-AO helps enhance the security of sessions created among LDP peers.

Figure 5: Topology for LDP Configuration



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 60](#)
- [Step-By-Step Procedure | 61](#)
- [Results | 63](#)

In this example, you associate the TCP-AO authentication keychain `new_auth_key` and authentication algorithm `ao` to both devices to authenticate their LDP session.

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, and then copy and paste the commands into the CLI.

R1

```
[edit]
set interfaces ge-0/0/1 description R1-to-R2-Link
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.11/32
set routing-options router-id 192.168.0.11
```

```

set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set protocols ldp session 192.168.0.12 authentication-algorithm ao
set protocols ldp session 192.168.0.12 authentication-key-chain new_auth_key
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

R2

```

[edit]
set interfaces ge-0/0/1 description R2-to-R1-Link
set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.2/30
set interfaces ge-0/0/1 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 192.168.0.12/32
set routing-options router-id 192.168.0.12
set protocols ldp interface ge-0/0/1.0
set protocols ldp interface lo0.0
set protocols ldp session 192.168.0.11 authentication-algorithm ao
set protocols ldp session 192.168.0.11 authentication-key-chain new_auth_key
set protocols ospf area 0.0.0.0 interface ge-0/0/1.0
set protocols ospf area 0.0.0.0 interface lo0.0

```

Step-By-Step Procedure

1. Enter configuration mode.
2. Configure basic setup such as device interface, loopback, interface description, router ID, AS number on R1 and R2.

R1

```

[edit]
user@R1# set interfaces ge-0/0/1 description R1-to-R2-Link
user@R1# set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.1/30
user@R1# set interfaces lo0 unit 0 family inet address 192.168.0.11/32
user@R1# set routing-options router-id 192.168.0.11

```

R2

```

[edit]
user@R2# set interfaces ge-0/0/1 description R2-to-R1-Link

```

```

user@R2# set interfaces ge-0/0/1 unit 0 family inet address 192.0.2.2/30
user@R2# set interfaces lo0 unit 0 family inet address 192.168.0.12/32
user@R2# set routing-options router-id 192.168.0.12

```

3. Configure MPLS and LDP on both devices.

R1

```

[edit]
user@R1# set interfaces ge-0/0/1 unit 0 family mpls
user@R1# set protocols ldp interface ge-0/0/1.0
user@R1# set protocols ldp interface lo0.0

```

R2

```

[edit]
user@R2# set interfaces ge-0/0/1 unit 0 family mpls
user@R2# set protocols ldp interface ge-0/0/1.0
user@R2# set protocols ldp interface lo0.0

```

4. Configure an interior gateway protocol (IGP) to advertise loopback address reachability. In this example, we configure OSPF.

R1

```

[edit protocols]
user@R1# set ospf area 0.0.0.0 interface ge-0/0/1.0
user@R1# set ospf area 0.0.0.0 interface lo0.0 passive

```

R2

```

[edit protocols]
user@R2# set ospf area 0.0.0.0 interface ge-0/0/1.0
user@R2# set ospf area 0.0.0.0 interface lo0.0 passive

```

5. Associate authentication-key-chain new_auth_key and authentication-algorithm ao with the label space ID of R1 and R2.

R1

```
[edit protocols]
user@R1# set ldap session 192.168.0.12 authentication-algorithm ao
user@R1# set ldap session 192.168.0.12 authentication-key-chain new_auth_key
```

R2

```
[edit protocols]
user@R2# set ldap session 192.168.0.11 authentication-algorithm ao
user@R2# set ldap session 192.168.0.11 authentication-key-chain new_auth_key
```

6. Enter `commit` from the configuration mode on both devices.

Results

Confirm your configuration by using the `show interfaces`, `show routing-options` and `show protocols` commands.

user@R1# **show interfaces**

```
ge-0/0/1 {
  description R1-to-R2-Link;
  unit 0 {
    family inet {
      address 192.0.2.1/30;
    }
  }
}
lo0 {
  unit 0 {
    family inet {
      address 192.168.0.11/32;
    }
  }
}
```

user@R1# **show routing-options**

```
router-id 192.168.0.11;
```

user@R1# **show protocols**

```
ldp {  
    interface ge-0/0/1.0;  
    interface lo0.0 passive;  
        authentication-algorithm ao;  
        authentication-key-chain new_auth_key;  
    {  
    {  
    ospf {  
        area 0.0.0.0 {  
            interface ge-0/0/1.0;  
            interface lo0.0;  
        {  
    {
```

Verification

IN THIS SECTION

- [Verify LDP Session | 64](#)

Verify LDP Session

Purpose

Verify LDP session Establishment with TCP-AO.

Action

Use the `show ldp session detail operational mode` command to verify the LDP session is correctly established.

```
user@R1> show ldp session detail
```

```
Address: 192.168.0.12, State: Operational, Connection: Open, Hold time: 22  
Session ID: 192.168.0.11:0--192.168.0.12:0
```

```

Next keepalive in 2 seconds
Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
Neighbor types: discovered
Keepalive interval: 10, Connect retry interval: 1
Local address: 192.168.0.11, Remote address: 192.168.0.12
Up for 01:11:59
Last down 01:13:12 ago; Reason: authentication key was changed
Number of session flaps: 2
Capabilities advertised: none
Capabilities received: none
Protection: disabled
Session flags: none
Authentication type: new_auth_key(ao key-chain, 192.168.0.12/32)
Local - Restart: disabled, Helper mode: enabled
Remote - Restart: disabled, Helper mode: enabled
Local maximum neighbor reconnect time: 120000 msec
Local maximum neighbor recovery time: 240000 msec
Local Label Advertisement mode: Downstream unsolicited
Remote Label Advertisement mode: Downstream unsolicited
Negotiated Label Advertisement mode: Downstream unsolicited
MTU discovery: disabled
Nonstop routing state: Not in sync
Next-hop addresses received:
    192.0.2.2
    192.168.0.12
    128.49.110.110

```

Meaning

The output indicates that LDP session is established.

Example: Use TCP-AO to Authenticate RPKI Validation Sessions

IN THIS SECTION

- [Overview | 66](#)
- [Requirements | 66](#)

- Topology | 66
- Configuration | 67

Overview

Resource Public Key Infrastructure (RPKI) is a public key infrastructure framework that is designed to secure the Internet's routing infrastructure, specifically the BGP. RPKI provides a way to connect Internet number resource information, such as IP Addresses, to a trust anchor. By using RPKI, legitimate holders of number resources are able to control the operation of Internet routing protocols to prevent route hijacking and other attacks.

Starting in Junos OS Release 22.2R1, you can authenticate RPKI sessions by using TCP Authentication Option (TCP-AO) and keychain.

This example shows you how to authenticate an RPKI validation session using a TCP-AO keychain. We'll be establishing an authenticated RPKI session between a client device (R1) and a server (R2).


Requirements

This example uses the following hardware and software components:

- 2 MX Series routers
- Junos OS Release 22.2R1 or later version.

Topology

Figure 6: Topology for Authenticated RPKI Session



Configuration

IN THIS SECTION

- [Verification | 72](#)

In this example, you must associate the TCP-AO authentication keychain `new_auth_key` and authentication algorithm `ao` on both devices to authenticate an RPKI connection.

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, and then copy and paste the commands into the CLI.

R1

```
[edit]
set system host-name R1
set interfaces ge-0/0/0 description R1-to-R2-Link
set interfaces ge-0/0/0 unit 0 family inet address 192.168.12.1/24
set interfaces lo0 unit 0 family inet address 172.16.255.1/32
set routing-options router-id 172.16.255.1
set routing-options autonomous-system 65500
set security authentication-key-chains key-chain new_auth_key key 0 secret "$ABC123"
set security authentication-key-chains key-chain new_auth_key key 0 start-time
"2022-5-18.04:00:00 -0700"
set security authentication-key-chains key-chain new_auth_key key 0 algorithm ao
set security authentication-key-chains key-chain new_auth_key key 0 ao-attribute send-id 3
set security authentication-key-chains key-chain new_auth_key key 0 ao-attribute rcv-id 8
set security authentication-key-chains key-chain new_auth_key key 1 secret "$ABC123"
set security authentication-key-chains key-chain new_auth_key key 1 start-time
"2022-6-18.04:00:00 -0700"
set security authentication-key-chains key-chain new_auth_key key 1 algorithm ao
set security authentication-key-chains key-chain new_auth_key key 1 ao-attribute send-id 1
set security authentication-key-chains key-chain new_auth_key key 1 ao-attribute rcv-id 2
set routing-options validation group to_servers session 192.168.12.2 port 8282
set routing-options validation group to_servers session 192.168.12.2 authentication-algorithm ao
```

```
set routing-options validation group to_servers session 192.168.12.2 authentication-key-chain
new_auth_key
```

R2

```
[edit]
set system host-name R2
set logical-systems rv_server_1 interfaces ge-0/0/0 unit 0 family inet address 192.168.12.2/24
set interfaces lo0 unit 0 family inet address 172.16.255.2/32
set routing-options router-id 172.16.255.2
set routing-options autonomous-system 65501
set logical-systems rv_server_1 routing-options validation local-cache listen-port 8282
set logical-systems rv_server_1 routing-options validation group to_dut session 192.168.12.1
local-cache
set logical-systems rv_server_1 routing-options validation group to_dut session 192.168.12.1
local-address 192.168.12.2
set security authentication-key-chains key-chain new_auth_key key 0 secret "$ABC123"
set security authentication-key-chains key-chain new_auth_key key 0 start-time
"2022-5-18.04:00:00 -0700"
set security authentication-key-chains key-chain new_auth_key key 0 algorithm ao
set security authentication-key-chains key-chain new_auth_key key 0 ao-attribute send-id 8
set security authentication-key-chains key-chain new_auth_key key 0 ao-attribute rcv-id 3
set security authentication-key-chains key-chain new_auth_key key 1 secret "$ABC123"
set security authentication-key-chains key-chain new_auth_key key 1 start-time
"2022-6-18.04:00:00 -0700"
set security authentication-key-chains key-chain new_auth_key key 1 algorithm ao
set security authentication-key-chains key-chain new_auth_key key 1 ao-attribute send-id 2
set security authentication-key-chains key-chain new_auth_key key 1 ao-attribute rcv-id 1
set logical-systems rv_server_1 routing-options validation group to_dut session 192.168.12.1
authentication-algorithm ao
set logical-systems rv_server_1 routing-options validation group to_dut session 192.168.12.1
authentication-key-chain new_auth_key
```

Step-By-Step Procedure

1. Configure basic settings such as, interfaces, a loopback address, router-ID, and AS number on both devices. On R2, we configure logical systems interface for the server.

R1

```
[edit]
user@R1# set system host-name R1
user@R1# set interfaces ge-0/0/0 description R1-to-R2-Link
user@R1# set interfaces ge-0/0/0 unit 0 family inet address 192.168.12.1/24
user@R1# set interfaces lo0 unit 0 family inet address 172.16.255.1/32
user@R1# set routing-options router-id 172.16.255.1
user@R1# set routing-options autonomous-system 65500
```

R2

```
[edit]
user@R2# set system host-name R2
user@R2# set logical-systems rv_server_1 interfaces ge-0/0/0 unit 0 family inet address
192.168.12.2/24
user@R2# set interfaces lo0 unit 0 family inet address 172.16.255.2/32
user@R2# set routing-options router-id 172.16.255.2
user@R2# set routing-options autonomous-system 65501
```

2. Configure a TCP session on the client device (R1) with the RPKI server (R2) with an alternative TCP port number.

R1

```
[edit]
user@R1# set routing-options validation group to_servers session 192.168.12.2 port 8282
```

3. On the server R2, configure an RPKI session with the client R1 for origin validation.

```
[edit]
user@R2# set logical-systems rv_server_1 routing-options validation group to_dut session
192.168.12.1 local-cache
user@R2# set logical-systems rv_server_1 routing-options validation group to_dut session
192.168.12.1 local-address 192.168.12.2
```

4. Create a keychain new_auth_key with the first key, (key 0):

R1

```
[edit security]
user@R1# set authentication-key-chains key-chain new_auth_key key 0 secret secretpassword
start-time 2022-5-18.04:00 algorithm ao ao-attribute send-id 3 rcv-id 8
```

R2 (with send-id and rcv-id values reversed)

```
[edit security]
user@R2# set authentication-key-chains key-chain new_auth_key key 0 secret secretpassword
start-time 2022-5-18.04:00 algorithm ao ao-attribute send-id 8 rcv-id 3
```

5. To add another key (key 1), after creating key 0:

R1

```
[edit security authentication-key-chains key-chain new_auth_key]
user@R1# set key 1 secret secretpassword start-time 2022-6-18.04:00 algorithm ao ao-attribute
send-id 1 rcv-id 2
```

R2 (with send-id and rcv-id values reversed)

```
[edit security authentication-key-chains key-chain new_auth_key]
user@R2# set key 1 secret secretpassword start-time 2022-6-18.04:00 algorithm ao ao-attribute
send-id 2 rcv-id 1
```

6. Apply the configured keychain `new_auth_key` and authentication algorithm `ao` on both R1 and R2.

R1

```
[edit]
user@R1# set routing-options validation group to_servers session 192.168.12.2 authentication-
algorithm ao
user@R1# set routing-options validation group to_servers session 192.168.12.2 authentication-
key-chain new_auth_key
```

R2

```
[edit]
user@R2# set logical-systems rv_server_1 routing-options validation group to_dut session
192.168.12.1 authentication-algorithm ao
user@R2# set logical-systems rv_server_1 routing-options validation group to_dut session
192.168.12.1 authentication-key-chain new_auth_key
```

7. Enter `commit` from configuration mode on both devices to activate your changes.
8. To verify the keychain `new_auth_key` with the two keys configured, use the `show security authentication-key-chains` command from configuration mode.

Results

Check the results of the keychain configuration on R1:

```
user@R1# show security authentication-key-chains
```

```
key-chain new_auth_key {
  key 0 {
    secret "$ABC123"; ## SECRET-DATA
    start-time "2022-5-18.04:00:00 -0700";
    algorithm ao;
    ao-attribute {
      send-id 3;
      rcv-id 8;
    }
  }
  key 1 {
    secret "$ABC123"; ## SECRET-DATA
    start-time "2022-6-18.04:00:00 -0700";
    algorithm ao;
    ao-attribute {
      send-id 1;
      rcv-id 2;
    }
  }
}
```

Confirm the remaining configurations applied on R1 by using the following commands:

user@R1# **show interfaces**

```
ge-0/0/0 {  
  description R1-to-R2-Link;  
  unit 0 {  
    family inet {  
      address 192.168.12.1/24;  
    }  
  }  
}  
lo0 {  
  unit 0 {  
    family inet {  
      address 172.16.255.1/32;  
    }  
  }  
}
```

user@R1# **show routing-options**

```
router-id 172.16.255.1;  
autonomous-system 65500;  
validation {  
  group to_servers {  
    session 192.168.12.2 {  
      authentication-algorithm ao;  
      authentication-key-chain new_auth_key;  
      port 8282;  
    }  
  }  
}
```

Verification

IN THIS SECTION

● Purpose | 73

● Action | 73

● Meaning | 73

Purpose

Verify the session is established with TCP-AO keychain and algorithm configured on both the peers.

Action

View a validated session by using the `show validation session 192.168.12.2 detail operational mode` command.

```
user@R1> show validation session 192.168.12.2 detail
Session 192.168.12.2, State: up, Session index: 2
  Group: to_servers, Preference: 100
  Port: 8282
  Refresh time: 300s
  Hold time: 600s
  Record Life time: 3600s
  Serial (Full Update): 6
  Serial (Incremental Update): 6
  Authentication key-chain: new_auth_key
    Session flaps: 1
    Session uptime: 2d 01:40:05
    Last PDU received: 00:04:59
    IPv4 prefix count: 0
    IPv6 prefix count: 0
```

Meaning

The output indicates the session is up with the configured keychain `new_auth_key`.

Release History Table

Release	Description
22.4R1	Starting in Junos OS Evolved Release 22.4R1, you can configure TCP-AO or TCP MD5 authentication with an IP subnet to include the entire range of addresses under that subnet.

22.4R1	Starting in Junos OS Evolved Release 22.4R1, TCP authentication is VRF aware.
--------	---

RELATED DOCUMENTATION

| [authentication-key-chains \(TCP-AO\)](#) | **101**

3

CHAPTER

Configure Port Security

[System Settings](#) | 76

[Password Authentication for Console Access to PICs](#) | 82

System Settings

IN THIS SECTION

- [Specifying the Physical Location of the Switch | 76](#)
- [Modifying the Default Time Zone for a Router or Switch Running Junos OS | 77](#)
- [Configuring Junos OS to Extend the Default Port Address Range | 78](#)
- [Configuring Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets | 79](#)
- [Rebooting and Halting a Device | 80](#)

Specifying 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.

The following example shows how to specify the physical location of the switch:

```
[edit system]
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
  lata service-area;
  latitude degrees;
  longitude degrees;
  npa-nxx number;
  postal-code postal-code;
  rack number;
  vcoord vertical-coordinate;
}
```

SEE ALSO

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

Modifying the Default Time Zone for a Router or Switch Running Junos OS

The default local time zone on the router or switch is UTC (Coordinated Universal Time, formerly known as Greenwich Mean Time, or GMT).

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

```
[edit 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

[Understanding NTP Time Servers](#)

[Updating the IANA Time Zone Database on Junos OS Devices](#)

Configuring Junos OS to 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 system internet-options]
source-port upper-limit upper-limit;
```

`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

[TCP | 34](#)

[ARP Learning and Aging Options | 6](#)

Configuring Junos OS to 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]
default-address-selection;
```

If you include the `default-address-selection` statement in the configuration, the Junos OS 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
```

<code>all-members</code>	Halt all virtual chassis members
<code>at</code>	Time at which to perform the operation
<code>backup-routing-engine</code>	Halt backup Routing Engine
<code>both-routing-engines</code>	Halt both Routing Engines
<code>in</code>	Number of minutes to delay before operation
<code>local</code>	Halt local virtual chassis member
<code>member</code>	Halt specific virtual chassis member (0..9)
<code>message</code>	Message to display to all users
<code>other-routing-engine</code>	Halt other Routing Engine
<code> </code>	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](#) | 23

Password Authentication for Console Access to PICs

By default, there is no password setting for console access. To configure console access to the Physical Interface Cards (PICs), include the `pic-console-authentication` statement at the `[edit system]` hierarchy level:

```
[edit system]
pic-console-authentication {
  (encrypted-password "password" | plain-text-password);
}
```

`encrypted-password "password"`—Use Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.

You cannot configure a blank password for `encrypted-password` using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

`plain-text-password`—Use a plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.

RELATED DOCUMENTATION

| [Configuring Junos OS to Set Console and Auxiliary Port Properties](#)

4

CHAPTER

Configuration Statements

[allow-6pe-traceroute](#) | 85

[allow-v4mapped-packets](#) | 86

[arp](#) | 88

[arp-max-cache](#) | 93

[arp-new-hold-limit](#) | 95

[arp-system-cache-limit](#) | 96

[auxiliary](#) | 99

[authentication-key-chains \(TCP-AO\)](#) | 101

[console \(System Ports\)](#) | 103

[default-address-selection](#) | 106

[diag-port-authentication](#) | 108

[extended-statistics](#) | 110

[icmp \(System\)](#) | 112

[icmp \(Error Message Rate Limit\)](#) | 114

[icmp6 \(Error Message Rate Limit\)](#) | 116

[internet-options](#) | 118

[nd-maxucast-retry](#) | 122

[no-multicast-echo](#) | 124

[non-subscriber-no-reply](#) | 125

[no-ping-record-route](#) | 127

no-ping-time-stamp | 128

path-mtu-discovery (Tunnel) | 130

proactive-arp-detection | 132

tcpao-auth-mismatch | 133

tcp-mss | 135

allow-6pe-traceroute

IN THIS SECTION

- [Syntax | 85](#)
- [Hierarchy Level | 85](#)
- [Description | 85](#)
- [Required Privilege Level | 86](#)

Syntax

```
allow-6pe-traceroute;
```

Hierarchy Level

```
[edit system]
```

Description

Allow IPv4-mapped IPv6 source addresses in an ICMPv6 echo request TTL expired packets.

In a dual-stack IPv6 network connected over an IPv4 MPLS network, the P routers in the IPv4 MPLS backbone do not have an IPv6 family. Consequently, the transit P routers are not shown in the output when you do an IPv6 traceroute. To generate an ICMPv6 echo request and a TTL expired response packet to and from the intermediate transit routers in the 6PE network, you must configure this statement along with the `allow-v4-mapped` statement.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[IPv6 Features | 29](#)

[allow-v4mapped-packets | 86](#)

allow-v4mapped-packets

IN THIS SECTION

- [Syntax | 86](#)
- [Hierarchy Level | 87](#)
- [Description | 87](#)
- [Options | 87](#)
- [Required Privilege Level | 87](#)
- [Release Information | 87](#)

Syntax

```
allow-v4mapped-packets;
```

Hierarchy Level

[edit system]

Description

Enable the processing of IPv4-mapped IPv6 packets.

Options

None

- **Default:** IPv4-mapped IPv6 address processing is disabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

[IPv6 Features | 29](#)

[allow-6pe-traceroute | 85](#)

arp

IN THIS SECTION

- [Syntax | 88](#)
- [Syntax \(EX Series\) | 89](#)
- [Hierarchy Level | 89](#)
- [Description | 89](#)
- [Options | 90](#)
- [Required Privilege Level | 92](#)
- [Release Information | 92](#)

Syntax

```
arp ip-address (mac | multicast-mac) mac-address publish;
```

```
arp {  
    aging-timer minutes;  
    arp-retries count;  
    gratuitous-arp-delay seconds;  
    gratuitous-arp-on-ifup;  
    interfaces {  
        interface-name {  
            aging-timer minutes;  
        }  
    }  
    non-subscriber-no-reply;  
    passive-learning;  
    purging;  
    unicast-mode-on-expire;  
}
```

Syntax (EX Series)

```
arp {
    aging-timer minutes;
}
```

Hierarchy Level

```
[edit system]
[edit interfaces interface-name unit logical-unit-number family inet address address]
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number
family inet address address]
```

NOTE: The edit `logical-systems` hierarchy is not available on QFabric systems.

Description

For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries mapping IP addresses to MAC addresses. IPv4 networks use ARP to map IP network addresses to physical (MAC) addresses. An address is resolved when the host device receives a proper ARP reply in response to the ARP request that it sent to a broadcast Ethernet address. The resolved addresses are stored in the ARP table for a configurable period of time. When an entry is close to expiration, it triggers the host to broadcast another ARP request to update the entry for that address. Only the intended receiver responds to the broadcast request; other recipients silently drop the request packet.

You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the `family inet` statement. By including the `arp` statement at the `[edit`

interfaces *interface-name* unit *logical-unit-number* family inet policer] hierarchy level, you can apply a specific ARP-packet policer to an interface. This feature is not available on EX Series switches.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the unnumbered-address statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] hierarchy level.

NOTE: For EX-Series switches, set only the time interval between ARP updates.

Options

aging-timer	Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.
arp-retries count	Configures the maximum number of retries the device attempts for ARP. The default retry attempts for ARP is 4.
ip-address	IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing address statement.
mac <i>mac-address</i>	MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0000.5e00.5355 or 00:00:5e:00:53:55.
multicast-mac <i>mac-address</i>	Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0000.5e00.5355 or 00:00:5e:00:53:55.
publish	(Optional) Have the router or switch reply to ARP requests for the specified IP address. If you omit this option, the router or switch uses the entry to reach the destination but does not reply to ARP requests.

NOTE: For unicast MAC addresses only, if you include the publish option, the router or switch replies to proxy ARP requests.

gratuitous-arp-delay	<p>Configure a delay for gratuitous ARP requests at the system level. By default, Junos OS sends gratuitous ARP requests immediately after network-related configuration changes are made on an interface (for example, a VLAN ID, MAC address, IP address change, or Aggregated Ethernet deployment). This might lead to the Packet Forwarding Engine dropping some initial request packets if the configuration updates have not been fully processed. To avoid such request packets being dropped, you can configure a delay in gratuitous ARP requests.</p> <ul style="list-style-type: none"> • Values: <ul style="list-style-type: none"> • <i>seconds</i>—Configure the ARP request delay in seconds. We recommend configuring a value in the range of 3 through 6 seconds.
gratuitous-arp-on-ifup	Add this statement to the [edit system arp] hierarchy to configure Junos OS to automatically issue a gratuitous ARP announcement when an interface is online.
interfaces	<p>Specify the ARP aging timer in minutes for a logical interface of family type inet.</p> <ul style="list-style-type: none"> • Values: aging-timer <i>minutes</i>—Time between ARP updates, in minutes. • Default: 20 • Range: 1 through 6,00,000
non-subscriber-no-reply	Configure the device to reply to ARP requests from subscribers only. Do not reply to ARP requests from non-subscribers.
passive-learning	Configure backup VRRP routers or switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the primary router fails, the backup router must learn all entries present in the ARP cache of the primary router. Configuring passive learning reduces transition delay when the backup router is activated. Learning of ARP mappings (IP-to-MAC address) by backup VRRP routers or switches for hosts sending the requests is disabled unless this statement is configured.
purging	Purge obsolete ARP entries from the cache when an interface or link goes offline.
unicast-mode-on-expire	<p>Send a unicast ARP request instead of the default broadcast request when an ARP cache entry ages out. When you include this option, the host device sends the requests only to the expected (currently cached) address. The ARP retry requests are unicast at intervals of 5 seconds. When you do not configure the unicast-mode-on-expire option, ARP retries are broadcast at intervals of 800 milliseconds.</p> <p>This option reduces the amount of broadcast traffic normally sent to resolve expiring addresses. It also supports a special use case where access nodes are configured not to</p>

forward broadcast ARP requests towards customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests.

NOTE: This option affects only the update requests. Initial ARP requests are broadcast as usual.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.0.

Statement introduced in Junos OS Release 11.1.

Statement introduced in Junos OS Release 14.1X53-D20.

arp-retries option introduced in Junos OS Evolved Release 22.1R1.

arp-retries option introduced in Junos OS Release 22.1R2.

RELATED DOCUMENTATION

[Configuring Static ARP Table Entries For Mapping IP Addresses to MAC Addresses](#)

[ARP Learning and Aging Options | 6](#)

[Adjusting the ARP Aging Timer | 8](#)

arp-max-cache

IN THIS SECTION

- [Syntax | 93](#)
- [Hierarchy Level | 93](#)
- [Description | 93](#)
- [Options | 94](#)
- [Required Privilege Level | 94](#)
- [Release Information | 94](#)

Syntax

```
arp-max-cache arp-max-cache;
```

Hierarchy Level

```
[edit interfaces name unit name family inet]
```

Description

The ARP cache limit for resolved next hops can be configured at an interface level. The benefit of configuring the ARP cache limit is to protect the device from DoS attacks.

NOTE: After modifying the default ARP next-hop cache limit on an interface, you must deactivate and then reactivate the interface for the newly configured values to take effect.

Options

arp-max-cache Indicates the maximum number of routes to be held in the ARP cache.

- **Default:** (ACX Series routers, EX2200, EX2200-C, EX3200, and EX3300 switches, SRX Series services gateways) 20,000
(EX4200, EX4300, EX4500, EX4550, and EX6210 switches, MX Series routers) 75,000
(Other platforms) 100,000
- **Range:** 1 through 2,000,000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Example: Configuring ARP Cache Protection | 10](#)

[arp | 88](#)

[arp | 88](#)

[arp-inspection \(MX Series\)](#)

[arp-new-hold-limit | 95](#)

[arp-system-cache-limit | 96](#)

arp-new-hold-limit

IN THIS SECTION

- [Syntax | 95](#)
- [Hierarchy Level | 95](#)
- [Description | 95](#)
- [Options | 96](#)
- [Required Privilege Level | 96](#)
- [Release Information | 96](#)

Syntax

```
arp-new-hold-limit number;
```

Hierarchy Level

```
[edit interfaces interface-name unit interface-unit-number family inet]
```

Description

The ARP cache limit for unresolved next hops can be configured at an interface level. The benefit of configuring the ARP cache limit is to protect the device from DoS attacks.

Options

number Indicates the new or maximum routes to be held, before getting added to the ARP cache.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Example: Configuring ARP Cache Protection | 10](#)

[arp | 88](#)

[arp-inspection \(MX Series\)](#)

[arp-system-cache-limit | 96](#)

[arp-max-cache | 93](#)

arp-system-cache-limit

IN THIS SECTION

● [Syntax | 97](#)

● [Hierarchy Level | 97](#)

● [Description | 97](#)

- Default | 97
- Options | 98
- Required Privilege Level | 98
- Release Information | 98

Syntax

```
arp-system-cache-limit number;
```

Hierarchy Level

```
[edit system]
```

Description

Specify the ARP cache next-hop limit at the system (global) level to restrict the number of next-hop routes. To configure the limit at the interface level, use the ["arp-max-cache" on page 93](#) statement.

Default

The default behavior of ARP is to remove the cache entry if it is not used within a certain period of time, and not to allow the cache to grow too large. You can also manage the number of ARP cache next-hop entries by configuring a limit to the maximum number of next hops that can be created.

Options

number Indicates the maximum number of routes to be held in the ARP cache.

- **Range:** 1 through 2,000,000
- **Default:** (ACX Series routers, EX2200, EX2200-C, EX3200, and EX3300 switches, SRX Series services gateways) 20,000
(EX4200, EX4300, EX4500, EX4550, and EX6210 switches, MX Series routers) 75,000
(Other platforms) 100,000

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Example: Configuring ARP Cache Protection | 10](#)

[arp | 88](#)

[arp-inspection \(MX Series\)](#)

[arp-max-cache | 93](#)

[arp-new-hold-limit | 95](#)

auxiliary

IN THIS SECTION

- [Syntax | 99](#)
- [Hierarchy Level | 99](#)
- [Description | 99](#)
- [Default | 100](#)
- [Options | 100](#)
- [Required Privilege Level | 100](#)
- [Release Information | 100](#)

Syntax

```
auxiliary {  
    disable;  
    insecure;  
    type terminal-type;  
    port-type (mini-usb | rj45);  
}
```

Hierarchy Level

```
[edit system ports]
```

Description

Configure the characteristics of the auxiliary port.

Default

The auxiliary port is disabled. `disable` is the default option.

Options

- | | |
|----------------------------------|---|
| disable | Disable the port. |
| insecure | Disable super user access or root logins to establish terminal connection. |
| type <i>terminal-type</i> | <p>Type of terminal that is connected to the port.</p> <ul style="list-style-type: none"> • Range: <code>ansi</code>, <code>vt100</code>, <code>small-xterm</code>, <code>xterm</code> • Default: The terminal type is unknown, and the user is prompted for the terminal type. |
| port-type | <p>Set the RJ-45 console port or the Mini-USB console port as the active console port.</p> <ul style="list-style-type: none"> • <code>mini-usb</code>—Set the Mini USB type-B console port as the active console port. • <code>rj45</code>—Set the RJ-45 console port as the active console port. • Default: The RJ-45 console port is the active port. |

Required Privilege Level

`system`—To view this statement in the configuration.

`system-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

`port-type` introduced in Junos OS Release 11.3 for EX Series switches.

RELATED DOCUMENTATION

[Configuring Junos OS to Set Console and Auxiliary Port Properties](#)

[Configuring Console and Auxiliary Port Properties](#)

[Configuring the Console Port Type \(CLI Procedure\)](#)

authentication-key-chains (TCP-AO)

IN THIS SECTION

- [Syntax | 101](#)
- [Hierarchy Level | 102](#)
- [Description | 102](#)
- [Required Privilege Level | 103](#)
- [Release Information | 103](#)

Syntax

```
authentication-key-chains {  
  key-chain keychain-name {  
    key key ID {  
      secret secretpassword;  
      start-time yyyy-mm-dd.hh:mm:ss;  
      algorithm ao;  
      ao-attribute {  
        send-id send ID;  
        rcv-id receive ID;  
        cryptographic-algorithm (aes-128-cmac-96 | hmac-sha-1-96);  
        tcp-ao-option (disabled | enabled);  
      }  
    }  
  }  
}
```

Hierarchy Level

[edit security]

Description

Configure authentication keychains for TCP Authentication Option (TCP-AO) .

Table 3: Options

Option	Description
key-chain	Enter a unique name for the keychain. For example, new-auth-key.
key	Enter a unique key ID for each key. In a key-chain, keys are numbered sequentially, from key0 through key63.
secret	Enter a unique secret key or password for each key. Use any alphanumeric characters without any space. Once configured, it will appear in an encrypted format.
start-time	Enter a time in <i>YYYY-MM-DD.HH:MM</i> format to specify the time when the control gets passed on from one key to the next. When a configured start time arrives (based on the device's clock), the key with that start time becomes active.
algorithm	Enter ao to indicate the authentication option.
send-id	Enter any two numbers between 0 and 255. You can also use the same number as the send-id and the rcv-id of the same key. You must not use this numbers for any other key inside that key chain. You can configure up to 64 keys in a key-chain.
rcv-id	Enter any two numbers between 0 and 255. You must not use this numbers for any other key inside that key chain. Reverse the send-id and rcv-id values when you configure the same key in the receiving device.

Table 3: Options *(Continued)*

Option	Description
cryptographic-algorithm	Choose a cryptographic algorithm. Starting in Junos OS Release 20.3R1, to be compliant to <i>RFC5925, The TCP Authentication Option</i> , we are supporting HMAC-SHA1 and AES-128 algorithms.
tcp-ao-option	Choose enable to include the TCP-AO option. NOTE: The default value is disabled.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 20.3R1.

RELATED DOCUMENTATION

[TCP Authentication Option \(TCP-AO\) | 45](#)

console (System Ports)

IN THIS SECTION

[Syntax | 104](#)

- Hierarchy Level | 104
- Description | 104
- Default | 104
- Options | 105
- Required Privilege Level | 105
- Release Information | 106

Syntax

```
console {  
    authentication-order [authentication-methods];  
    disable;  
    insecure;  
    log-out-on-disconnect;  
    type terminal-type;  
}
```

Hierarchy Level

```
[edit system ports]
```

Description

Configure the characteristics of the console port.

Default

The console port is enabled and its speed is 9600 baud.

Options

authentication-order Specify the order in which the authentication methods such as **password** (for local password authentication), **radius** (for RADIUS server authentication), or **tacplus** (for TACACS+ server authentication) should be attempted.

authentication-methods—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:

- **password**—Use the password configured for the user with the authentication statement at the [edit system login user] hierarchy level.
- **radius**—Use RADIUS authentication services.
- **tacplus**—Use TACACS+ authentication services.

disable—Disable console login connections.

insecure—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode. This option can be used to prevent a user from attempting password recovery by booting into single-user mode, if the user does not know the root password.

log-out-on-disconnect—Log out the session when the data carrier on the console port is lost.

NOTE:

- The **log-out-on-disconnect** option is not operational on MX80 routers. On MX80 routers you must manually log out from the console with the `request system logout u0` command.
- The **log-out-on-disconnect** option is not operational on guest network functions (GNFs), which are managed using Juniper Device Manager (JDM). You must use the `exit` command to log out from the GNF console. For more information, see [Junos Node Slicing user guide](#).

type terminal-type—Type of terminal that is connected to the port.

- **Range:** `ansi`, `vt100`, `small-xterm`, `xterm`
- **Default:** The terminal type is unknown, and the user is prompted for the terminal type.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

disable option added in Junos OS Release 7.6.

authentication-order option added in Junos OS Release 12.2R3.

RELATED DOCUMENTATION

[Configuring Junos OS to Set Console and Auxiliary Port Properties](#)

[Configuring Console and Auxiliary Port Properties](#)

[Determine the Authentication Order for LDAPS, RADIUS, TACACS+, and Password Authentication](#)

default-address-selection

IN THIS SECTION

- [Syntax | 106](#)
- [Hierarchy Level | 107](#)
- [Description | 107](#)
- [Required Privilege Level | 108](#)
- [Release Information | 108](#)

Syntax

```
default-address-selection;
```

Hierarchy Level

[edit system]

Description

Use the loopback interface, `lo0`, as the source address for all locally generated IP packets when the packet is sent through a routed interface, and also when the packet is sent through a local interface such as `fxp0`. The `lo0` interface is the interface to the router's or switch's Routing Engine.

The default address is used as the source address for all locally generated IP packets on outgoing interfaces that are unnumbered. If an outgoing interface is numbered, the default address is chosen using the following sequence:

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface `lo0` that is not `127.0.0.1` is used.
2. The primary address on the primary interface is used.
3. When there are multiple interfaces with "primary" and "preferred" addresses, the interface with the lowest interface index is selected, and the primary address is used. In the case that none of the interface's IP addresses are explicitly marked with the `primary` statement, the numerically lowest address on that interface is used as the system default address.
4. Any remaining interface with an IP address may be selected. This includes the router's management or internal interfaces. For this reason, it's recommended that you assign a loopback address, or explicitly configure a primary interface, to control default address selection.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out through the interface. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

An interface's *preferred address* is the default local address used for packets sourced by the local router or switch to destinations on the subnet. By default, the numerically lowest local address configured for the interface is chosen as the preferred address on the subnet.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[TCP | 34](#)

Configure Default, Primary, and Preferred Addresses and Interfaces

diag-port-authentication

IN THIS SECTION

- [Syntax | 109](#)
- [Hierarchy Level | 109](#)
- [Description | 109](#)
- [Default | 109](#)
- [Options | 109](#)
- [Required Privilege Level | 110](#)
- [Release Information | 110](#)

Syntax

```
diag-port-authentication (encrypted-password "password" | plain-text-password);
```

Hierarchy Level

```
[edit system]
```

Description

Configure a password for performing diagnostics on the router's System Control Board (SCB), System and Switch Board (SSB), Switching and Forwarding Module (SFM), or Forwarding Engine Board (FEB) port.

For routers that have more than one SSB, the same password is used for both SSBs.

NOTE: Do not run diagnostics on the SCB, SSB, SFM, or FEB unless you have been instructed to do so by Customer Support personnel.

Default

No password is configured on the diagnostics port.

Options

`encrypted-password password`—Use MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.

You cannot configure a blank password for encrypted-password using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

`plain-text-password`—Use a plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password for each user.

Required Privilege Level

`system`—To view this statement in the configuration.

`system-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Diagnostic Tools Overview](#)

[show interfaces diagnostics optics](#)

extended-statistics

IN THIS SECTION

- [Syntax | 111](#)
- [Hierarchy Level | 111](#)
- [Description | 111](#)
- [Default | 111](#)
- [Required Privilege Level | 111](#)

Syntax

```
extended-statistics;
```

Hierarchy Level

```
[edit chassis]
```

Description

(MX Series routers only) Enable accounting of system statistics for IPv4 and IPv6 traffic.

Default

Accounting of system statistics is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3.

RELATED DOCUMENTATION

| [show system statistics](#)

icmp (System)

IN THIS SECTION

- [Syntax | 112](#)
- [Hierarchy Level | 113](#)
- [Description | 113](#)
- [Options | 113](#)
- [Required Privilege Level | 113](#)
- [Release Information | 113](#)

Syntax

```
icmp {  
  ttl-expired-source-address ttl-expired-source-address;  
}
```

Hierarchy Level

[edit system]

Description

Configure Internet Control Message Protocol (ICMP) features at the system level.

Options

ttl-expired-source-address

Use an IPv4 address as the source address for ICMP time-to-live (TTL) expiry error messages.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.4R1.

RELATED DOCUMENTATION

| [ICMP Features](#) | 20

icmp (Error Message Rate Limit)

IN THIS SECTION

- [Syntax | 114](#)
- [Hierarchy Level | 114](#)
- [Description | 114](#)
- [Options | 115](#)
- [Required Privilege Level | 115](#)
- [Release Information | 115](#)

Syntax

```
icmp {  
    rate-limit rate-limit;  
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the rate at which ICMP messages are generated for IPv4 packet errors for non-ttl-expired packets. 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.

NOTE: This statement does not apply to ICMP messages for ttl-expired packets; for these errors, the rate is fixed at 500 pps.

Options

`rate-limit` *rate-limit*—Rate in packets per second (pps). Starting in Junos OS Release 18.4R1, the maximum rate is 1000 pps. In earlier releases, the maximum rate is 50 pps.

- **Range:** 1 through 1000
- **Default:** 1

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Rate Limit ICMP Error Messages](#) | 25

icmp6 (Error Message Rate Limit)

IN THIS SECTION

- [Syntax | 116](#)
- [Hierarchy Level | 116](#)
- [Description | 116](#)
- [Options | 117](#)
- [Required Privilege Level | 117](#)
- [Release Information | 117](#)

Syntax

```
icmp6 {  
    rate-limit rate-limit;  
}
```

Hierarchy Level

```
[edit chassis]
```

Description

Configure the rate at which ICMP messages are generated for IPv6 packet errors for non-ttl-expired packets. 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.

NOTE: This statement does not apply to ICMP messages for ttl-expired packets; for these errors, the rate is fixed at 500 pps.

Options

`rate-limit` *rate-limit*—Rate in packets per second (pps). Starting in Junos OS Release 18.4R1, the maximum rate is 1000 pps. In earlier releases, the maximum rate is 50 pps.

- **Range:** 1 through 1000
- **Default:** 1

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Rate Limit ICMP Error Messages](#) | 25

internet-options

IN THIS SECTION

- [Syntax | 118](#)
- [Hierarchy Level | 119](#)
- [Description | 119](#)
- [Options | 119](#)
- [Required Privilege Level | 121](#)
- [Release Information | 122](#)

Syntax

```
internet-options {
    (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
    icmpv4-rate-limit bucket-size <bucket-size seconds> <packet-rate packet-rate>;
    icmpv6-rate-limit bucket-size <bucket-size seconds> <packet-rate packet-rate>;
    (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
    ipv6-duplicate-addr-detection-transmits ipv6-duplicate-addr-detection-transmits;
    (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);
    (ipv6-reject-zero-hop-limit | no-ipv6-reject-zero-hop-limit);
    ipv6-path-mtu-discovery-timeout minutes;
    no-tcp-reset (drop-all-tcp | drop-tcp-with-syn-only);
    no-tcp-rfc1323;
    no-tcp-rfc1323-paws;
    (path-mtu-discovery | no-path-mtu-discovery);
    source-port {
        upper-limit upper-limit;
    }
    (source-quench | no-source-quench);
    tcp-drop-synfin-set;
    tcp-mss mss-value;
}
```

Hierarchy Level

[edit system]

Description

Configure system IP options to protect against certain types of DoS attacks.

Options

gre-path-mtu-discovery

(ACX Series, EX Series, Junos Fusion, M Series, MX Series, OCX Series, PTX Series, QFX Series, SRX Series, T Series) Configure path MTU discovery for outgoing GRE tunnel connections. By default, path MTU discovery is enabled.

- no-gre-path-mtu-discovery—Path MTU discovery is disabled.

icmpv4-rate-limit

Configure rate-limiting parameters for ICMPv4 messages sent.

- Values:
 - bucket-size *seconds*—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds. Default: 5.
 - packet-rate *pps*—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps. Default: 1000.

icmpv6-rate-limit

(ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, SRX Series) Configure rate-limiting parameters for ICMPv6 messages sent.

- Values:
 - bucket-size *seconds*—Number of seconds in the rate-limiting bucket. Range: 0 through 4294967295 seconds. Default: 5.
 - packet-rate *pps*—Rate-limiting packets earned per second. Range: 0 through 4294967295 pps. Default: 1000.

ipip-path-mtu-discovery	<p>(ACX Series, EX Series, Junos Fusion, M Series, MX Series, OCX Series, PTX Series, QFX Series, SRX Series, T Series) Configure path MTU discovery for outgoing IP-IP tunnel connections. By default, path MTU discovery is enabled.</p> <ul style="list-style-type: none"> • no-ipip-path-mtu-discovery—Path MTU discovery is disabled.
ipv6-duplicate-addr-detection-transmits	<p>Control the number of attempts for IPv6 duplicate address detection.</p> <ul style="list-style-type: none"> • Range: 0 to 20 • Default: 3
ipv6-path-mtu-discovery	<p>(ACX Series, EX Series, Junos Fusion, M Series, MX Series, OCX Series, PTX Series, QFX Series, SRX Series, T Series) Configure path MTU discovery for IPv6 packets. By default, IPv6 path MTU discovery is enabled.</p> <ul style="list-style-type: none"> • no-ipv6-path-mtu-discovery—IPv6 path MTU discovery is disabled.
ipv6-path-mtu-discovery-timeout	<p>(ACX Series, EX Series, Junos Fusion, M Series, MX Series, OCX Series, PTX Series, QFX Series, SRX Series, T Series) Set the IPv6 path MTU discovery time-out interval.</p> <ul style="list-style-type: none"> • Values: <i>minutes</i>—IPv6 path MTU discovery timeout. • Default: 10 minutes.
ipv6-reject-zero-hop-limit	<p>Reject incoming IPv6 packets with a zero hop-limit value in their header. This is enabled by default.</p> <ul style="list-style-type: none"> • no-ipv6-reject-zero-hop-limit—Allow incoming IPv6 packets with a zero hop-limit value in their header.
no-tcp-reset	<p>Do not send an RST TCP packet (a packet with the reset flag set) in response to a TCP packet received on a non-listening port.</p> <p>By default, when a TCP packet is received on a non-listening port, a device sends a TCP packet with the RST flag set and drops the connection. This might lead to a security risk. Configuring this statement prevents the sending of RST TCP packets to non-listening ports.</p> <p>You must configure this statement with one of two options:</p> <ul style="list-style-type: none"> • drop-all-tcp—When a TCP segment is received on a closed port, the device drops the packet and does not send back a RST segment. This helps to protect against stealth port scans.

- **drop-tcp-with-syn-only**—When a TCP packet with a SYN bit is received on a non-listening port, the device drops the packet and does not send back a RST segment, which makes the device appear as a null route. For all other TCP packets, the device sends back a RST segment and does not drop the packet.

no-tcp-rfc1323 Configure the Junos OS to disable RFC 1323 TCP extensions.

no-tcp-rfc1323-paws Configure the Junos OS to disable the RFC 1323 Protection Against Wrapped Sequence (PAWS) number extension.

path-mtu-discovery Configure path MTU discovery for outgoing Transmission Control Protocol (TCP) connections. By default, path MTU discovery is enabled.

- **no-path-mtu-discovery**—Path MTU discovery is disabled.

source-port Configure the range of port addresses.

- **Values:**
 - **upper-limit** *upper-limit*—(Optional) The range of port addresses can be a value from 5000 through 65,355.

source-quench Configure how the Junos OS handles Internet Control Message Protocol (ICMP) source quench messages. By default, the Junos OS reacts to ICMP source quench messages.

- **no-source-quench**—Do not react to incoming ICMP source quench messages.

tcp-drop-synfin-set Configure the device to drop packets that have both the SYN and FIN bits set.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

no-tcp-reset introduced in Junos OS Release 9.4.

no-tcp-reset introduced in Junos OS Release 11.1 for SRX Series and vSRX devices.

icmpv4-rate-limit and source-port introduced in Junos OS Release 11.1 for the QFX Series and Junos OS Release 14.1X53-D20 for the OCX Series.

RELATED DOCUMENTATION

[ICMP Features | 20](#)

[IPv6 Features | 29](#)

[Path MTU Discovery | 32](#)

[TCP | 34](#)

[Configuring Junos OS to Extend the Default Port Address Range | 78](#)

[Understanding Traffic Processing on Security Devices](#)

nd-maxucast-retry

IN THIS SECTION

- [Syntax | 123](#)
- [Hierarchy Level | 123](#)
- [Description | 123](#)
- [Default | 123](#)
- [Options | 123](#)
- [Required Privilege Level | 123](#)
- [Release Information | 124](#)

Syntax

```
nd-maxucast-retry count
```

Hierarchy Level

```
[edit system]
```

Description

Configures how many times the device attempts to reach the IPv6 neighbor entries.

Default

The default number of retry attempts is 3.

Options

count Indicates the maximum number of retries for unicast neighbor solicitation.

Required Privilege Level

admin-control

Release Information

Statement introduced in Junos OS Evolved Release 19.4R1.

RELATED DOCUMENTATION

[Example: Configuring ARP Cache Protection | 10](#)
[arp | 88](#)

no-multicast-echo

IN THIS SECTION

- [Syntax | 124](#)
- [Hierarchy Level | 124](#)
- [Description | 125](#)
- [Default | 125](#)
- [Required Privilege Level | 125](#)
- [Release Information | 125](#)

Syntax

```
no-multicast-echo
```

Hierarchy Level

```
[edit system]
```

Description

Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

Default

The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

| [ICMP Features](#) | [20](#)

non-subscriber-no-reply

IN THIS SECTION

- [Syntax](#) | [126](#)
- [Hierarchy Level](#) | [126](#)
- [Description](#) | [126](#)

- Options | 126
- Required Privilege Level | 126
- Release Information | 127

Syntax

```
non-subscriber-no-reply;
```

Hierarchy Level

```
[edit system arp]
```

Description

Enable this option to drop ARP requests from non-subscribers when a user route is dynamically added for a subscriber. Configuring this statement suppresses the ARP response from the kernel when there is an ARP request for a loopback interface from static DHCP subscribers using a common LAN segment between two devices. However, this configuration might not be effective if the subscriber configuration has suppressed either a destination Layer 2 route or an access Layer 3 route.

Options

This command has no options.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3R9.

RELATED DOCUMENTATION

| *autoinstallation*

no-ping-record-route

IN THIS SECTION

- [Syntax | 127](#)
- [Hierarchy Level | 128](#)
- [Description | 128](#)
- [Required Privilege Level | 128](#)
- [Release Information | 128](#)

Syntax

```
no-ping-record-route;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the Junos OS to disable the reporting of the IP address in ping responses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

| [ICMP Features](#) | 20

no-ping-time-stamp

IN THIS SECTION

- [Syntax](#) | 129
- [Hierarchy Level](#) | 129

- [Description | 129](#)
- [Required Privilege Level | 129](#)
- [Release Information | 129](#)

Syntax

```
no-ping-time-stamp;
```

Hierarchy Level

```
[edit system]
```

Description

Configure the Junos OS to disable the recording of timestamps in ping responses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

RELATED DOCUMENTATION

| [ICMP Features](#) | [20](#)

path-mtu-discovery (Tunnel)

IN THIS SECTION

- [Syntax](#) | [130](#)
- [Hierarchy Level](#) | [130](#)
- [Description](#) | [130](#)
- [Default](#) | [131](#)
- [Required Privilege Level](#) | [131](#)
- [Release Information](#) | [131](#)

Syntax

```
(path-mtu-discovery | no-path-mtu-discovery);
```

Hierarchy Level

```
[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]  
[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]
```

Description

Configure path MTU discovery for outgoing tunnel connections:

- `path-mtu-discovery`—Path MTU discovery is enabled.
- `no-path-mtu-discovery`—Path MTU discovery is disabled.

NOTE: Starting in Junos OS Release 17.2R1, the `no-path-mtu-discovery` configuration statement in the `[edit interfaces ip-fpc/pic/port unit logical-unit-number tunnel]` and `[edit interfaces gr-fpc/pic/port unit logical-unit-number tunnel]` hierarchies is no longer available for `ipip6` tunnels.

Default

Path MTU discovery is enabled.

Required Privilege Level

`system`—To view this statement in the configuration.

`system-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Path MTU Discovery](#) | 32

proactive-arp-detection

IN THIS SECTION

- [Syntax | 132](#)
- [Hierarchy Level | 132](#)
- [Description | 132](#)
- [Default | 133](#)
- [Required Privilege Level | 133](#)
- [Release Information | 133](#)

Syntax

```
proactive-arp-detection;
```

Hierarchy Level

```
[edit system arp]
```

Description

Enable proactive ARP detection to check the reachability of connected devices (within an IP subnet range) on a specified interface. After enabling this statement, you can set the ARP configurations at the interface level by the setting the: `host-discovery address-range`, `ageing-time-out seconds`, and `discovery-time-interval seconds` options at the `[edit interfaces interface-name family inet address ip-address]` hierarchy level. Likewise, you can delete the ARP configuration settings by using the `delete interfaces interface-name unit unit family inet address ip-address host-discovery address-range` command.

Default

This configuration statement is disabled by default.

Required Privilege Level

admin

Release Information

Statement introduced in Junos OS Release 18.3R1-S3, 18.3R2 and 19.3R1.

tcpao-auth-mismatch

IN THIS SECTION

- [Syntax | 133](#)
- [Hierarchy Level | 134](#)
- [Description | 134](#)
- [Options | 134](#)
- [Required Privilege Level | 134](#)
- [Release Information | 134](#)

Syntax

```
tcpao-auth-mismatch {  
    allow-without-tcpao;  
}
```

Hierarchy Level

```
[edit fabric protocols bgp]
[edit logical-systems name protocols bgp]
[edit logical-systems name routing-instances name protocols bgp]
[edit logical-systems name tenants name routing-instances name protocols bgp]
[edit routing-instances name protocols bgp]
[edit tenants name routing-instances name protocols bgp]
[edit protocols bgp group group neighbor neighbor]
[edit protocols bgp]
[edit protocols ldp session session]
```

Description

Continue without TCP-AO if any one TCP endpoint does not have TCP-AO configured.

Options

allow-without-tcpao Allow the connection establishment without TCP-AO.

Required Privilege Level

routing—To view this statement in the configuration.

routing-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

[TCP Authentication Option \(TCP-AO\) | 45](#)

tcp-mss

IN THIS SECTION

- [Syntax | 135](#)
- [Hierarchy Level | 135](#)
- [Description | 136](#)
- [Options | 137](#)
- [Required Privilege Level | 137](#)
- [Release Information | 137](#)

Syntax

```
tcp-mss mss-value;
```

Hierarchy Level

```
[edit system internet-options]  
[edit interfaces name unit number family protocol]
```

Description

Enable and specify the TCP maximum segment size (TCP MSS) to be used to replace that of TCP SYN packets whose maximum segment size (MSS) option is set to a higher value than the value you choose.

If the router receives a TCP packet with the SYN bit and MSS option set and the MSS option specified in the packet is larger than the MSS specified by the `tcp-mss` command, the router replaces the MSS value in the packet with the lower value specified by the `tcp-mss` statement.

There are multiple factors which defines the MSS value for TCP packets in Junos, which are reflected in the MSS value displayed in the output of the `show system connection extensive`.

- The MSS value offered by the peer in the SYN packet
- Rounding the MSS off to the nearest multiple of 2048
- The MTU value of the interface
- The configured path MTU value
- Whether TCP sessions that are not directly-connected, and the path MTU are disabled
- Whether the TCP sessions are on directly-connected network

This statement enables you to specify the MSS size in TCP SYN packets used during session establishment. Decreasing the MSS size helps to limit packet fragmentation and to protect against packet loss that can occur when a packet must be fragmented to meet the MTU size but the packet's DF (don't fragment) bit is set.

Use the `tcp-mss` statement to specify a lower TCP MSS value than the value in the TCP SYN packets.

If you configure this statement under the `[edit interfaces]` hierarchy, keep in mind:

- This statement only takes effect on `lt-` interfaces, `gr-` interfaces, and service-related interfaces like L2TP network server (LNS) and service sets.
- TCP MSS adjustment only takes effect for packets entering the interface. This statement has no effect on packets exiting an interface.

NOTE: When an SRX Series device is running in packet mode with MPLS, TCP MSS is not supported.

Options

mss-value—TCP MSS value for SYN packets with a higher MSS value set.

- **Range:** 64 through 65535 bytes.
- **Default:** TCP MSS is disabled.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[TCP](#) | 34

5

CHAPTER

Operational Commands

`clear arp` | 139

`clear multicast snooping statistics` | 141

`show arp` | 143

`show system statistics arp` | 150

`show system statistics icmp` | 160

`show system statistics icmp6` | 167

`show system statistics igmp` | 177

`show system statistics ip` | 183

`show system statistics ip6` | 195

`show system statistics tcp` | 205

clear arp

IN THIS SECTION

- [Syntax | 139](#)
- [Description | 139](#)
- [Options | 140](#)
- [Required Privilege Level | 140](#)
- [Output Fields | 140](#)
- [Sample Output | 140](#)
- [Release Information | 141](#)

Syntax

```
clear arp
<all>
<hostname hostname>
<interface interface-name>
<logical-system logical-system-name>
<tenant name>
<vpn vpn>
```

Description

Remove entries from the Address Resolution Protocol (ARP) table for the current CLI view. To clear entries for a specific logical system, you must first enter the `set cli logical-system logical-system-name` command, and then issue the `clear arp` command.

Options

all	Clear all entries from the ARP table.
hostname <i>hostname</i>	(Optional) Clear only the specified host entry from the ARP table.
interface <i>interface-name</i>	(Optional) Clear entries only for the specified interface from the ARP table.
logical-system <i>logical-system-name</i>	(Optional) Clear entries for only the specified logical system from the ARP table (only available in main router context).
tenant <i>name</i>	(Optional) Clear entries for only the specified tenant from the ARP table (only available in main router context).
vpn <i>vpn</i>	(Optional) Clear entries from the ARP table for the specified virtual private network (VPN).

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear arp all

```
user@host> clear arp all
192.168.71.254    deleted
192.168.65.46    deleted
192.168.64.10    deleted
```

```
10.0.12.14      deleted
10.0.17.14      deleted
```

clear arp logical-system ls1

```
user@host> clear arp logical-system ls1
192.168.71.254  deleted
192.168.65.46   deleted
192.168.64.10   deleted
10.0.12.14      deleted
10.0.17.14      deleted
```

Release Information

Command introduced before Junos OS Release 7.4.

all option introduced in Junos OS Release 14.2.

tenant option added in Junos OS Release 18.3.

RELATED DOCUMENTATION

[set cli logical-system](#)

[show arp](#) | [143](#)

[show dhcp-security arp inspection statistics](#)

[Port Security Features](#)

clear multicast snooping statistics

IN THIS SECTION

● [Syntax](#) | [142](#)

- [Description | 142](#)
- [Options | 142](#)
- [Required Privilege Level | 143](#)
- [Output Fields | 143](#)
- [Sample Output | 143](#)
- [Release Information | 143](#)

Syntax

```
clear multicast snooping statistics
<instance instance-name>
<interface interface-name>
<logical-system (all | logical-system-name)>
```

Description

Clear IP multicast snooping statistics.

Options

none	Clear multicast snooping statistics for all supported address families on all interfaces.
instance <i>instance-name</i>	(Optional) Clear multicast snooping statistics for the specified instance.
interface <i>interface-name</i>	(Optional) Clear multicast snooping statistics on a specific interface.
logical-system (all <i>logical-system-name</i>)	(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear multicast snooping statistics

```
user@host> clear multicast snooping statistics
```

Release Information

Command introduced in Junos OS Release 8.5.

show arp

IN THIS SECTION

- [Syntax | 144](#)
- [Description | 144](#)
- [Options | 145](#)
- [Required Privilege Level | 146](#)
- [Output Fields | 146](#)
- [Sample Output | 147](#)

Syntax

```
show arp
<expiration-time>
<hostname host-name>
<interface interface-name>
<logical-system logical-system-name>
<no-resolve>
<reference-count count>
<tenant name>
<state state>
<vpn vpn-name>
```

Description

Display all entries in the Address Resolution Protocol (ARP) table. To display entries for a particular logical system only, first enter the `set cli logical-system logical-system-name` command, and then enter the `show arp` command.

NOTE: Starting with Junos OS Release 14.2, the following enhancements have been made to the output of the `show arp interfaces` command:

- For integrated routing and bridging (IRB) interfaces, in the output of the `show arp` command, the IRB interface name is displayed under the Interface field of the output and the Layer 2 interface identifier is specified in square brackets following the IRB name. Until Release 14.1 and earlier, only the layer 2 interface name and not the IRB name was displayed.
- Starting with release 14.2, if you do not specify a subinterface or a logical unit of the interface with the `show arp interface interface-name` command, an error message is shown. Until Release 14.1 and earlier, if you did not specify the subinterface for a physical interface, the system

considered the supplied command to be for subinterface 0 and displayed the output. For example, if you entered `ge-2/2/5`, it was processed by the system as `ge-2/2/5.0`.

- When IRB interfaces are configured and if you attempt to specify an interface name that is not configured on the system, an error message is displayed stating the particular interface is not defined on the system. Until release 14.1 and earlier, unrelated and incorrect entries were displayed even for interface names that did not exist.
- Starting with Release 14.2, you can enter the `show arp interface` command with the IRB name and retrieve the statistical details for the IRB interface. This functionality was not available previously. However, you could previously obtain the ARP details of an IRB interface that had a Layer 2 interface configured.

NOTE: Starting with Junos OS Release 16.1, `show arp no-resolve` command does not display the underlying `ifl` information if `enhanced-convergence` statement at `[edit irb unit unit-number hierarchy level` and `enhanced-ip` statement at `[edit chassis network-services]` hierarchy level is configured for the destination interface IRB.

Options

none	Display the entries in the ARP table.
expiration-time	(Optional) Display the amount of time, in seconds, until each ARP entry is set to expire.
hostname <i>host-name</i>	(Optional) Display the hostname.
interface <i>interface-name</i>	(Optional) Display information about ARP for the specified logical interface
logical-system <i>logical-system-name</i>	(Optional) Display ARP entries for the specified logical system; only available on the main router context.
no-resolve	(Optional) Do not attempt to determine the hostname that corresponds to the IP address.
reference-count <i>count</i>	(Optional) Display the next-hop reference count.
state <i>state</i>	(Optional) Display the next-hop current state.

- tenant *name*** (Optional) Display the ARP entries for the specified tenant. Only available in the main router context.
- vpn *vpn-name*** (Optional) Display entries in the ARP table for the specified virtual private network's (VPN) routing table.

NOTE:

- The `vpn` option displays only entries for non-subscriber interfaces. You must omit this option if you want to view ARP entries for subscriber interfaces.
- The `show arp vpn vpn-name no-resolve` command does not display entries for subscriber interfaces in the non-default routing-instance, that is when the subscriber interface is coming up in the non-default routing-instance.

Required Privilege Level

view

Output Fields

[Table 4 on page 146](#) describes the output fields for the `show arp` command. Output fields are listed in the approximate order in which they appear.

Table 4: show arp Output Fields

Field Name	Field Description
MAC Address	Media access control (MAC) address that corresponds to the IP address.
Address	IP address that corresponds to the hostname.
Name	Hostname.

Table 4: show arp Output Fields (Continued)

Field Name	Field Description
Interface	Interface name.
Flags	<p>(no-resolve option only) Indicates how mappings between IP and MAC addresses are defined:</p> <ul style="list-style-type: none"> • Permanent—Static mapping. • Permanent and published—Static mapping that is published. • None—Dynamic mapping.
TTE	(expiration-time option only) Amount of time, in seconds, until ARP entry is set to expire.

Sample Output

show arp

```
user@host> show arp
MAC Address      Address          Name              Interface
00:e0:81:22:fd:74 192.168.64.10    firewall.my.net    fxp0.0
00:04:5a:65:78:e1 192.168.65.13    lab.my net         fxp0.0
```

show arp no-resolve

```
user@host> show arp no-resolve
MAC Address      Address          Interface         Flags
00:90:69:96:00:01 10.10.45.5       fe-0/0/1.0       none
00:00:00:00:00:01 200.200.200.1    fe-0/0/0.0       permanent published
00:00:00:00:00:02 200.200.200.2    fe-0/0/0.0       permanent
00:90:69:91:b0:00 200.200.200.3    fe-0/0/0.0       none
Total entries: 4
```

show arp no-resolve

```
user@host> show arp no-resolve (QFX Series)
```

The command is displaying the MAC and MAC-IP routes are programmed over the .local interface.

MAC Address	Address	Interface	Flags
00:00:10:00:11:00	10.0.1.1	irb.100 [.local..5]	permanent remote
00:00:10:00:22:00	10.0.1.2	irb.100 [.local..5]	permanent remote
00:00:10:00:33:00	10.0.1.3	irb.100 [.local..5]	permanent remote
00:00:10:00:44:00	10.0.1.4	irb.100 [.local..5]	permanent remote
00:00:10:11:22:00	10.0.1.12	irb.100 [.local..5]	permanent remote
00:00:10:11:33:00	10.0.1.13	irb.100 [.local..5]	permanent remote
00:00:10:11:44:00	10.0.1.14	irb.100 [.local..5]	permanent remote
00:00:10:11:66:00	10.0.1.16	irb.100 [.local..5]	permanent remote
40:00:10:11:11:00	10.0.1.254	irb.100	permanent published gateway
00:00:10:00:11:01	10.0.2.1	irb.200 [.local..5]	permanent remote
00:00:10:00:22:01	10.0.2.2	irb.200 [.local..5]	permanent remote
00:00:10:00:33:01	10.0.2.3	irb.200 [.local..5]	permanent remote
00:00:10:00:44:01	10.0.2.4	irb.200 [.local..5]	permanent remote
00:00:10:11:22:01	10.0.2.12	irb.200 [.local..5]	permanent remote
00:00:10:11:33:01	10.0.2.13	irb.200 [.local..5]	permanent remote
00:00:10:11:44:01	10.0.2.14	irb.200 [.local..5]	permanent remote
00:00:10:11:66:01	10.0.2.16	irb.200 [.local..5]	permanent remote
Total entries: 26			

command-name

```
user@host> show arp no-resolve
```

The command displaying the underlying l2 ifl information when enhanced-convergence statement and enhanced-ip statement is not configured.

```
show arp no-resolve
```

MAC Address	Address	Interface	Flags
02:01:00:00:00:05	10.0.0.5	em1.0	none
00:00:5e:00:01:1b	91.91.91.50	irb.0[xe-2/1/0.0]	none >>> underlying l2 ifl associated
02:01:00:00:00:05	128.0.0.5	em1.0	none
02:01:00:00:00:05	128.0.0.6	em1.0	none

```
02:00:00:00:00:12 128.0.0.18      em0.0      none
00:26:88:6a:c6:80 192.168.237.126 fxp0.0     none
Total entries: 6
```

The command not displaying the underlying l2 ifl information when enhanced-convergence statement and enhanced-ip statement is configured.

```
MAC Address      Address      Interface    Flags
02:01:00:00:00:05 10.0.0.5     em1.0        none
00:00:5e:00:01:1b 91.91.91.50  irb.0        none >>> underlying l2 ifl
association is removed.
02:01:00:00:00:05 128.0.0.5    em1.0        none
02:01:00:00:00:05 128.0.0.6    em1.0        none
02:00:00:00:00:12 128.0.0.18   em0.0        none
00:26:88:6a:c6:80 192.168.237.126 fxp0.0     none
Total entries: 6
```

show arp expiration-time

```
user@host> show arp expiration-time
MAC Address      Address      Name          Interface    Flags TTE
00:a0:a5:12:3e:d4 10.0.0.5     10.0.0.5      fxp1.0     none
00:e0:81:22:fd:74 192.168.64.10 supernova.englab.juniper. fxp0.0 none 1491
00:30:48:84:03:56 192.168.65.46 kgb.englab.juniper.net   fxp0.0 none 1279
00:03:ba:12:f7:5e 192.168.65.226 nmssun1-eri0.englab.junip fxp0.0 none 452
00:90:69:8e:b0:fc 192.168.71.254 stonewall-ge-200.englab.j fxp0.0 none 1421
Total entries: 5
```

Release Information

Command introduced before Junos OS Release 7.4.

expiration-time option added in Junos OS Release 8.1.

logical-system and vpn options added in Junos OS Release 10.1.

reference-count, tenant, and state options added in Junos OS Release 18.3.

RELATED DOCUMENTATION

[*clear arp*](#)

[set cli logical-system](#)

show system statistics arp

IN THIS SECTION

- [Syntax | 150](#)
- [Syntax \(EX Series Switches\) | 150](#)
- [Syntax \(TX Matrix Router\) | 151](#)
- [Syntax \(TX Matrix Plus Router\) | 151](#)
- [Description | 151](#)
- [Options | 151](#)
- [Additional Information | 152](#)
- [Required Privilege Level | 152](#)
- [Sample Output | 153](#)
- [Release Information | 159](#)

Syntax

```
show system statistics arp
```

Syntax (EX Series Switches)

```
show system statistics arp  
<all-members>
```

```
<local>  
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics arp  
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics arp  
<all-chassis | all-lcc | lcc number / sfc number>
```

Description

Display system-wide Address Resolution Protocol (ARP) statistics.

Options

none	Display system-wide ARP statistics.
all-chassis	(Optional) Display ARP statistics for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system-wide ARP statistics for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system-wide ARP statistics for all routers connected to the TX Matrix Plus router
all-members	(Optional) Display ARP statistics for all members of the Virtual Chassis configuration.

lcc *number* (Optional) On a TX Matrix router, display ARP statistics for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display ARP statistics for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local (Optional) Display ARP statistics for the local Virtual Chassis member.

member *member-id* (Optional) Display ARP statistics for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value from 0 through 9.

scc (Optional) Display ARP statistics for the TX Matrix router (or switch-card chassis).

sfc *number* (Optional) Display ARP statistics for the TX Matrix Plus router. Replace *number* with 0.

Additional Information

By default, when you issue the `show system statistics arp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics arp

```
user@host> show system statistics arp
arp:
    184710 datagrams received
    2886 ARP requests received
    684 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 unrestricted proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast target address
    0 datagrams with my own hardware address
    0 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    0 datagrams with source address duplicate to mine
    181140 datagrams which were not for me
    0 packets discarded waiting for resolution
    4 packets sent after waiting for resolution
    703 ARP requests sent
    2886 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor
```

show system statistics arp (EX Series Switches)

```
user@host> show system statistics arp
arp:
    186423 datagrams received
    88 ARP requests received
    88 ARP replies received
    0 resolution request received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 proxy requests not proxied
    0 restricted proxy requests not proxied
    0 datagrams with bogus interface
    0 datagrams with incorrect length
    0 datagrams for non-IP protocol
    0 datagrams with unsupported op code
    0 datagrams with bad protocol address length
    0 datagrams with bad hardware address length
    0 datagrams with multicast source address
    0 datagrams with multicast source address
    0 datagrams with my own hardware address
    164 datagrams for an address not on the interface
    0 datagrams with a broadcast source address
    0 datagrams with source address duplicate to mine
    186075 datagrams which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    50 ARP requests sent
    88 ARP replies sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor
```

show system statistics arp (TX Matrix Plus Router)

```

user@host> show system statistics arp
sfc0-re0:
-----
arp:
    487 datagrams received
    8 ARP requests received
    438 ARP replys received
    438 resolution requests received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 proxy requestss not proxied
    0 restricted-proxy requestss not proxied
    0 with bogus interface
    0 with incorrect length
    0 for non-IP protocol
    0 with unsupported op code
    0 with bad protocol address length
    0 with bad hardware address length
    0 with multicast source address
    0 with multicast target address
    0 with my own hardware address
    0 for an address not on the interface
    0 with a broadcast source address
    0 with source address duplicate to mine
    41 which were not for me
    0 packets discarded waiting for resolution
    438 packets sent after waiting for resolution
    1282 ARP requests sent
    8 ARP replys sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

```

lcc0-re0:

arp:

```

    19 datagrams received
    0 ARP requests received
    1 ARP reply received
    0 resolution requests received
    0 unrestricted proxy requests
    0 restricted proxy requests
    0 received proxy requests
    0 proxy requestss not proxied
    0 restricted-proxy requestss not proxied
    0 with bogus interface
    0 with incorrect length
    0 for non-IP protocol
    0 with unsupported op code
    0 with bad protocol address length
    0 with bad hardware address length
    0 with multicast source address
    0 with multicast target address
    0 with my own hardware address
    0 for an address not on the interface
    0 with a broadcast source address
    0 with source address duplicate to mine
    18 which were not for me
    0 packets discarded waiting for resolution
    0 packets sent after waiting for resolution
    8 ARP requests sent
    0 ARP replys sent
    0 requests for memory denied
    0 requests dropped on entry
    0 requests dropped during retry
    0 requests dropped due to interface deletion
    0 requests on unnumbered interfaces
    0 new requests on unnumbered interfaces
    0 replies for from unnumbered interfaces
    0 requests on unnumbered interface with non-subnetted donor
    0 replies from unnumbered interface with non-subnetted donor

```

lcc1-re0:

arp:

```

    17 datagrams received

```

```

0 ARP requests received
1 ARP reply received
0 resolution requests received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requestss not proxied
0 restricted-proxy requestss not proxied
0 with bogus interface
0 with incorrect length
0 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
0 with multicast target address
0 with my own hardware address
0 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
16 which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
9 ARP requests sent
0 ARP replys sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

```
lcc2-re0:
```

```
-----
arp:
```

```

18 datagrams received
1 ARP request received
1 ARP reply received
0 resolution requests received
0 unrestricted proxy requests

```

```

0 restricted proxy requests
0 received proxy requests
0 proxy requestss not proxied
0 restricted-proxy requestss not proxied
0 with bogus interface
0 with incorrect length
0 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
0 with multicast target address
0 with my own hardware address
0 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
16 which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
9 ARP requests sent
1 ARP reply sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

lcc3-re0:

arp:

```

13 datagrams received
0 ARP requests received
1 ARP reply received
0 resolution requests received
0 unrestricted proxy requests
0 restricted proxy requests
0 received proxy requests
0 proxy requestss not proxied
0 restricted-proxy requestss not proxied

```

```

0 with bogus interface
0 with incorrect length
0 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
0 with multicast target address
0 with my own hardware address
0 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
12 which were not for me
0 packets discarded waiting for resolution
0 packets sent after waiting for resolution
8 ARP requests sent
0 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

[Example: Configuring Proxy ARP on an EX Series Switch](#)

[Verifying That Proxy ARP Is Working Correctly](#)

show system statistics icmp

IN THIS SECTION

- [Syntax | 160](#)
- [Syntax \(EX Series Switches\) | 160](#)
- [Syntax \(TX Matrix Router\) | 161](#)
- [Syntax \(TX Matrix Plus Router\) | 161](#)
- [Description | 161](#)
- [Options | 161](#)
- [Additional Information | 162](#)
- [Required Privilege Level | 162](#)
- [Sample Output | 163](#)
- [Release Information | 167](#)

Syntax

```
show system statistics icmp
```

Syntax (EX Series Switches)

```
show system statistics icmp  
<all-members>  
<local>  
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics icmp
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics icmp
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Internet Control Message Protocol (ICMP) statistics.

Options

none	Display system statistics for ICMP.
all-chassis	(Optional) Display system statistics for ICMP for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system statistics for ICMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMP for all connected T1600 or T4000 LCCs.
all-members	(Optional) Display ICMP statistics for all members of the Virtual Chassis configuration.
lcc <i>number</i>	(Optional) On a TX Matrix router, display system statistics for ICMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local	(Optional) Display ICMP statistics for the local Virtual Chassis member.
member <i>member-id</i>	(Optional) Display ICMP statistics for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value from 0 through 9.
scc	(Optional) Display system statistics for ICMP for the TX Matrix router (or switch-card chassis).
sfc number	(Optional) Display system statistics for ICMP for the TX Matrix Plus router. Replace <i>number</i> with 0.

Additional Information

By default, when you issue the `show system statistics icmp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics icmp

```

user@host> show system statistics icmp
icmp:
    16783 drops due to rate limit
    9998 calls to icmp_error
    0 errors not generated because old message was icmp
Output Histogram
    38877 echo reply
    1 destination unreachable
    1 routing redirect
    163 echo
    5000 time exceeded
    4996 parameter problem
    5000 time stamp reply
    0 messages with bad code fields
    0 messages less than the minimum length
    0 messages with bad checksum
    0 messages with bad source address
    20000 messages with bad length
    0 echo drops with broadcast or multicast destination address
    0 timestamp drops with broadcast or multicast destination address
Input Histogram
    5093 echo reply
    5000 destination unreachable
    5000 source quench
    5000 routing redirect
    5000 alternate host address
    38877 echo
    5000 router advertisement
    5000 router solicitation
    5000 time exceeded
    5000 parameter problem
    5000 time stamp
    5000 time stamp reply
    5000 information request
    5000 information request reply
    5000 address mask request
    5000 address mask reply

```

```

5000 traceroute
5000 data conversion error
5000 mobile host redirect
5000 IPv6 where-are-you
5000 IPv6 i-am-here
5000 mobile registration request
5000 mobile registration reply
5000 skip
5000 photuris
43877 message responses generated

```

show system statistics icmp (EX Series Switches)

```

user@host> show system statistics icmp
icmp:
    0 drops due to rate limit
    12 calls to icmp_error
    0 errors not generated because old message was icmp
Output histogram:
    297 echo reply
    12 destination unreachable
    0 messages with bad code fields
    0 messages less than the minimum length
    0 messages with bad checksum
    0 messages with bad source address
    0 messages with bad length
    0 echo drops with broadcast or multicast destination address
    0 timestamp drops with broadcast or multicast destination address
Input histogram:
    297 echo
    297 message responses generated

```

show system statistics icmp (TX Matrix Plus Router)

```

user@host> show system statistics icmp
sfc0-re0:
-----
icmp:
    0 drops due to rate limit
    0 calls to icmp_error

```

```

0 errors not generated because old message was icmp
Output histogram:
    echo reply: 21
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 21
21 message responses generated

```

lcc0-re0:

```

icmp:
    0 drops due to rate limit
    1 call to icmp_error
    0 errors not generated because old message was icmp
Output histogram:
    echo reply: 24
    destination unreachable: 1
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 24
24 message responses generated

```

lcc1-re0:

```

icmp:
    0 drops due to rate limit
    0 calls to icmp_error
    0 errors not generated because old message was icmp
Output histogram:
    echo reply: 23
0 messages with bad code fields

```

```

0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 23
23 message responses generated

```

lcc2-re0:

```

icmp:
    0 drops due to rate limit
    0 calls to icmp_error
    0 errors not generated because old message was icmp
Output histogram:
    echo reply: 22
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address
0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 22
22 message responses generated

```

lcc3-re0:

```

icmp:
    0 drops due to rate limit
    0 calls to icmp_error
    0 errors not generated because old message was icmp
Output histogram:
    echo reply: 22
0 messages with bad code fields
0 messages less than the minimum length
0 messages with bad checksum
0 messages with bad source address
0 messages with bad length
0 echo drops with broadcast or multicast destination address

```

```

0 timestamp drops with broadcast or multicast destination address
Input histogram:
    echo: 22
22 message responses generated

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

show system statistics icmp6

IN THIS SECTION

- [Syntax \(EX Series Switches\) | 168](#)
- [Syntax \(MX Series Routers, SRX Series\) | 168](#)
- [Syntax \(TX Matrix Router\) | 168](#)
- [Syntax \(TX Matrix Plus Router\) | 168](#)
- [Description | 168](#)
- [Options | 169](#)
- [Additional Information | 170](#)
- [Required Privilege Level | 170](#)
- [Sample Output | 170](#)
- [Sample Output | 173](#)
- [Release Information | 177](#)

Syntax (EX Series Switches)

```
show system statistics icmp6  
<all-members>  
<local>  
<member member-id>
```

Syntax (MX Series Routers, SRX Series)

```
show system statistics icmp6
```

Syntax (TX Matrix Router)

```
show system statistics icmp6  
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics icmp6  
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Internet Control Message Protocol for IPv6 (ICMPv6) statistics.

Options

none	Display system statistics for ICMPv6.
all-chassis	(Optional) Display system statistics for ICMPv6 for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system statistics for ICMPv6 for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMPv6 for all connected T1600 or T4000 LCCs.
all-members	(Optional) Display ICMPv6 statistics for all members of the Virtual Chassis configuration.
lcc <i>number</i>	<p>(Optional) On a TX Matrix router, display system statistics for ICMPv6 for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for ICMPv6 for a specific router that is connected to the TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
local	(Optional) Display ICMPv6 statistics for the local Virtual Chassis member.
member <i>member-id</i>	(Optional) Display ICMPv6 statistics for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value from 0 through 9.
scc	(Optional) Display system statistics for ICMPv6 for the TX Matrix router (or switch-card chassis).
sfc <i>number</i>	(Optional) Display system statistics for ICMPv6 for the TX Matrix Plus router. Replace <i>number</i> with 0.

Additional Information

By default, when you issue the `show system statistics icmp6` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

`show system statistics icmp6` (MX Series Routers)

```
user@host> show system statistics icmp6
icmp6:
    79 Calls to icmp_error
    0 Errors not generated because old message was icmp error
    0 Errors not generated because rate limitation
Output histogram:
    79 unreachable
    30 echo
    163 multicast listener query
    6 multicast listener report
    940 neighbor solicitation
    694184 neighbor advertisement
    0 Messages with bad code fields
    0 Messages < minimum length
    0 Bad checksums
    0 Messages with bad length
Input histogram:
    10 echo reply
    6 multicast listener report
    693975 neighbor solicitation
Histogram of error messages to be generated:
```

```

    0 No route
    0 Administratively prohibited
    0 Beyond scope
  79 Address unreachable
    0 Port unreachable
    0 Time exceed transit
    0 Time exceed reassembly
    0 Erroneous header field
    0 Unrecognized next header
    0 Unrecognized option
    0 Unknown
0 Message responses generated
0 Messages with too many ND options
100000 Max System ND nh cache limit
79840 Max Public ND nh cache limit
200 Max IRI ND nh cache limit
19960 Max Management intf ND nh cache limit
79840 Current Public ND nexthops present
4 Current IRI ND nexthops present
0 Current Management ND nexthops present
909266 Total ND nexthops creation failed as limit reached
909266 Public ND nexthops creation failed as public limit reached
0 IRI ND nexthops creation failed as iri limit reached
0 Management ND nexthops creation failed as mgt limit reached

```

show system statistics icmp6 (EX Series Switches)

```

user@host> show system statistics icmp6
icmp6:
    0 Calls to icmp_error
    0 Errors not generated because old message was icmp error
    0 Errors not generated because rate limitation
    0 Messages with bad code fields
    0 Messages < minimum length
    0 Bad checksums
    0 Messages with bad length
        0 No route
        0 Administratively prohibited
        0 Beyond scope
        0 Address unreachable
        0 Port unreachable

```

```

0 packet too big
0 Time exceed transit
0 Time exceed reassembly
0 Erroneous header field
0 Unrecognized next header
0 Unrecognized option
0 redirect
0 Unknown
0 Message responses generated
0 Messages with too many ND options

```

show system statistics icmp6 (SRX Series and vSRX 3.0)

```

0 Calls to icmp_error
0 Errors not generated because old message was icmp error
0 Errors not generated because rate limitation
0 Messages with bad code fields
0 Messages < minimum length
0 Bad checksums
0 Messages with bad length
Histogram of error messages to be generated:
0 No route
0 Administratively prohibited
0 Beyond scope
0 Address unreachable
0 Port unreachable
0 Time exceed transit
0 Time exceed reassembly
0 Erroneous header field
0 Unrecognized next header
0 Unrecognized option
0 Unknown
0 Message responses generated
0 Messages with too many ND options
100000 Max System ND nh cache limit
79840 Max Public ND nh cache limit
200 Max IRI ND nh cache limit
19960 Max Management intf ND nh cache limit
0 Current Public ND nexthops present
0 Current IRI ND nexthops present
0 Current Management ND nexthops present

```

```

0 Total ND nexthops creation failed as limit reached
0 Public ND nexthops creation failed as public limit reached
0 IRI ND nexthops creation failed as iri limit reached
0 Management ND nexthops creation failed as mgt limit reached
0 interface-restricted ndp proxy requests
0 interface-restricted dad proxy requests
0 interface-restricted ndp proxy responses
0 interface-restricted dad proxy conflicts
0 interface-restricted dad proxy duplicates
0 interface-restricted ndp proxy resolve requests
0 interface-restricted dad proxy resolve requests
0 interface-restricted dad packets from same node dropped
0 interface-restricted proxy packets dropped with nomac
0 interface-unrestricted ndp proxy requests
0 interface-unrestricted dad proxy requests
0 interface-unrestricted ndp proxy responses
0 interface-unrestricted dad proxy conflicts
0 interface-unrestricted dad proxy duplicates
0 interface-unrestricted ndp proxy resolve requests
0 interface-unrestricted dad proxy resolve requests
0 interface-unrestricted dad packets from same port dropped
0 interface-unrestricted proxy packets dropped with nomac
0 ND hold nexthops dropped on entry by RED mark
0 ND hold nexthops dropped on timer expire by RED mark

```

Sample Output

show system statistics icmp6 (TX Matrix Plus Router)

```

user@host> show system statistics icmp6
sfc0-re0:
-----
icmp6:
  0 calls to icmp_error
  0 errors not generated because old message was icmp error or so
  0 errors not generated because rate limitation
Output histogram:
  neighbor solicitation: 12
  neighbor advertisement: 4

```

```

0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length

```

Histogram of error messages to be generated:

```

    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options

```

lcc0-re0:

icmp6:

```

0 calls to icmp_error
0 errors not generated because old message was icmp error or so
0 errors not generated because rate limitation

```

Output histogram:

```

    neighbor solicitation: 12
    neighbor advertisement: 4
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length

```

Histogram of error messages to be generated:

```

    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly

```

```

    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options

```

lcc1-re0:

icmp6:

```

    0 calls to icmp_error
    0 errors not generated because old message was icmp error or so
    0 errors not generated because rate limitation

```

Output histogram:

```

    neighbor solicitation: 12
    neighbor advertisement: 4

```

```

0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length

```

Input histogram:

```

    neighbor advertisement: 2

```

Histogram of error messages to be generated:

```

    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown

```

```

0 message responses generated
0 messages with too many ND options

```

lcc2-re0:

icmp6:

```

0 calls to icmp_error
0 errors not generated because old message was icmp error or so
0 errors not generated because rate limitation
Output histogram:
    neighbor solicitation: 12
    neighbor advertisement: 4
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length
Input histogram:
    neighbor advertisement: 2
Histogram of error messages to be generated:
    0 no route
    0 administratively prohibited
    0 beyond scope
    0 address unreachable
    0 port unreachable
    0 packet too big
    0 time exceed transit
    0 time exceed reassembly
    0 erroneous header field
    0 unrecognized next header
    0 unrecognized option
    0 redirect
    0 unknown
0 message responses generated
0 messages with too many ND options

```

lcc3-re0:

icmp6:

```

0 calls to icmp_error
0 errors not generated because old message was icmp error or so
0 errors not generated because rate limitation
Output histogram:
    neighbor solicitation: 12
    neighbor advertisement: 4
0 messages with bad code fields
0 messages < minimum length
0 bad checksums
0 messages with bad length
Input histogram:

```

```

neighbor advertisement: 2
Histogram of error messages to be generated:
  0 no route
  0 administratively prohibited
  0 beyond scope
  0 address unreachable
  0 port unreachable
  0 packet too big
  0 time exceed transit
  0 time exceed reassembly
  0 erroneous header field
  0 unrecognized next header
  0 unrecognized option
  0 redirect
  0 unknown
  0 message responses generated
  0 messages with too many ND options

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

show system statistics igmp

IN THIS SECTION

- [Syntax | 178](#)
- [Syntax \(EX Series Switches\) | 178](#)

- [Syntax \(TX Matrix Router\) | 178](#)
- [Syntax \(TX Matrix Plus Router\) | 179](#)
- [Description | 179](#)
- [Options | 179](#)
- [Additional Information | 180](#)
- [Required Privilege Level | 180](#)
- [Sample Output | 180](#)
- [Release Information | 183](#)

Syntax

```
show system statistics igmp
```

Syntax (EX Series Switches)

```
show system statistics igmp  
<all-members>  
<local>  
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics igmp  
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics igmp
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Internet Group Management Protocol (IGMP) statistics.

Options

none	Display system statistics for IGMP.
all-chassis	(Optional) Display system statistics for IGMP for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system statistics for IGMP for all T640 routers (or line-card chassis) connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for all connected T1600 or T4000 LCCs.
all-members	(Optional) Display IGMP statistics for all members of the Virtual Chassis configuration.
lcc <i>number</i>	(Optional) On a TX Matrix router, display system statistics for IGMP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IGMP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local	(Optional) Display IGMP statistics for the local Virtual Chassis member.
member <i>member-id</i>	(Optional) Display IGMP statistics for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value from 0 through 9.
scc	(Optional) Display system statistics for IGMP for the TX Matrix router (or switch-card chassis).
sfc number	(Optional) Display system statistics for IGMP for the TX Matrix Plus router. Replace <i>number</i> with 0.

Additional Information

By default, when you issue the `show system statistics igmp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics igmp

```
user@host> show system statistics igmp
igmp:
    17178 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
```

```

0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

```

show system statistics igmp (EX Series Switches)

```

user@host> show system statistics igmp
igmp:
    0 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid fields
    0 membership reports received
    0 membership reports received with invalid fields
    0 membership reports received for groups to which we belong
    0 Membership reports sent

```

show system statistics igmp (TX Matrix Plus Router)

```

user@host> show system statistics igmp
sfc0-re0:
-----
igmp:
    0 messages received
    0 messages received with too few bytes
    0 messages received with bad checksum
    0 membership queries received
    0 membership queries received with invalid field(s)
    0 membership reports received
    0 membership reports received with invalid field(s)
    0 membership reports received for groups to which we belong
    0 membership reports sent

lcc0-re0:
-----
igmp:

```

```

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

```

lcc1-re0:

igmp:

```

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

```

lcc2-re0:

igmp:

```

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received
0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

```

lcc3-re0:

igmp:

```

0 messages received
0 messages received with too few bytes
0 messages received with bad checksum
0 membership queries received

```

```

0 membership queries received with invalid field(s)
0 membership reports received
0 membership reports received with invalid field(s)
0 membership reports received for groups to which we belong
0 membership reports sent

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

show system statistics ip

IN THIS SECTION

- [Syntax | 184](#)
- [Syntax \(EX Series Switches\) | 184](#)
- [Syntax \(TX Matrix Router\) | 184](#)
- [Syntax \(TX Matrix Plus Router\) | 184](#)
- [Description | 184](#)
- [Options | 185](#)
- [Additional Information | 186](#)
- [Required Privilege Level | 186](#)
- [Sample Output | 186](#)
- [Release Information | 194](#)

Syntax

```
show system statistics ip
```

Syntax (EX Series Switches)

```
show system statistics ip  
<all-members>  
<local>  
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics ip  
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics ip  
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide IPv4 statistics.

Some of the statistics include, the total number of IPv4 packets received, number of packets destined to dead next hop, number of fragments received and dropped, number of multicast packets dropped, number of Time-triggered Protocol over IP (TTPoIP) packets sent, received, and dropped, and number of raw packets dropped.

Options

none	Display system statistics for IPv4.
all-chassis	(Optional) Display system statistics for IPv4 for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system statistics for IPv4 for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv4 for all T1600 or T4000 routers connected to the TX Matrix Plus router.
all-members	(Optional) Display IPv4 statistics for all members of the Virtual Chassis configuration.
lcc <i>number</i>	<p>(Optional) On a TX Matrix router, display system statistics for IPv4 for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv4 for a specific router that is connected to the TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
local	(Optional) Display IPv4 statistics for the local Virtual Chassis member.
member <i>member-id</i>	(Optional) Display IPv4 statistics for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value from 0 through 9.
scc	(Optional) Display system statistics for IPv4 for the TX Matrix router (or switch-card chassis).
sfc <i>number</i>	(Optional) Display system statistics for IPv4 for the TX Matrix Plus router. Replace <i>number</i> with 0.

Additional Information

By default, when you issue the `show system statistics ip` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

`show system statistics ip`

```
user@host> show system statistics ip
ip:
    1752658 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragment sessions dropped (queue overflow)
    0 fragments dropped after timeout
    0 packets reassembled ok
    1709456 packets for this host
    10494 packets for unknown/unsupported protocol
    546 packets forwarded
    0 packets not forwardable
    546 redirects sent
    1340179 packets sent from this host
```

```

0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
10494 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
10494 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
0 raw packets dropped. no space in socket recv buffer

```

show system statistics ip (EX Series Switches)

```

user@host> show system statistics ip
ip:
    74121 total packets received
    0 bad header checksums
    0 with size smaller than minimum
    0 with data size < data length
    0 with header length < data size
    0 with data length < header length
    0 with incorrect version number
    0 packets destined to dead next hop
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragment sessions dropped (queue overflow)

```

```

0 fragments dropped after timeout
0 packets reassembled ok
1134061 packets for this host
0 packets for unknown/unsupported protocol
40177 packets forwarded
0 packets not forwardable
40177 redirects sent
1122558 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
0 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
0 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped

```

show system statistics ip (TX Matrix Plus Router)

```

user@host> show system statistics ip
sfc0-re0:

```

```

-----
ip:

```

```

47695035 total packets received
0 bad header checksums

```

```

0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
42350 fragments received
0 fragments dropped (dup or out of space)
0 fragment sessions dropped (queue overflow)
0 fragments dropped after timeout
21175 packets reassembled ok
47674941 packets for this host
146 packets for unknown/unsupported protocol
0 packets forwarded
0 packets not forwardable
0 redirects sent
61304579 packets sent from this host
8496 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
6746344 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
2400 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
2400 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
12995412 incoming ttpoip packets received
0 incoming ttpoip packets dropped
16959177 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
0 raw packets dropped. no space in socket recv buffer

```

lcc0-re0:

ip:

```

12990061 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragment sessions dropped (queue overflow)
0 fragments dropped after timeout
0 packets reassembled ok
12989979 packets for this host
82 packets for unknown/unsupported protocol
0 packets forwarded
0 packets not forwardable
0 redirects sent
9318381 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
3440 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
82 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
82 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received

```

```

0 incoming ttpoip packets dropped
548071 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
0 raw packets dropped. no space in socket recv buffer

```

lcc1-re0:

ip:

```

12849723 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragment sessions dropped (queue overflow)
0 fragments dropped after timeout
0 packets reassembled ok
12849641 packets for this host
82 packets for unknown/unsupported protocol
0 packets forwarded
0 packets not forwardable
0 redirects sent
7676351 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
82 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
82 router alert options

```

```

0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
0 raw packets dropped. no space in socket recv buffer

```

lcc2-re0:

ip:

```

16926850 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragment sessions dropped (queue overflow)
0 fragments dropped after timeout
0 packets reassembled ok
16926768 packets for this host
82 packets for unknown/unsupported protocol
0 packets forwarded
0 packets not forwardable
0 redirects sent
10039747 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets with bad options
82 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options

```

```

0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
82 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
0 raw packets dropped. no space in socket recv buffer

```

lcc3-re0:

ip:

```

18025026 total packets received
0 bad header checksums
0 with size smaller than minimum
0 with data size < data length
0 with header length < data size
0 with data length < header length
0 with incorrect version number
0 packets destined to dead next hop
0 fragments received
0 fragments dropped (dup or out of space)
0 fragment sessions dropped (queue overflow)
0 fragments dropped after timeout
0 packets reassembled ok
18024944 packets for this host
82 packets for unknown/unsupported protocol
0 packets forwarded
0 packets not forwardable
0 redirects sent
10456545 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented

```

```

0 packets with bad options
82 packets with options handled without error
0 strict source and record route options
0 loose source and record route options
0 record route options
0 timestamp options
0 timestamp and address options
0 timestamp and prespecified address options
0 option packets dropped due to rate limit
82 router alert options
0 multicast packets dropped (no iflist)
0 packets dropped (src and int don't match)
0 transit re packets dropped on mgmt i/f
0 packets used first nexthop in ecmp unilist
0 incoming ttpoip packets received
0 incoming ttpoip packets dropped
0 outgoing TTPoIP packets sent
0 outgoing TTPoIP packets dropped
0 raw packets dropped. no space in socket recv buffer

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

show system statistics ip6

IN THIS SECTION

- [Syntax | 195](#)
- [Syntax \(EX Series Switches\) | 195](#)
- [Syntax \(TX Matrix Router\) | 196](#)
- [Syntax \(TX Matrix Plus Router\) | 196](#)
- [Description | 196](#)
- [Options | 196](#)
- [Additional Information | 197](#)
- [Required Privilege Level | 197](#)
- [Sample Output | 198](#)
- [Release Information | 205](#)

Syntax

```
show system statistics ip6
```

Syntax (EX Series Switches)

```
show system statistics ip6  
<all-members>  
<local>  
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics ip6
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics ip
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide IPv6 statistics.

Options

none	Display system statistics for IPv6.
all-chassis	(TX Matrix routers and TX Matrix Plus routers only) (Optional) Display system statistics for IPv6 for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system statistics for IPv6 for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv6 for all connected T1600 or T4000 LCCs.
all-members	(Optional) Display IPv6 statistics for all members of the Virtual Chassis configuration.
lcc <i>number</i>	(Optional) On a TX Matrix router, display system statistics for IPv6 for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for IPv6 for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local	(Optional) Display IPv6 statistics for the local Virtual Chassis member.
member <i>member-id</i>	(Optional) Display IPv6 statistics for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value from 0 through 9.
scc	(Optional) Display system statistics for IPv6 for the TX Matrix router (or switch-card chassis).
sfc number	(Optional) Display system statistics for IPv6 for the TX Matrix Plus router. Replace <i>number</i> with 0.

Additional Information

By default, when you issue the `show system statistics ip6` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics ip6

```
user@host> show system statistics ip6
ip6:
    0 total packets received
    0 with size smaller than minimum
    0 with data size < data length
    0 with bad options
    0 with incorrect version number
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped after timeout
    0 fragment sessions dropped (queue overflow)
    0 packets reassembled ok
    0 packets for this host
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    0 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs, etc.
    0 output packets discarded due to no route
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets that violated scope rules
    0 multicast packets which we don't join
Mbuf statistics:
    0 packets whose headers are not continuous
    0 tunneling packets that can't find gif
    0 packets discarded due to too many headers
    0 failures of source address selection
    0 forward cache hit
    0 forward cache miss
    0 packets destined to dead next hop
    0 option packets dropped due to rate limit
    0 packets dropped (src and int don't match)
    0 packets dropped due to bad protocol
```

show system statistics ip6 (EX Series Switches)

```
user@host> show system statistics ip6
ip6:
    0 total packets received
    0 packets with size smaller than minimum
    0 packets with data size < data length
    0 packets with bad options
    0 packets with incorrect version number
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped after timeout
    0 fragment sessions dropped (queue overflow)
    0 packets reassembled ok
    0 packets for this host
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    0 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs, etc.
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets that violated scope rules
    0 multicast packets which we don't join
    0 packets whose headers are not continuous
    0 tunneling packets that can't find gif
    0 packets discarded due to too may headers
    0 failures of source address selection
    0 forward cache hit
    0 forward cache miss
    0 Packets destined to dead next hop
    0 option packets dropped due to rate limit
    0 Packets dropped (src and int don't match)
    0 packets dropped due to bad protocol
    0 transit re packet(null) dropped on mgmt i/f
```

show system statistics ip6 (TX Matrix Router)

```

user@host> show system statistics ip6
sfc0-re0:
-----
ip6:
    0 total packets received
    0 with size smaller than minimum
    0 with data size < data length
    0 with bad options
    0 with incorrect version number
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped after timeout
    0 fragment sessions dropped (queue overflow)
    0 packets reassembled ok
    0 packets for this host
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    0 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs, etc.
    0 output packets discarded due to no route
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets that violated scope rules
    0 multicast packets which we don't join
Mbuf statistics:
    0 packets whose headers are not continuous
    0 tunneling packets that can't find gif
    0 packets discarded due to too may headers
    0 failures of source address selection
source addresses on an outgoing I/F
    4 link-locals
source addresses of same scope
    4 link-locals
    0 forward cache hit
    0 forward cache miss
    0 packets destined to dead next hop
    0 option packets dropped due to rate limit

```

```

0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
0 packet(null) used first nexthop in ecmp unilist

```

lcc0-re0:

ip6:

```

0 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragment sessions dropped (queue overflow)
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join

```

Mbuf statistics:

```

0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection

```

source addresses on an outgoing I/F

```

    4 link-locals

```

source addresses of same scope

```

    4 link-locals

```

```

0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop

```

```

0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
0 packet(null) used first nexthop in ecmp unilist

```

lcc1-re0:

ip6:

```

2 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragment sessions dropped (queue overflow)
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join

```

Input histogram:

ICMP6: 2

Mbuf statistics:

```

0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
source addresses on an outgoing I/F

```

4 link-locals

source addresses of same scope

4 link-locals

```

0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
0 packet(null) used first nexthop in ecmp unilist

```

lcc2-re0:

ip6:

```

2 total packets received
0 with size smaller than minimum
0 with data size < data length
0 with bad options
0 with incorrect version number
0 fragments received
0 fragments dropped (dup or out of space)
0 fragments dropped after timeout
0 fragment sessions dropped (queue overflow)
0 packets reassembled ok
0 packets for this host
0 packets forwarded
0 packets not forwardable
0 redirects sent
0 packets sent from this host
0 packets sent with fabricated ip header
0 output packets dropped due to no bufs, etc.
0 output packets discarded due to no route
0 output datagrams fragmented
0 fragments created
0 datagrams that can't be fragmented
0 packets that violated scope rules
0 multicast packets which we don't join

```

Input histogram:

ICMP6: 2

Mbuf statistics:

```

0 packets whose headers are not continuous
0 tunneling packets that can't find gif
0 packets discarded due to too many headers
0 failures of source address selection
source addresses on an outgoing I/F

```

```

    4 link-locals
source addresses of same scope
    4 link-locals
0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
0 packet(null) used first nexthop in ecmp unilist

```

lcc3-re0:

ip6:

```

    2 total packets received
    0 with size smaller than minimum
    0 with data size < data length
    0 with bad options
    0 with incorrect version number
    0 fragments received
    0 fragments dropped (dup or out of space)
    0 fragments dropped after timeout
    0 fragment sessions dropped (queue overflow)
    0 packets reassembled ok
    0 packets for this host
    0 packets forwarded
    0 packets not forwardable
    0 redirects sent
    0 packets sent from this host
    0 packets sent with fabricated ip header
    0 output packets dropped due to no bufs, etc.
    0 output packets discarded due to no route
    0 output datagrams fragmented
    0 fragments created
    0 datagrams that can't be fragmented
    0 packets that violated scope rules
    0 multicast packets which we don't join
Input histogram:
    ICMP6: 2
Mbuf statistics:
    0 packets whose headers are not continuous
    0 tunneling packets that can't find gif

```

```

0 packets discarded due to too may headers
0 failures of source address selection
source addresses on an outgoing I/F
    4 link-locals
source addresses of same scope
    4 link-locals
0 forward cache hit
0 forward cache miss
0 packets destined to dead next hop
0 option packets dropped due to rate limit
0 packets dropped (src and int don't match)
0 packets dropped due to bad protocol
0 transit re packet(null) dropped on mgmt i/f
0 packet(null) used first nexthop in ecmp unilist

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Routing Matrix with a TX Matrix Plus Router Solutions Page](#)

show system statistics tcp

IN THIS SECTION

- [Syntax | 206](#)
- [Syntax \(EX Series Switches\) | 206](#)
- [Syntax \(TX Matrix Router\) | 206](#)
- [Syntax \(TX Matrix Plus Router\) | 207](#)

- [Description | 207](#)
- [Options | 207](#)
- [Additional Information | 208](#)
- [Required Privilege Level | 208](#)
- [Sample Output | 208](#)
- [Release Information | 218](#)

Syntax

```
show system statistics tcp
```

Syntax (EX Series Switches)

```
show system statistics tcp  
<all-members>  
<local>  
<member member-id>
```

Syntax (TX Matrix Router)

```
show system statistics tcp  
<all-chassis | all-lcc | lcc number / scc>
```

Syntax (TX Matrix Plus Router)

```
show system statistics tcp
<all-chassis | all-lcc | lcc number | sfc number>
```

Description

Display system-wide Transmission Control Protocol (TCP) statistics.

Options

none	Display system statistics for TCP.
all-chassis	(Optional) Display system statistics for TCP for all the routers in the chassis.
all-lcc	(Optional) On a TX Matrix router, display system statistics for TCP for all T640 routers connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for TCP for all connected T1600 or T4000 LCCs.
all-members	(Optional) Display TCP statistics for all members of the Virtual Chassis configuration.
lcc <i>number</i>	(Optional) On a TX Matrix router, display system statistics for TCP for a specific T640 router that is connected to the TX Matrix router. On a TX Matrix Plus router, display system statistics for TCP for a specific router that is connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

local	(Optional) Display TCP statistics for the local Virtual Chassis member.
member <i>member-id</i>	(Optional) Display TCP statistics for the specified member of the Virtual Chassis configuration. Replace <i>member-id</i> with a value from 0 through 9.
scc	(Optional) Display system statistics for TCP for the TX Matrix router (or switch-card chassis).
sfc number	(Optional) Display system statistics for TCP for the TX Matrix Plus router. Replace <i>number</i> with 0.

Additional Information

By default, when you issue the `show system statistics tcp` command on the primary Routing Engine of a TX Matrix router or a TX Matrix Plus router, the command is broadcast to all the primary Routing Engines of the LCCs connected to it in the routing matrix. Likewise, if you issue the same command on the backup Routing Engine of a TX Matrix or a TX Matrix Plus router, the command is broadcast to all backup Routing Engines of the LCCs that are connected to it in the routing matrix.

Required Privilege Level

view

Sample Output

show system statistics tcp

```
user@host> show system statistics tcp
tcp:
    3844 packets sent
        3618 data packets (1055596 bytes)
        0 data packets (0 bytes) retransmitted
```

```

    0 resends initiated by MTU discovery
    205 ack-only packets (148 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    1079 control packets
5815 packets received
    3377 acks (for 1055657 bytes)
    24 duplicate acks
    0 acks for unsent data
    2655 packets (15004 bytes) received in-sequence
    1 completely duplicate packet (0 bytes)
    0 old duplicate packets
    0 packets with some dup. data (0 bytes duped)
    0 out-of-order packets (0 bytes)
    0 packets (0 bytes) of data after window
    0 window probes
    7 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
1 connection request
32 connection accepts
0 bad connection attempts
0 listen queue overflows
33 connections established (including accepts)
30 connections closed (including 0 drops)
    27 connections updated cached RTT on close
    27 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
3374 segments updated rtt (of 3220 attempts)
0 retransmit timeouts
    0 connections dropped by rexmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
344 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
1096 correct ACK header predictions
1314 correct data packet header predictions
32 syncache entries added

```

```

    0 retransmitted
    0 dupsyn
    0 dropped
    32 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
1058 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors

```

show system statistics tcp (EX Series Switches)

```

user@host> show system statistics tcp
Tcp:
    572724 packets sent
        21936 data packets (1887657 bytes)
        2 data packets retransmitted (20 bytes)
        0 resends initiated by MTU discovery
        3724 ack only packets (537 packets delayed)
        0 URG only packets
        1 window probe packets
        1 window update packets
        1094083 control packets
    1134258 packets received
        21371 acks(for 1886660 bytes)
        5870 duplicate acks
        0 acks for unsent data
        19908 packets received in-sequence(267794 bytes)

```

```

    3022 completely duplicate packets(0 bytes)
    0 old duplicate packets
    4 packets with some duplicate data(4 bytes duped)
    2 out-of-order packets(2 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    40 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
547027 connection requests
80 connection accepts
0 bad connection attempts
0 listen queue overflows
103 connections established (including accepts)
547106 connections closed (including 6 drops)
    47 connections updated cached RTT on close
    47 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
547004 embryonic connections dropped
20862 segments updated rtt(of 567830 attempts)
2 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
3032 keepalive timeouts
    3031 keepalive probes sent
    1 connections dropped by keepalive
7823 correct ACK header predictions
12533 correct data packet header predictions
80 syncache entries added
    0 retransmitted
    0 dupsyn
    4 dropped
    80 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach

```

```

    0 zone failures
0 cookies sent
0 cookies received
1 SACK recovery episodes
1 segment retransmits in SACK recovery episodes
1 byte retransmits in SACK recovery episodes
71 SACK options (SACK blocks) received
1 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
547024 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

show system statistics tcp lcc (TX Matrix Router)

```

user@host> show system statistics tcp lcc 2
lcc2-re0:
-----
tcp:
    21271 packets sent
        11069 data packets (12044 bytes)
        0 data packets (0 bytes) retransmitted
        0 resends initiated by MTU discovery
        10198 ack-only packets (10194 packets delayed)
        0 URG only packets
        0 window probe packets
        0 window update packets
        4 control packets
    13363 packets received
        11073 acks (for 12044 bytes)
        0 duplicate acks
        0 acks for unsent data
        12895 packets (2400874 bytes) received in-sequence
        0 completely duplicate packets (0 bytes)
        0 old duplicate packets

```

```

    0 packets with some dup. data (0 bytes duped)
    0 out-of-order packets (0 bytes)
    0 packets (0 bytes) of data after window
    0 window probes
    0 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
4 connection requests
0 connection accepts
0 bad connection attempts
0 listen queue overflows
4 connections established (including accepts)
33 connections closed (including 0 drops)
    0 connections updated cached RTT on close
    0 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
11073 segments updated rtt (of 11073 attempts)
0 retransmit timeouts
    0 connections dropped by rexmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
464 correct ACK header predictions
2172 correct data packet header predictions
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 out-of-sequence segment drops due to insufficient memory
0 RST packets
0 ICMP packets ignored by TCP

```

show system statistics tcp (TX Matrix Plus Router)

```

user@host> show system statistics tcp
sfc0-re0:

```

```

-----
Tcp:

```

```

10420 packets sent
    10203 data packets (2374613 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    202 ack only packets (120 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    30 control packets
16635 packets received
    9468 acks(for 2374674 bytes)
    32 duplicate acks
    0 acks for unsent data
    7764 packets received in-sequence(38286 bytes)
    20 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    356 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
10 connection requests
33 connection accepts
0 bad connection attempts
0 listen queue overflows
34 connections established (including accepts)
50 connections closed (including 0 drops)
    24 connections updated cached RTT on close
    24 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
9 embryonic connections dropped
9468 segments updated rtt(of 9256 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
14 keepalive timeouts
    14 keepalive probes sent
    0 connections dropped by keepalive

```

```

6220 correct ACK header predictions
6625 correct data packet header predictions
33 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    33 completed
    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
15 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

lcc0-re0:

 Tcp:

```

1306 packets sent
    1251 data packets (161855 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    51 ack only packets (1 packets delayed)
    0 URG only packets
    0 window probe packets

```

```

    0 window update packets
    6 control packets
1397 packets received
    1218 acks(for 161904 bytes)
    2 duplicate acks
    0 acks for unsent data
    612 packets received in-sequence(12495 bytes)
    0 completely duplicate packets(0 bytes)
    0 old duplicate packets
    0 packets with some duplicate data(0 bytes duped)
    0 out-of-order packets(0 bytes)
    0 packets of data after window(0 bytes)
    0 window probes
    22 window update packets
    0 packets received after close
    0 discarded for bad checksums
    0 discarded for bad header offset fields
    0 discarded because packet too short
1 connection requests
24 connection accepts
0 bad connection attempts
0 listen queue overflows
25 connections established (including accepts)
27 connections closed (including 0 drops)
    24 connections updated cached RTT on close
    24 connections updated cached RTT variance on close
    0 connections updated cached ssthresh on close
0 embryonic connections dropped
1218 segments updated rtt(of 1192 attempts)
0 retransmit timeouts
    0 connections dropped by retransmit timeout
0 persist timeouts
    0 connections dropped by persist timeout
0 keepalive timeouts
    0 keepalive probes sent
    0 connections dropped by keepalive
196 correct ACK header predictions
119 correct data packet header predictions
24 syncache entries added
    0 retransmitted
    0 dupsyn
    0 dropped
    24 completed

```

```

    0 bucket overflow
    0 cache overflow
    0 reset
    0 stale
    0 aborted
    0 badack
    0 unreach
    0 zone failures
0 cookies sent
0 cookies received
0 SACK recovery episodes
0 segment retransmits in SACK recovery episodes
0 byte retransmits in SACK recovery episodes
0 SACK options (SACK blocks) received
0 SACK options (SACK blocks) sent
0 SACK scoreboard overflow
0 ACKs sent in response to in-window but not exact RSTs
0 ACKs sent in response to in-window SYNs on established connections
0 rcv packets dropped by TCP due to bad address
0 out-of-sequence segment drops due to insufficient memory
2 RST packets
0 ICMP packets ignored by TCP
0 send packets dropped by TCP due to auth errors
0 rcv packets dropped by TCP due to auth errors
0 outgoing segments dropped due to policing

```

lcc1-re0:

Tcp:

```

1118 packets sent
    1066 data packets (131896 bytes)
    0 data packets retransmitted (0 bytes)
    0 resends initiated by MTU discovery
    48 ack only packets (2 packets delayed)
    0 URG only packets
    0 window probe packets
    0 window update packets
    6 control packets
1215 packets received

```

show system statistics tcp (Junos OS Evolved)

```

user@host> show system statistics tcp
Tcp:
    2635574 packets sent
        0 window probe packets
    1124324 packets received
        0 discarded for bad checksums
        0 discarded for bad header offset fields
    3495 connection requests
    2371 bad connection attempts
    0 listen queue overflows
    1574 connections established (including accepts)
    1 embryonic connections dropped
        22 connections dropped by retransmit timeout
        3677 keepalive probes sent
        8 retransmitted
        17 reset
        0 aborted
    0 cookies sent
    0 cookies received
    199 SACK recovery episodes
    13743 segment retransmits in SACK recovery episodes
    0 out-of-sequence segment drops due to insufficient memory
    2365 RST packets

```

Release Information

Command introduced before Junos OS Release 7.4.

sfc option introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

| [Routing Matrix with a TX Matrix Plus Router Solutions Page](#)