JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

IPv6 Neighbor Discovery User Guide

Published
2020-09-29

JUNOS

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 IPv6 Neighbor Discovery User Guide
Copyright © 2020 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 the Documentation | vii

Documentation and Release Notes | **vii**

Using the Examples in This Manual | **vii**

- Merging a Full Example | **viii**
- Merging a Snippet | **ix**

Documentation Conventions | **ix**

Documentation Feedback | **xii**

Requesting Technical Support | **xii**

- Self-Help Online Tools and Resources | **xiii**
- Creating a Service Request with JTAC | **xiii**

1 **Configuring IPv6 Neighbor Discovery**

IPv6 Neighbor Discovery | 15

IPv6 Neighbor Discovery Overview | **15**

- Improvements Over Ipv4 Protocols | **16**
- Router Discovery | **17**
- Address Resolution | **17**
- Redirect | **17**
- SLAAC | **17**

Supported ICMP Router Discovery and IPv6 Neighbor Discovery Standards | **18**

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery | **19**

Secure IPv6 Neighbor Discovery | 30

- Understanding Secure IPv6 Neighbor Discovery | **31**
- Example: Configuring Secure IPv6 Neighbor Discovery | **31**

NDP Proxy and DAD Proxy | 36

- Configuring NDP Proxy | **36**
- Configuring DAD Proxy | **37**

Neighbor Discovery Cache Protection | 37

Neighbor Discovery Cache Protection Overview | **38**

Configuring Neighbor Discovery Cache Protection | **39**

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks | **40**

2 **Troubleshooting**

Troubleshooting Network Issues | 47

Working with Problems on Your Network | **47**

Isolating a Broken Network Connection | **48**

Identifying the Symptoms of a Broken Network Connection | **49**

Isolating the Causes of a Network Problem | **51**

Taking Appropriate Action for Resolving the Network Problem | **52**

Evaluating the Solution to Check Whether the Network Problem Is Resolved | **53**

Checklist for Tracking Error Conditions | **54**

Configure Routing Protocol Process Tracing | **57**

Configure Routing Protocol Tracing for a Specific Routing Protocol | **60**

Monitor Trace File Messages Written in Near-Real Time | **62**

Stop Trace File Monitoring | **63**

3 **Configuration Statements**

autonomous | 67

cryptographic-address | 68

current-hop-limit | 69

default-lifetime | 70

interface (Protocols IPv6 Neighbor Discovery) | 71

key-length | 73

key-pair | 74

link-mtu | 75

managed-configuration | 76

max-advertisement-interval (Protocols IPv6 Neighbor Discovery) | 77

min-advertisement-interval (Protocols IPv6 Neighbor Discovery) | 78

nd-retransmit-timer | 79

nd-system-cache-limit | 80

nd6-max-cache | 81

nd6-new-hold-limit | 82

neighbor-discovery | 84

on-link | 86

onlink-subnet-only | 87

other-stateful-configuration | 88

preference (IPv6 Router Advertisement) | 89

preferred-lifetime | 90

prefix (Protocols IPv6 Neighbor Discovery) | 91

reachable-time | 92

retransmit-timer | 93

router-advertisement | 94

secure | 95

security-level | 96

solicit-router-advertisement-unicast | 97

timestamp | 98

traceoptions (Protocols IPv6 Neighbor Discovery) | 100

traceoptions (Protocols Secure Neighbor Discovery) | 102

valid-lifetime | 104

4 Operational Commands

clear ipv6 neighbors | 106

clear ipv6 router-advertisement | 108

monitor interface | 109

monitor start | 125

monitor stop | 127

ping | 129

show ipv6 neighbors | 137

show ipv6 router-advertisement | 140

show log | 144

traceroute | 152

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | vii
- Using the Examples in This Manual | vii
- Documentation Conventions | ix
- Documentation Feedback | xii
- Requesting Technical Support | xii

Use this guide to configure, monitor, and troubleshoot the IPv6 neighbor discovery on your Juniper Network devices.

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at https://www.juniper.net/documentation/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at https://www.juniper.net/books.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
       scripts {
           commit {
               file ex-script.xsl;
           }
       }
   }
   interfaces {
       fxp0 {
           disable;
           unit 0 {
               family inet {
                   address 10.0.0.1/24;
               }
           }
       }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

   ```
   [edit]
   user@host# edit system scripts
   [edit system scripts]
   ```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

   ```
   [edit system scripts]
   user@host# load merge relative /var/tmp/ex-script-snippet.conf
   load complete
   ```

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

Table 1 on page x defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page x defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

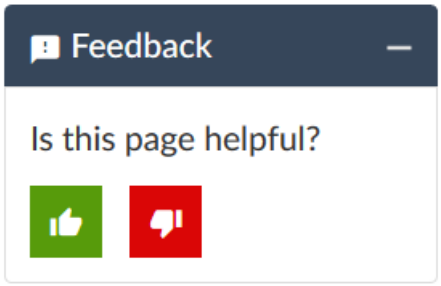

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://myjuniper.juniper.net

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/.

CHAPTER 1

Configuring IPv6 Neighbor Discovery

IPv6 Neighbor Discovery | **15**

Secure IPv6 Neighbor Discovery | **30**

NDP Proxy and DAD Proxy | **36**

Neighbor Discovery Cache Protection | **37**

IPv6 Neighbor Discovery

SUMMARY

Neighbor discovery is a protocol used for IPv6 traffic that allows different nodes on the same link to advertise their existence to their neighbors, and to learn about the existence of their neighbors.

IN THIS SECTION

- IPv6 Neighbor Discovery Overview | **15**
- Supported ICMP Router Discovery and IPv6 Neighbor Discovery Standards | **18**
- Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery | **19**

IPv6 Neighbor Discovery Overview

IN THIS SECTION

- Improvements Over Ipv4 Protocols | **16**
- Router Discovery | **17**
- Address Resolution | **17**
- Redirect | **17**
- SLAAC | **17**

Neighbor discovery is a protocol that allows different nodes on the same link to advertise their existence to their neighbors, and to learn about the existence of their neighbors.

Routers and hosts (nodes) use Neighbor Discovery (ND) messages to determine the link-layer addresses of neighbors that reside on attached links and to overwrite invalid cache entries. Hosts also use ND to find neighboring routers that can forward packets on their behalf.

In addition, nodes use ND to actively track the ability to reach neighbors. When a router (or the path to a router) fails, nodes actively search for alternatives to reach the destination.

This section discusses the following topics:

Improvements Over Ipv4 Protocols

IPv6 Neighbor Discovery corresponds to a number of the IPv4 protocols — ARP, ICMP Router Discovery, and ICMP Redirect. However, Neighbor Discovery provides many improvements over the IPv4 set of protocols. These improvements address the following:

- Router discovery—How a host locates routers residing on an attached link.
- Prefix discovery—How a host discovers address prefixes for destinations residing on an attached link. Nodes use prefixes to distinguish between destinations that reside on an attached link and those destinations that it can reach only through a router.
- Parameter discovery—How a node learns various parameters (link parameters or Internet parameters) that it places in outgoing packets.
- Address resolution—How a node uses only a destination IPv6 address to determine a link-layer address for destinations on an attached link.
- Next-hop determination—The algorithm that a node uses for mapping an IPv6 destination address into a neighbor IPv6 address (either the next router hop or the destination itself) to which it plans to send traffic for the destination.
- Neighbor unreachability detection—How a node determines that it can no longer reach a neighbor.
- Duplicate address detection—How a node determines whether an address is already in use by another node.

A router periodically multicasts a router advertisement from each of its multicast interfaces, announcing its availability. Hosts listen for these advertisements for address autoconfiguration and discovery of link-local addresses of the neighboring routers. When a host starts, it multicasts a router solicitation to ask for immediate advertisements.

The router discovery messages do not constitute a routing protocol. They enable hosts to discover the existence of neighboring routers, but are not used to determine which router is best to reach a particular destination.

Neighbor discovery uses the following Internet Control Message Protocol version 6 (ICMPv6) messages: router solicitation, router advertisement, neighbor solicitation, neighbor advertisement, and redirect.

Neighbor discovery for IPv6 replaces the following IPv4 protocols: router discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

Junos OS Release 9.3 and later supports Secure Neighbor Discovery (SEND). SEND enables you to secure Neighbor Discovery protocol (NDP) messages. It is applicable in environments where physical security on a link is not assured and attacks on NDP messages are a concern. The Junos OS secures NDP messages through cryptographically generated addresses (CGAs).

Router Discovery

Router advertisements can contain a list of prefixes. These prefixes are used for address autoconfiguration, to maintain a database of onlink (on the same data link) prefixes, and for duplication address detection. If a node is onlink, the router forwards packets to that node. If the node is not onlink, the packets are sent to the next router for consideration. For IPv6, each prefix in the prefix list can contain a prefix length, a valid lifetime for the prefix, a preferred lifetime for the prefix, an onlink flag, and an autoconfiguration flag. This information enables address autoconfiguration and the setting of link parameters such as maximum transmission unit (MTU) size and hop limit.

Address Resolution

For IPv6, ICMPv6 neighbor discovery replaces Address Resolution Protocol (ARP) for resolving network addresses to link-level addresses. Neighbor discovery also handles changes in link-layer addresses, inbound load balancing, anycast addresses, and proxy advertisements.

Nodes requesting the link-layer address of a target node multicast a neighbor solicitation message with the target address. The target sends back a neighbor advertisement message containing its link-layer address.

Neighbor solicitation and advertisement messages are used for detecting duplicate unicast addresses on the same link. Autoconfiguration of an IP address depends on whether there is a duplicate address on that link. Duplicate address detection is a requirement for autoconfiguration.

Neighbor solicitation and advertisement messages are also used for neighbor unreachability detection. Neighbor unreachability detection involves detecting the presence of a target node on a given link.

Redirect

Redirect messages are sent to inform a host of a better next-hop router to a particular destination or an onlink neighbor. This is similar to ICMPv4 redirect. Very similar to the ICMPv4 Redirect feature, the ICMPv6 redirect message is used by routers to inform on-link hosts of a better next-hop for a given destination. The intent is to allow the routers to help hosts make the most efficient local routing decisions possible.

SLAAC

In addition to all the other improvements it brings to the networking world, Neighbor Discovery also enables address autoconfiguration, namely Stateless Address Autoconfiguration (SLAAC). IPv6 maintains the capability for stateful address assignment through DHCPv6 (and static assignment), but SLAAC provides a lightweight address configuration method that might be desirable in many circumstances.

SLAAC provides plug-and-play IP connectivity in two phases: Phase 1: Link-local address assignment; and then, in Phase 2: Global address assignment.

- Phase 1—Steps for local connectivity:

  1. Link-Local Address Generation: Any time that a multicast-capable IPv6-enabled interface is turned up, the node generates a link-local address for that interface. This is done by appending an interface identifier to the link-local prefix (FE80::/10). The auto generated link-local address cannot be deleted. However, a new link-local address can also be manually entered, which overwrites the auto generated link-local address.
  2. Duplicate Detection: Before assigning the new link-local address to its interface, the node verifies that the address is unique. This is accomplished by sending a Neighbor Solicitation message destined to the new address. If there is a reply, then the address is a duplicate and the process stops, requiring operator intervention.
  3. Link-Local Address Assignment: If the address is unique, the node assigns it to the interface for which it was generated.

  At this point, the node has IPv6 connectivity to all other nodes on the same link. Phase 2 can only be completed by hosts. The router's interface addresses must be configured by other means.

- Phase 2—Steps for global connectivity:

  1. Router Advertisement: The node sends a Router Solicitation to prompt all on-link routers to send it router advertisements. When the router is enabled to provide stateless autoconfiguration support, the router advertisement contains a subnet prefix for use by neighboring hosts.
  2. Global Address Generation: Once it receives a subnet prefix from a router, the host generates a global address by appending the interface id to the supplied prefix.
  3. Duplicate Address Detection: The host again performsDuplicate Address Detection (DAD), this time for the new global address. 4. Global Address Assignment: Assuming that the address is not a duplicate, the host assigns it to the interface.

  This process ensures full IPv6 global connectivity with no manual host configuration and very little router configuration.

Supported ICMP Router Discovery and IPv6 Neighbor Discovery Standards

Junos OS substantially supports the following RFCs, which define standards for the Internet Control Message Protocol (ICMP for IP version 4 [IPv4]) and neighbor discovery (for IP version 6 [IPv6]).

- RFC 1256, *ICMP Router Discovery Messages*
- RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
- RFC 2462, *IPv6 Stateless Address Autoconfiguration*
- RFC 2463, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*

- RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
- RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*
- RFC 4862, *IPv6 Stateless Address Autoconfiguration*
- RFC 8335, *PROBE: A Utility for Probing Interfaces*

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

IN THIS SECTION

- Requirements | **19**
- Overview | **19**
- Configuration | **21**
- Verification | **25**

This example shows how to configure the router or switch to send IPv6 neighbor discovery messages.

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, all of the interfaces in the sample topology are configured with IPv6 addresses. If you plan to extend IPv6 functionality into your LAN, datacenter, or customer networks, you might want to use Stateless Address Auto-Configuration (SLAAC) and that means configuring router advertisements. SLAAC is an IPv6 protocol that provides some similar functionality to DHCP in IPv4. Using SLAAC, network hosts can autoconfigure a globally unique IPv6 address based on the prefix provided by a nearby router in a router advertisement. This removes the need to explicitly configure every interface in a given section of the network. Router advertisement messages are disabled by default, and you must enable them to take advantage of SLAAC.

To configure the router to send router advertisement messages, you must include at least the following statements in the configuration. All other router advertisement configuration statements are optional.

```
protocols {
    router-advertisement {
        interface interface-name {
            prefix prefix;
        }
    }
}
```

To configure neighbor discovery, include the following statements. You configure router advertisement on a per-interface basis.

```
protocols {
    router-advertisement {
        interface interface-name {
            current-hop-limit number;
            default-lifetime seconds;
            (link-mtu | no-link-mtu);
            (managed-configuration | no-managed-configuration);
            max-advertisement-interval seconds;
            min-advertisement-interval seconds;
            (other-stateful-configuration | no-other-stateful-configuration);
            prefix prefix {
                (autonomous | no-autonomous);
                (on-link | no-on-link);
                preferred-lifetime seconds;
                valid-lifetime seconds;
            }
            reachable-time milliseconds;
            retransmit-timer milliseconds;
            solicit-router-advertisement-unicast;
            virtual-router-only;
        }
        traceoptions {
            file filename <files number> <size maximum-file-size> <world-readable | no-world-readable>;
            flag flag;
        }
    }
}
```

Figure 1 on page 21 shows a simplified sample topology.

Figure 1: ICMP Router Discover Topology

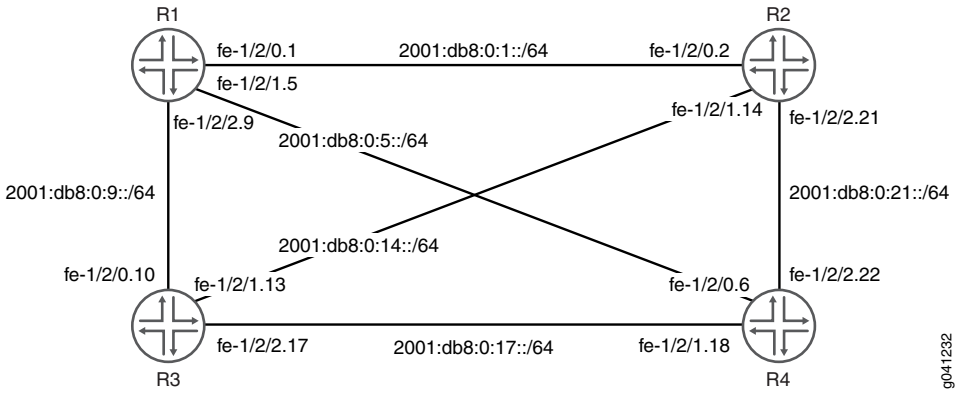

This example shows how to make sure that all of the IPv6 hosts attached to the subnets in the sample topology can auto-configure a local EUI-64 address.

"CLI Quick Configuration" on page 21 shows the configuration for all of the devices in Figure 1 on page 21. "Step-by-Step Procedure" on page 23 describes the steps on Device R1.

Configuration

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 at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-P2
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces fe-1/2/1 unit 5 description to-P4
set interfaces fe-1/2/1 unit 5 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces fe-1/2/2 unit 9 description to-P3
set interfaces fe-1/2/2 unit 9 family inet6 address 2001:db8:0:9::/64 eui-64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set protocols router-advertisement interface fe-1/2/0.1 prefix 2001:db8:0:1::/64
set protocols router-advertisement interface fe-1/2/1.5 prefix 2001:db8:0:5::/64
set protocols router-advertisement interface fe-1/2/2.9 prefix 2001:db8:0:9::/64
```

Device R2

```
set interfaces fe-1/2/0 unit 2 description to-P1
set interfaces fe-1/2/0 unit 2 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces fe-1/2/1 unit 14 description to-P3
set interfaces fe-1/2/1 unit 14 family inet6 address 2001:db8:0:14::/64 eui-64
set interfaces fe-1/2/2 unit 21 description to-P4
set interfaces fe-1/2/2 unit 21 family inet6 address 2001:db8:0:21::/64 eui-64
set interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set protocols router-advertisement interface fe-1/2/0.2 prefix 2001:db8:0:1::/64
set protocols router-advertisement interface fe-1/2/1.14 prefix 2001:db8:0:14::/64
set protocols router-advertisement interface fe-1/2/2.21 prefix 2001:db8:0:21::/64
```

Device R3

```
set interfaces fe-1/2/0 unit 10 description to-P1
set interfaces fe-1/2/0 unit 10 family inet6 address 2001:db8:0:9::/64 eui-64
set interfaces fe-1/2/1 unit 13 description to-P2
set interfaces fe-1/2/1 unit 13 family inet6 address 2001:db8:0:14::/64 eui-64
set interfaces fe-1/2/2 unit 17 description to-P4
set interfaces fe-1/2/2 unit 17 family inet6 address 2001:db8:0:17::/64 eui-64
set interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set protocols router-advertisement interface fe-1/2/0.10 prefix 2001:db8:0:9::/64
set protocols router-advertisement interface fe-1/2/1.13 prefix 2001:db8:0:14::/64
set protocols router-advertisement interface fe-1/2/2.17 prefix 2001:db8:0:17::/64
```

Device R4

```
set interfaces fe-1/2/0 unit 6 description to-P1
set interfaces fe-1/2/0 unit 6 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces fe-1/2/1 unit 18 description to-P3
set interfaces fe-1/2/1 unit 18 family inet6 address 2001:db8:0:17::/64 eui-64
set interfaces fe-1/2/2 unit 22 description to-P2
set interfaces fe-1/2/2 unit 22 family inet6 address 2001:db8:0:21::/64 eui-64
set interfaces lo0 unit 4 family inet6 address 2001:db8::4/128
set protocols router-advertisement interface fe-1/2/0.6 prefix 2001:db8:0:5::/64
set protocols router-advertisement interface fe-1/2/1.18 prefix 2001:db8:0:17::/64
set protocols router-advertisement interface fe-1/2/2.22 prefix 2001:db8:0:21::/64
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a IPv6 neighbor discovery:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 description to-P2
   user@R1# set fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
   user@R1# set fe-1/2/1 unit 5 description to-P4
   user@R1# set fe-1/2/1 unit 5 family inet6 address 2001:db8:0:5::/64 eui-64
   user@R1# set fe-1/2/2 unit 9 description to-P3
   user@R1# set fe-1/2/2 unit 9 family inet6 address 2001:db8:0:9::/64 eui-64
   user@R1# set lo0 unit 1 family inet6 address 2001:db8::1/128
   ```

2. Enable neighbor discovery.

   ```
   [edit protocols router-advertisement]
   user@R1# set interface fe-1/2/0.1 prefix 2001:db8:0:1::/64
   user@R1# set interface fe-1/2/1.5 prefix 2001:db8:0:5::/64
   user@R1# set interface fe-1/2/2.9 prefix 2001:db8:0:9::/64
   ```

Results

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

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        description to-P2;
        family inet6 {
            address 2001:db8:0:1::/64 {
                eui-64;
            }
        }
    }
}
fe-1/2/1 {
```

```
    unit 5 {
        description to-P4;
        family inet6 {
            address 2001:db8:0:5::/64 {
                eui-64;
            }
        }
    }
}
fe-1/2/2 {
    unit 9 {
        description to-P3;
        family inet6 {
            address 2001:db8:0:9::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 1 {
        family inet6 {
            address 2001:db8::1/128;
        }
    }
}
```

```
user@R1# show protocols
router-advertisement {
    interface fe-1/2/0.1 {
        prefix 2001:db8:0:1::/64;
    }
    interface fe-1/2/1.5 {
        prefix 2001:db8:0:5::/64;
    }
    interface fe-1/2/2.9 {
        prefix 2001:db8:0:9::/64;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- Checking the Interfaces | 25
- Pinging the Interfaces | 26
- Checking the IPv6 Neighbor Cache | 26
- Verifying IPv6 Router Advertisements | 27
- Tracing Neighbor Discovery Events | 29

To confirm that the configuration is working properly, perform this task:

Checking the Interfaces

Purpose

Verify that the interfaces are up, and view the assigned EUI-64 addresses.

Action

From operational mode, enter the **show interfaces terse** command.

```
user@R1> show interfaces terse
```

```
Interface               Admin Link Proto    Local                 Remote
fe-1/2/0
fe-1/2/0.1              up    up   inet6    2001:db8:0:1:2a0:a514:0:14c/64
                                            fe80::2a0:a514:0:14c/64
fe-1/2/1.5              up    up   inet6    2001:db8:0:5:2a0:a514:0:54c/64
                                            fe80::2a0:a514:0:54c/64
fe-1/2/2.9              up    up   inet6    2001:db8:0:9:2a0:a514:0:94c/64
                                            fe80::2a0:a514:0:94c/64
lo0
lo0.1                   up    up   inet6    2001:db8::1
                                            fe80::2a0:a50f:fc56:14c
```

Meaning

The output shows that all interfaces are configured with the IPv6 (inet6) address family. Each IPv6-enabled interface has two IPv6 addresses; one link-local address, and one global address. The global addresses match those shown in Figure 1 on page 21. Junos OS automatically creates a link-local address for any interface that is enabled for IPv6 operation. All link-local addresses begin with the fe80::/64 prefix. The

host portion of the address is a full 64 bits long and matches the link-local interface identifier. When an interface address is configured using the **eui-64** statement, its interface identifier matches the interface identifier of the link-local address. This is because link-local addresses are coded according to the EUI-64 specification.

Pinging the Interfaces

Purpose

Verify connectivity between the directly connected interfaces.

Action

1. Determine the remote router's IPv6 interface address.

   On Device R2, run the **show interfaces terse** command for the interface that is directly connected to Device R1, and copy the global address into the capture buffer of your terminal emulator.

   `user@R2>` **show interfaces fe-1/2/0.2 terse**

```
Interface               Admin Link Proto    Local                 Remote
fe-1/2/0.2              up    up   inet6    2001:db8:0:1:2a0:a514:0:24c/64
                                            fe80::2a0:a514:0:24c/64
```

2. On Device R1, run the **ping** command, using the global address that you copied.

   `user@R1>` **ping 2001:db8:0:1:2a0:a514:0:24c**

```
PING6(56=40+8+8 bytes) 2001:db8:0:1:2a0:a514:0:14c --> 2001:db8:0:1:2a0:a514:0:24c
16 bytes from 2001:db8:0:1:2a0:a514:0:24c, icmp_seq=0 hlim=64 time=20.412 ms
16 bytes from 2001:db8:0:1:2a0:a514:0:24c, icmp_seq=1 hlim=64 time=18.897 ms
16 bytes from 2001:db8:0:1:2a0:a514:0:24c, icmp_seq=2 hlim=64 time=1.389 ms
```

Meaning

Junos OS uses the same ping command for both IPv4 and IPv6 testing. The lack of any interior gateway protocol (IGP) in the network limits the ping testing to directly-connected neighbors. Repeat the ping test for other directly connected neighbors.

Checking the IPv6 Neighbor Cache

Purpose

Display information about the IPv6 neighbors.

After conducting ping testing, you can find an entries for interface addresses in the IPv6 neighbor cache.

Action

From operational mode, enter the **show ipv6 neighbors** command.

```
user@R1> show ipv6 neighbors
```

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale         546 yes no
fe-1/2/0.1
fe80::2a0:a514:0:24c         00:05:85:8f:c8:bd  stale         258 yes no
fe-1/2/0.1
fe80::2a0:a514:0:64c         00:05:85:8f:c8:bd  stale         111 yes no
fe-1/2/1.5
fe80::2a0:a514:0:a4c         00:05:85:8f:c8:bd  stale         327 yes no
fe-1/2/2.9
```

Meaning

In IPv6, the Address Resolution Protocol (ARP) has been replaced by the Neighbor Discovery Protocol (NDP). The IPv4 command **show arp** is replaced by the IPv6 command **show ipv6 neighbors**. The key pieces of information displayed by this command are the IP address, the MAC (Link Layer) address, and the interface.

Verifying IPv6 Router Advertisements

Purpose

Confirm that devices can be added to the network using SLAAC by ensuring that router advertisements are working properly.

Action

From operational mode, enter the **show ipv6 router-advertisement** command.

```
user@R1> show ipv6 router-advertisement
```

```
Interface: fe-1/2/0.1
  Advertisements sent: 37, last sent 00:01:41 ago
  Solicits received: 0
  Advertisements received: 38
  Advertisement from fe80::2a0:a514:0:24c, heard 00:05:46 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 1800 sec
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:0:1::/64
```

```
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
Interface: fe-1/2/1.5
  Advertisements sent: 36, last sent 00:05:49 ago
  Solicits received: 0
  Advertisements received: 37
  Advertisement from fe80::2a0:a514:0:64c, heard 00:00:54 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 1800 sec
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:0:5::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
Interface: fe-1/2/2.9
  Advertisements sent: 36, last sent 00:01:37 ago
  Solicits received: 0
  Advertisements received: 38
  Advertisement from fe80::2a0:a514:0:a4c, heard 00:01:00 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 1800 sec
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:0:9::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
```

Meaning

The output shows that router advertisements are being sent and received on Device R1's interfaces, indicating that both Device R1 and its directly connected neighbors are configured to generate router-advertisements.

Tracing Neighbor Discovery Events

Purpose

Perform additional validation by tracing router advertisements.

Action

1. Configure trace operations.

```
[edit protocols router-advertisement traceoptions]
user@R1# set file ipv6-nd-trace
user@R1# set traceoptions flag all
user@R1# commit
```

2. Run the **show log** command.

```
user@R1> show log ipv6-nd-trace
```

```
Mar 29 14:07:16 trace_on: Tracing to "/var/log/P1/ipv6-nd-trace" started
Mar 29 14:07:16.287229 background dispatch running job
ipv6_ra_delete_interface_config_job for task Router-Advertisement
Mar 29 14:07:16.287452 task_job_delete: delete background job
ipv6_ra_delete_interface_config_job for task Router-Advertisement
Mar 29 14:07:16.287505 background dispatch completed job
ipv6_ra_delete_interface_config_job for task Router-Advertisement
Mar 29 14:07:16.288288 ipv6_ra_iflchange(Router-Advertisement): ifl 0xb904378
ifl fe-1/2/2.9 104 change 0, intf 0xba140d8
Mar 29 14:07:16.288450 ipv6_ra_iflchange(Router-Advertisement): ifl 0xb904250
ifl fe-1/2/0.1 85 change 0, intf 0xba14000
Mar 29 14:07:16.288656 ipv6_ra_iflchange(Router-Advertisement): ifl 0xb9044a0
ifl fe-1/2/1.5 80 change 0, intf 0xba1406c
Mar 29 14:07:16.289293 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba002bc
fe80::2a0:a514:0:54c ifl fe-1/2/1.5 80 change 0, intf 0xba1406c
Mar 29 14:07:16.289358 -- nochange/add
Mar 29 14:07:16.289624 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba00230
2001:db8:0:5:2a0:a514:0:54c ifl fe-1/2/1.5 80 change 0, intf 0xba1406c
Mar 29 14:07:16.289682 -- nochange/add
Mar 29 14:07:16.289950 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba001a4
fe80::2a0:a514:0:14c ifl fe-1/2/0.1 85 change 0, intf 0xba14000
Mar 29 14:07:16.290009 -- nochange/add
Mar 29 14:07:16.290302 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba00118
2001:db8:0:1:2a0:a514:0:14c ifl fe-1/2/0.1 85 change 0, intf 0xba14000
Mar 29 14:07:16.290365 -- nochange/add
Mar 29 14:07:16.290634 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba003d4
fe80::2a0:a514:0:94c ifl fe-1/2/2.9 104 change 0, intf 0xba140d8
Mar 29 14:07:16.290694 -- nochange/add
```

```
Mar 29 14:07:16.290958 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba00348
2001:db8:0:9:2a0:a514:0:94c ifl fe-1/2/2.9 104 change 0, intf 0xba140d8
Mar 29 14:07:16.291017 -- nochange/add
Mar 29 14:07:20.808516 task_job_create_foreground: create job ipv6 ra for task
 Router-Advertisement
Mar 29 14:07:20.808921 foreground dispatch running job ipv6 ra for task
Router-Advertisement
Mar 29 14:07:20.809027 ipv6_ra_send_advertisement: sending advertisement for
ifl 104 to ff02::1
Mar 29 14:07:20.809087 (4810916) sending advertisement for ifl 104
Mar 29 14:07:20.809170  ifa 0xba00348 2001:db8:0:9:2a0:a514:0:94c/64
Mar 29 14:07:20.809539  --> sent 56 bytes
Mar 29 14:07:20.809660 task_timer_reset: reset Router-Advertisement_ipv6ra
Mar 29 14:07:20.809725 task_timer_set_oneshot_latest: timer
Router-Advertisement_ipv6ra interval set to 7:07
Mar 29 14:07:20.809772 foreground dispatch completed job ipv6 ra for task
Router-Advertisement
```

RELATED DOCUMENTATION

- *Supported IPv4, TCP, and UDP Standards*
- *Supported IPv6 Standards*
- *Accessing Standards Documents on the Internet*

Secure IPv6 Neighbor Discovery

SUMMARY

The Secure Neighbor Discovery (SEND) Protocol for IPv6 traffic prevents an attacker who has access to the broadcast segment from abusing NDP or ARP to trick hosts into sending the attacker traffic destined for someone else, a technique known as ARP poisoning.

IN THIS SECTION

- Understanding Secure IPv6 Neighbor Discovery | **31**
- Example: Configuring Secure IPv6 Neighbor Discovery | **31**

Understanding Secure IPv6 Neighbor Discovery

One of the functions of the IPv6 Neighbor Discovery Protocol (NDP) is to resolve network layer (IP) addresses to link layer (for example, Ethernet) addresses, a function performed in IPv4 by Address Resolution Protocol (ARP). The Secure Neighbor Discovery (SEND) Protocol prevents an attacker who has access to the broadcast segment from abusing NDP or ARP to trick hosts into sending the attacker traffic destined for someone else, a technique known as ARP poisoning.

To protect against ARP poisoning and other attacks against NDP functions, SEND should be deployed where preventing access to the broadcast segment might not be possible.

SEND uses RSA key pairs to produce cryptographically generated addresses, as defined in RFC 3972, *Cryptographically Generated Addresses (CGA)*. This ensures that the claimed source of an NDP message is the owner of the claimed address.

Example: Configuring Secure IPv6 Neighbor Discovery

IN THIS SECTION

- Requirements | **31**
- Overview | **32**
- Configuration | **32**
- Verification | **34**

This example shows how to configure IPv6 Secure Neighbor Discovery (SEND).

Requirements

This example has the following requirements:

- Junos OS Release 9.3 or later
- IPv6 deployed in your network
- If you have not already done so, you must generate or install an RSA key pair.

  To generate a new RSA key pair, enter the following command:

  `user@host>` **request security pki generate-key-pair type rsa certificate-id *certificate-id-name* size *size***

Overview

To configure SEND, include the following statements:

```
protocols {
    neighbor-discovery {
        onlink-subnet-only;
        secure {
            security-level {
                (default | secure-messages-only);
            }
            cryptographic-address {
                key-length number;
                key-pair pathname;
            }
            timestamp {
                clock-drift number;
                known-peer-window seconds;
                new-peer-window seconds;
            }
            traceoptions {
                file filename <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
                flag flag;
                no-remote-trace;
            }
        }
    }
}
```

Specify **default** to send and receive both secure and unsecured Neighbor Discovery Protocol (NDP) packets. To configure SEND to accept secured NDP messages only and to drop unsecured ones. specify **secure-messages-only**.

All nodes on the segment need to be configured with SEND if the **secure-messages-only** option is used, which is recommended unless only a small subset of devices require increased protection. Failure to configure SEND for all nodes might result in loss of connectivity.

Configuration

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 at the **[edit]** hierarchy level.

```
set protocols neighbor-discovery secure security-level secure-messages-only
set protocols neighbor-discovery secure cryptographic-address key-length 1024
set protocols neighbor-discovery secure cryptographic-address key-pair /var/etc/rsa_key
set protocols neighbor-discovery secure timestamp
```

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* in the *CLI User Guide*.

To configure a secure IPv6 neighbor discovery:

1. Configure the security level.

   ```
   [edit protocols neighbor-discovery secure]
   user@host# set security-level secure-messages-only
   ```

2. (Optional) Enable the key length.

   The default key length is 1024.

   ```
   [edit protocols neighbor-discovery secure]
   user@host# set cryptographic-address key-length 1024
   ```

3. (Optional) Specify the directory path of the public-private key file generated for the cryptographic address.

   The default location of the file is the **/var/etc/rsa_key** directory.

   ```
   [edit protocols neighbor-discovery secure]
   user@host# set cryptographic-address key-pair /var/etc/rsa_key
   ```

4. (Optional) Configure a timestamp to ensure that solicitation and redirect messages are not being replayed.

   ```
   [edit protocols neighbor-discovery secure]
   user@host# set timestamp
   ```

Results

From configuration mode, confirm your configuration by entering the **show protocols** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show protocols
neighbor-discovery {
    secure {
        security-level {
            secure-messages-only;
        }
        cryptographic-address {
            key-length 1024;
            key-pair /var/etc/rsa_key;
        }
        timestamp;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

IN THIS SECTION

- Checking the IPv6 Neighbor Cache | **34**
- Tracing Neighbor Discovery Events | **35**

Confirm that the configuration is working properly.

Checking the IPv6 Neighbor Cache

Purpose

Display information about the IPv6 neighbors.

Action

From operational mode, enter the **show ipv6 neighbors** command.

Meaning

In IPv6, the Address Resolution Protocol (ARP) has been replaced by the NDP. The IPv4 command **show arp** is replaced by the IPv6 command **show ipv6 neighbors**. The key pieces of information displayed by this command are the IP address, the MAC (Link Layer) address, and the interface.

Tracing Neighbor Discovery Events

Purpose

Perform additional validation by tracing SEND.

Action

1. Configure trace operations.

```
[edit protocols neighbor-discovery secure]
user@host# set traceoptions file send-log
user@host# set traceoptions flag all
```

2. Run the **show log** command.

```
user@host> show log send-log
Apr 11 06:21:26 proto: outgoing pkt on idx 68 does not have CGA
(fe80::2a0:a514:0:14c), dropping pkt
Apr 11 06:26:44 proto: sendd_msg_handler: recv outgoing 96 bytes on idx 70 with
 offset 40
Apr 11 06:26:44 dbg: sendd_proto_handler: Modifier (16)
        00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00


Apr 11 06:26:44 cga: snd_is_lcl_cga: BEFORE overriding cc, cc:0, ws->col:0
Apr 11 06:26:44 proto: outgoing pkt on idx 70 does not have CGA
(fe80::2a0:a514:0:24c), dropping pkt
Apr 11 06:26:47 proto: sendd_msg_handler: recv outgoing 96 bytes on idx 68 with
 offset 40
Apr 11 06:26:47 dbg: sendd_proto_handler: Modifier (16)
        00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00
```

Meaning

The output shows that because the packet does not have a cryptographically generated address, the packet is dropped.

NDP Proxy and DAD Proxy

SUMMARY

The Neighbor Discovery Protocol (NDP) proxy functionality enables packet forwarding among the hosts that are in the same subnet and are restricted from communicating directly with each other. The Duplicate Address Detection (DAD) proxy functionality enables a device to respond to DAD queries for a node that is prevented from communicating directly with other nodes in the same subnet.

IN THIS SECTION

- Configuring NDP Proxy | **36**
- Configuring DAD Proxy | **37**

Configuring NDP Proxy

The Neighbor Discovery Protocol (NDP) proxy functionality enables packet forwarding among the hosts that are in the same subnet and are restricted from communicating directly with each other. This functionality is primarily used in a scenario where the proxy node needs to apply access control and intercept traffic flowing among the hosts. When you configure NDP proxy in an SRX Series device, the device sends Neighbor Discovery (ND) advertisements and responds to ND solicitation requests from devices seeking MAC addresses of IPv6 prefixes assigned to hosts inside the SRX Series device.

To configure neighbor discovery proxy on an interface:

1. Set ndp proxy restricted to an interface.

```
set interfaces interface-name family inet6 ndp-proxy interface-restricted
```

To disable NDP proxy for an address that is not present in neighbor cache, execute the following command:

```
set protocols neighbor-discovery ndp-proxy no-proxy-on-resolve
```

Use the **show system statistics icmp6** command to get the statistics of events such as NDP proxy requests, NDP proxy conflicts, NDP proxy duplicates, NDP proxy resolve requests and dropped NDP packets.

Configuring DAD Proxy

The Duplicate Address Detection (DAD) feature detects the usage of duplicate addresses on a local link by using Neighbor Solicitation (NS) messages. The DAD feature is intended for IPv6 address and functions similar to gratuitous ARP in IPv4. The DAD proxy functionality enables an SRX Series device to respond to DAD queries for a node that is prevented from communicating directly with other nodes in the same subnet.

To configure DAD proxy on an interface:

1. Set dad proxy restricted to an interface.

```
set interfaces interface-name family inet6 dad-proxy interface-restricted
```

To disable DAD proxy for an address that is not present in a neighbor cache, execute the following command:

```
set protocols neighbor-discovery dad-proxy no-proxy-on-resolve
```

Use the **show system statistics icmp6** command to get the statistics of events such as DAD proxy requests, DAD proxy conflicts, DAD proxy duplicates, DAD proxy resolve requests and dropped DAD packets.

Neighbor Discovery Cache Protection

SUMMARY

NDP Cache Protection enables you to protect the routing engine from certain types of denial-of-service (DoS) attacks in IPv6 deployment scenarios.

IN THIS SECTION

- Neighbor Discovery Cache Protection Overview | **38**
- Configuring Neighbor Discovery Cache Protection | **39**
- Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks | **40**

Neighbor Discovery Cache Protection Overview

Routing Engines can be susceptible to certain denial-of-service (DoS) attacks in IPv6 deployment scenarios. IPv6 subnets in general tend be very large—for example, a **/64** subnet might have a high number of unassigned addresses. The control plane of the Routing Engine performs the address resolution for unknown addresses. An attacker can quickly overwhelm the control plane of the Routing Engine by generating resolution requests for this unassigned address space, resulting in a cache overflow. The attacker relies on both the number of requests generated and the rate at which requests are queued up. Such scenarios can tie up router resources and prevent the Routing Engine from answering valid neighbor solicitations and maintaining existing neighbor cache entries, effectively resulting in a DoS attack for legitimate users.

The strategies for mitigating such DoS attacks are as follows:

- Filter unused address space.
- Minimize the size of subnets.
- Configure discard routes for subnets.
- Enforce limits to the size and rate of resolution for entries in the neighbor discovery cache.

Neighbor discovery cache impact can be minimized by restricting the number of IPv6 neighbors and new unresolved next-hop addresses that can be added to the cache. You can set limits per interface by using the **nd6-max-cache** and the **nd6-new-hold-limit** configuration statements or system-wide by using the **nd-system-cache-limit** configuration statement.

NOTE:

- For small sized platforms such as ACX, EX22XX, EX3200, EX33XX, and SRX, default is 20,000.
- For medium sized platforms such as EX4200, EX45XX, EX4300, EX62XX, QFX, and MX, default is 75,000.
- For rest of the platforms, default is 100,000.

Configuring Neighbor Discovery Cache Protection

Routing Engines can be susceptible to certain types of denial-of-service (DoS) attacks in IPv6 deployment scenarios. IPv6 subnets in general tend be very large; for example, a **/64** subnet might have a high number of unassigned addresses. The control plane of the Routing Engine performs the address resolution for unknown addresses. An attacker can quickly overwhelm the control plane of the Routing Engine by generating resolution requests for this unassigned address space, resulting in a cache overflow. An attacker relies on both the number of requests generated and the rate at which requests are queued up.

The neighbor discovery process is that part of the control plane that implements the Neighbor Discovery Protocol. It is responsible for performing address resolution and maintaining the entries in the neighbor cache. One way to mitigate the DoS attacks is by enforcing limits to the size of the neighbor discovery cache and the rate of resolution of new next-hop entries, and by prioritizing certain categories of neighbor discovery traffic. You can configure limits to the neighbor discovery cache per interface and systemwide.

Before you begin, ensure that you are running Junos OS Release 15.1 or later.

Local limits apply to individual interfaces and are defined for resolved and unresolved entries in the neighbor discovery queue, while global limits apply systemwide.

To configure neighbor discovery cache protection on an interface:

1. Configure IPv6 family for the interface.

   ```
   [edit interfaces interface-name unit unit number family]
   user@host# set inet6
   ```

2. Configure the maximum size of the neighbor discovery cache for the interface.

   ```
   [edit interfaces interface-name unit unit number family inet6]
   user@host# set nd6-max-cache limit
   ```

3. Configure the maximum number of unresolved entries in the neighbor discovery cache that can be attached to the interface.

   ```
   [edit interfaces interface-name unit unit number family inet6]
   user@host# set nd6-new-hold-limit limit
   ```

To verify the configuration, execute the **show interfaces *interface-name*** operational command.

To configure neighbor discovery cache protection systemwide:

- Configure the systemwide limit for the neighbor discovery cache.

```
[edit]
user@host# set system nd-system-cache-limit limit
```

To verify the configured system-wide limits, execute the **show system statistics icmp6** operational command.

NOTE:

- For small sized platforms such as ACX, EX22XX, EX3200, EX33XX, and SRX, default is 20,000.
- For medium sized platforms such as EX4200, EX45XX, EX4300, EX62XX, QFX, and MX, default is 75,000.
- For rest of the platforms, default is 100,000.

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks

IN THIS SECTION

- Requirements | 41
- Overview | 41
- Configuration | 41
- Verification | 43

This example shows how to configure a limit to the number of IPv6 neighbor entries that can be added to the neighbor discovery. Enforcing limits to the number of entries in the cache mitigates denial-of-service (DoS) attacks. The neighbor discovery cache feature supports two types of limits:

- Local—Local limits are configured per interface and are defined for resolved and unresolved entries in the neighbor discovery cache.
- Global—Global limits apply systemwide. 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 systemwide cap on entries for the neighbor discovery cache, including for the loopback interface for the internal routing instance, as well as management interfaces and the public interfaces.

Requirements

This example requires MX Series routers running Junos OS Release 15.1 or later.

Overview

Routing Engines can be susceptible to certain types of DoS attacks in IPv6 deployment scenarios. IPv6 subnets in general tend be very large—for example, a **/64** subnet might have a high number of unassigned addresses, which can be used to perform DoS attacks. The control plane of the Routing Engine performs the address resolution for unknown addresses. An attacker can quickly overwhelm the control plane of the Routing Engine by generating resolution requests for this unassigned address space and overflow the queue. The attacker relies on both the number of requests generated and the rate at which requests are queued up.

The neighbor discovery process is that part of the control plane that implements the Neighbor Discovery Protocol. It is responsible for performing address resolution and maintaining the neighbor cache. One way to mitigate DoS attacks is by enforcing limits ot the neighbor discovery queue limits, which can be done by restricting queue size and the rate of resolution, and by prioritizing certain categories of neighbor discovery traffic.

Configuration

IN THIS SECTION

- Configuring Neighbor Discovery Cache Protection | **42**
- Results | **42**

To configure neighbor discovery cache protection, perform these tasks:

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 copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

```
set interfaces ge-0/3/0 unit 5 family inet6 nd6-max-cache 100
set interfaces ge-0/3/0 unit 5 family inet6 nd6-new-hold-limit 100
```

You can also configure a systemwide limit to the number of IPv6 neighbor entries in the neighbor discovery cache. This limit also includes the loopback interface, management interfaces, and the public interfaces.

```
set system nd-system-cache-limit 100
```

The limit distribution from the **nd-system-cache-limit** statement for different interface types is performed according to certain fixed percentages. When **nd-system-cache-limit** is defined as *X* and the internal routing interface neighbor discovery cache limit is *Y* (default is 200), then:

- Public maximum cache limit, *Z* = 80% of (*X*- *Y*)
- Management interface maximum cache limit (for example, fxp0), *M* = 20% of (*X*-*Y*)

Configuring Neighbor Discovery Cache Protection

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* in the *CLI User Guide*.

To configure neighbor discovery cache protection per interface:

- Configure the **nd6-max-cache** and **nd6-new-hold-limit**.

```
[edit]
user@host# set interfaces ge-0/3/0 unit 5 family inet6 nd6-max-cache 100
user@host# set interfaces ge-0/3/0 unit 5 family inet6 nd6-new-hold-limit 100
```

Results

To confirm neighbor discovery cache protection locally, enter **show interfaces ge-0/3/0** from configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show interfaces ge-0/3/0
unit 5{
   family inet6 {
      nd6-max-cache 100;
      nd6-new-hold-limit 100;
   }
}
```

Verification

IN THIS SECTION

- Verifying Neighbor Discovery Cache Protection Globally | 43
- Verifying Neighbor Discovery Cache Protection Locally | 44

Confirm that the configuration is working properly.

Verifying Neighbor Discovery Cache Protection Globally

Purpose

Verify that the output reflects the systemwide limit for the neighbor discovery cache.

Action

From operational mode, run the **show system statistics icmp6** command.

```
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 unreach
                 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
```

Meaning

The systemwide cap enforced on the neighbor discovery cache entries is **100000**.

Management ND nexthops creation failed as mgt limit reached indicates the drop count for the management interface when the systemwide limit is reached. **Total ND nexthops creation failed as limit reached** indicates failure for management, public, or Internal routing instance interfaces, and **Public ND nexthops creation failed as public limit reached** indicates the drop count for public interfaces when the systemwide limit to the number of entries is reached.

Verifying Neighbor Discovery Cache Protection Locally

Purpose

Verify that the output reflects the configured interface limits.

Action

From operational mode, run the **show interfaces ge-0/3/0** command.

```
user@host> show interfaces ge-0/3/0
```

```
Logical interface ge-0/2/0.8 (Index 348) (SNMP ifIndex 690)
    Flags: Up SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.8 ]  Encapsulation: ENET2
    Input packets : 181628
    Output packets: 79872
    Protocol inet6, MTU: 1500
    Max nh cache: 100000, New hold nh limit: 100000, Curr nh cnt: 79840, Curr new hold cnt:
 0, NH drop cnt: 0
      Flags: Is-Primary
      Addresses, Flags: Is-Preferred Is-Primary
        Destination: 8001:1::/64, Local: 8001:1::1:1
      Addresses, Flags: Is-Preferred
        Destination: fe80::/64, Local: fe80::56e0:3200:8c6:e0a4
    Protocol multiservice, MTU: Unlimited
```

Meaning

The maximum number of total entries and the maximum number of entries for new unresolved next-hop addresses that can be attached to interface ge-0/3/0 is **100000**.

NH drop cnt refers to the number of neighbor discovery requests not serviced because the interface maximum queue size limits have been reached.

RELATED DOCUMENTATION

- IPv6 Neighbor Discovery Overview | **15**
- nd-system-cache-limit | **80**
- nd6-max-cache | **81**
- nd6-new-hold-limit | **82**

2
CHAPTER

Troubleshooting

Troubleshooting Network Issues | 47

Troubleshooting Network Issues

IN THIS SECTION

- Working with Problems on Your Network | 47
- Isolating a Broken Network Connection | 48
- Identifying the Symptoms of a Broken Network Connection | 49
- Isolating the Causes of a Network Problem | 51
- Taking Appropriate Action for Resolving the Network Problem | 52
- Evaluating the Solution to Check Whether the Network Problem Is Resolved | 53
- Checklist for Tracking Error Conditions | 54
- Configure Routing Protocol Process Tracing | 57
- Configure Routing Protocol Tracing for a Specific Routing Protocol | 60
- Monitor Trace File Messages Written in Near-Real Time | 62
- Stop Trace File Monitoring | 63

Working with Problems on Your Network

Problem

Description: This checklist provides links to troubleshooting basics, an example network, and includes a summary of the commands you might use to diagnose problems with the router and network.

Solution

Table 3: Checklist for Working with Problems on Your Network

Tasks	Command or Action
"Isolating a Broken Network Connection" on page 48	
1. Identifying the Symptoms of a Broken Network Connection on page 49	**ping (*ip-address* | *hostname*)** **show route (*ip-address* | *hostname*)** **traceroute (*ip-address* | *hostname*)**
2. Isolating the Causes of a Network Problem on page 51	**show < configuration | interfaces | protocols | route >**

Table 3: Checklist for Working with Problems on Your Network *(continued)*

Tasks	Command or Action
3. Taking Appropriate Action for Resolving the Network Problem on page 52	**[edit]** **delete routing options static route *destination-prefix*** **commit and-quit** **show route *destination-prefix***
4. Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 53	**show route (*ip-address* | *hostname*)** **ping (*ip-address* | *hostname*) count 3** **traceroute (*ip-address* | *hostname*)**

Isolating a Broken Network Connection

By applying the standard four-step process illustrated in Figure 2 on page 48, you can isolate a failed node in the network. Note that the functionality described in this section is not supported in versions 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Figure 2: Process for Diagnosing Problems in Your Network

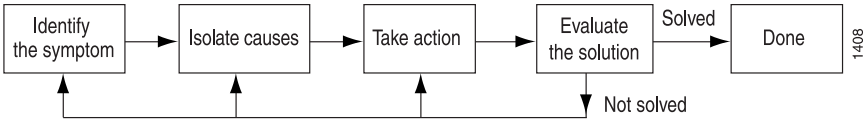

Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic in your approach to the maintenance and monitoring of your network. To prepare for problems on your network, understand how the network functions under normal conditions, have records of baseline network activity, and carefully observe the behavior of your network during a problem situation.

Figure 3 on page 49 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 3: Network with a Problem

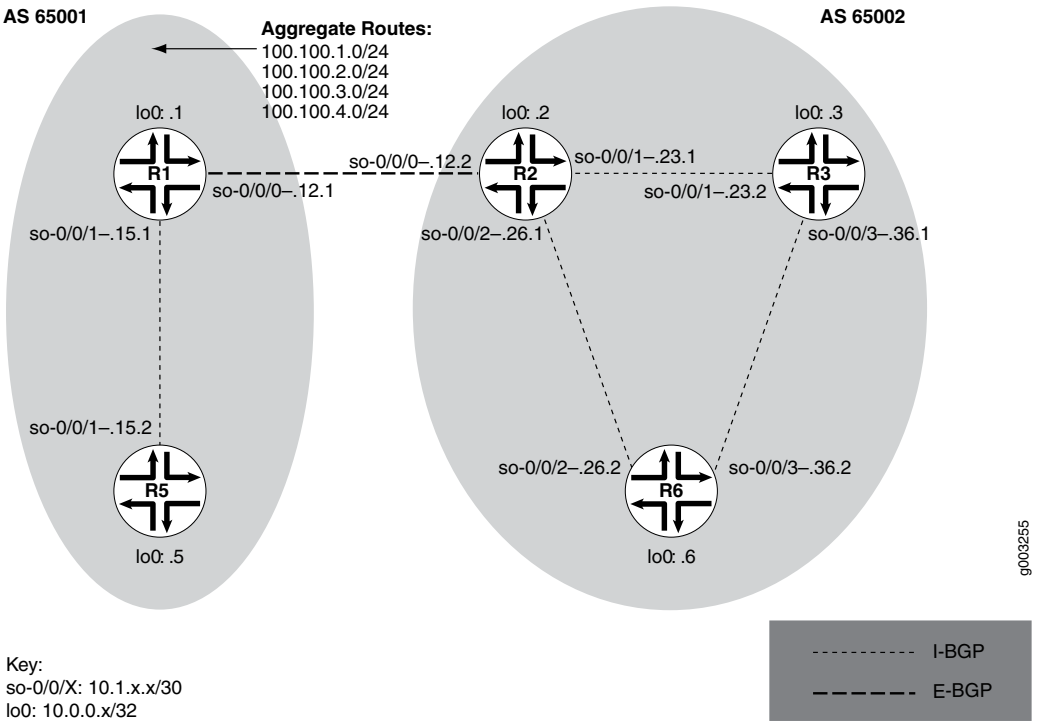

The network in Figure 3 on page 49 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (**R1**) in AS 65001 announces aggregated prefixes **100.100/24** to the AS 65002 network. The problem in this network is that **R6** does not have access to **R5** because of a loop between **R2** and **R6**.

To isolate a failed connection in your network, follow the steps in these topics:

- Isolating the Causes of a Network Problem on page 51
- Taking Appropriate Action for Resolving the Network Problem on page 52
- Taking Appropriate Action for Resolving the Network Problem on page 52
- Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 53

Identifying the Symptoms of a Broken Network Connection

Problem

Description: The symptoms of a problem in your network are usually quite obvious, such as the failure to reach a remote host.

Solution

To identify the symptoms of a problem on your network, start at one end of your network and follow the routes to the other end, entering all or one of the following Junos OS command-line interfaces (CLI) operational mode commands:

```
user@host> ping (ip-address | host-name)
user@host> show route (ip-address | host-name)
user@host> traceroute (ip-address | host-name)
```

Sample Output

```
user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2db   0 0000  01  01 a8c6 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2de   0 0000  01  01 a8c3 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2e2   0 0000  01  01 a8bf 10.1.26.2  10.0.0.5

^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[IS-IS/165] 00:02:39, metric 10
                     > to 10.1.26.1  via so-0/0/2.0

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.649 ms  0.521 ms  0.490 ms
 2  10.1.26.2 (10.1.26.2)  0.521 ms  0.537 ms  0.507 ms
 3  10.1.26.1 (10.1.26.1)  0.523 ms  0.536 ms  0.514 ms
 4  10.1.26.2 (10.1.26.2)  0.528 ms  0.551 ms  0.523 ms
 5  10.1.26.1 (10.1.26.1)  0.531 ms  0.550 ms  0.524 ms
```

Meaning

The sample output shows an unsuccessful **ping** command in which the packets are being rejected because the time to live is exceeded. The output for the **show route** command shows the interface (**10.1.26.1**) that you can examine further for possible problems. The **traceroute** command shows the loop between **10.1.26.1** (**R2**) and **10.1.26.2** (**R6**), as indicated by the continuous repetition of the two interface addresses.

Isolating the Causes of a Network Problem

Problem

Description: A particular symptom can be the result of one or more causes. Narrow down the focus of your search to find each individual cause of the unwanted behavior.

Solution

To isolate the cause of a particular problem, enter one or all of the following Junos OS CLI operational mode command:

```
user@host> show < configuration | bgp | interfaces | isis | ospf | route >
```

Your particular problem may require the use of more than just the commands listed above. See the appropriate command reference for a more exhaustive list of commonly used operational mode commands.

Sample Output

```
user@R6> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.56.2/30
                                  iso
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.26.2/30
                                  iso
so-0/0/3                up    up
so-0/0/3.0              up    up   inet  10.1.36.2/30
                                  iso
[...Output truncated...]
```

The following sample output is from **R2**:

```
user@R2> show route 10.0.0.5
```

```
inet.0: 22 destinations, 25 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[Static/5] 00:16:21
                     > to 10.1.26.2 via so-0/0/2.0
                     [BGP/170] 3d 20:23:35, MED 5, localpref 100
                       AS path: 65001 I
                     > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows that all interfaces on **R6** are up. The output from **R2** shows that a static route [**Static/5**] configured on **R2** points to **R6** (**10.1.26.2**) and is the preferred route to **R5** because of its low preference value. However, the route is looping from **R2** to **R6**, as indicated by the missing reference to **R5** (**10.1.15.2**).

Taking Appropriate Action for Resolving the Network Problem

Problem

Description: The appropriate action depends on the type of problem you have isolated. In this example, a static route configured on **R2** is deleted from the [**routing-options**] hierarchy level. Other appropriate actions might include the following:

Solution

- Check the local router's configuration and edit it if appropriate.
- Troubleshoot the intermediate router.
- Check the remote host configuration and edit it if appropriate.
- Troubleshoot routing protocols.
- Identify additional possible causes.

To resolve the problem in this example, enter the following Junos OS CLI commands:

```
[edit]
user@R2# delete routing-options static route destination-prefix
user@R2# commit and-quit
user@R2# show route destination-prefix
```

Sample Output

```
[edit]
user@R2# delete routing-options static route 10.0.0.5/32

[edit]
user@R2# commit and-quit
commit complete
Exiting configuration mode

user@R2> show route 10.0.0.5

inet.0: 22 destinations, 24 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[BGP/170] 3d 20:26:17, MED 5, localpref 100
                        AS path: 65001 I
                     > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows the static route deleted from the [**routing-options**] hierarchy and the new configuration committed. The output for the **show route** command now shows the BGP route as the preferred route, as indicated by the asterisk (*).

Evaluating the Solution to Check Whether the Network Problem Is Resolved

Problem

Description: If the problem is solved, you are finished. If the problem remains or a new problem is identified, start the process over again.

You can address possible causes in any order. In relation to the network in "Isolating a Broken Network Connection" on page 48, we chose to work from the local router toward the remote router, but you might start at a different point, particularly if you have reason to believe that the problem is related to a known issue, such as a recent change in configuration.

Solution

To evaluate the solution, enter the following Junos OS CLI commands:

```
user@host> show route (ip-address  |host-name)
user@host> ping (ip-address | host-name)
```

```
user@host> traceroute (ip-address | host-name)
```

Sample Output

```
user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[BGP/170]  00:01:35, MED 5, localpref 100, from 10.0.0.2
                      AS path: 65001 I
                    > to 10.1.26.1 via so-0/0/2.0

user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=0 ttl=253 time=0.866 ms
64 bytes from 10.0.0.5: icmp_seq=1 ttl=253 time=0.837 ms
64 bytes from 10.0.0.5: icmp_seq=2 ttl=253 time=0.796 ms
^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.796/0.833/0.866/0.029 ms

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.629 ms  0.538 ms  0.497 ms
 2  10.1.12.1 (10.1.12.1)  0.534 ms  0.538 ms  0.510 ms
 3  10.0.0.5 (10.0.0.5)  0.776 ms  0.705 ms  0.672 ms
```

Meaning

The sample output shows that there is now a connection between **R6** and **R5.** The **show route** command shows that the BGP route to **R5** is preferred, as indicated by the asterisk (*). The **ping** command is successful and the **traceroute** command shows that the path from **R6** to **R5** is through **R2** (**10.1.26.1**), and then through **R1** (**10.1.12.1**).

Checklist for Tracking Error Conditions

Problem

Description: Table 4 on page 55 provides links and commands for configuring routing protocol daemon tracing, Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS) protocol, and Open Shortest Path First (OSPF) protocol tracing to diagnose error conditions.

Solution

Table 4: Checklist for Tracking Error Conditions

Tasks	Command or Action
Configure Routing Protocol Process Tracing	
1. Configure Routing Protocol Process Tracing on page 57	**[edit]** **edit routing-options traceoptions** **set file *filename* size *size* files *number*** **show** **commit** **run show log *filename***
2. Configure Routing Protocol Tracing for a Specific Routing Protocol on page 60	**[edit]** **edit protocol *protocol-name* traceoption** **set file *filename* size *size* files *number*** **show** **commit** **run show log *filename***
3. Monitor Trace File Messages Written in Near-Real Time on page 62	**monitor start *filename***
4. Stop Trace File Monitoring on page 63	**monitor stop *filename***
Configure BGP-Specific Options	
1. *Display Detailed BGP Protocol Information*	**[edit]** **edit protocol bgp traceoptions** **set flag update detail** **show** **commit** **run show log *filename***
2. *Display Sent or Received BGP Packets*	**[edit]** **edit protocol bgp traceoptions** **set flag update (send | receive)** **show** **commit** **run show log *filename***

Table 4: Checklist for Tracking Error Conditions *(continued)*

Tasks	Command or Action
3. *Diagnose BGP Session Establishment Problems*	**[edit]** **edit protocol bgp** **set traceoptions flag open detail** **show** **commit** **run show log *filename***
Configure IS-IS-Specific Options	
1. *Displaying Detailed IS-IS Protocol Information*	**[edit]** **edit protocol isis traceoptions** **set flag hello detail** **show** **commit** **run show log *filename***
2. *Displaying Sent or Received IS-IS Protocol Packets*	**[edit]** **edit protocols isis traceoptions** **set flag hello (send | receive)** **show** **commit** **run show log *filename***
3. *Analyzing IS-IS Link-State PDUs in Detail*	**[edit]** **edit protocols isis traceoptions** **set flag lsp detail** **show** **commit** **run show log *filename***
Configure OSPF-Specific Options	
1. *Diagnose OSPF Session Establishment Problems*	**[edit]** **edit protocols ospf traceoptions** **set flag hello detail** **show** **commit** **run show log *filename***

Table 4: Checklist for Tracking Error Conditions *(continued)*

Tasks	Command or Action
2. *Analyze OSPF Link-State Advertisement Packets in Detail*	**[edit]** **edit protocols ospf traceoptions** **set flag lsa update detail** **show** **commit** **run show log *filename***

Configure Routing Protocol Process Tracing

Action

To configure routing protocol process (rpd) tracing, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit routing-options traceoptions
```

2. Configure the file, file size, number, and flags:

```
[edit routing-options traceoptions]
user@host#  set file filename size size file number
[edit routing-options traceoptions]
user@host#  set flag flag
```

For example:

```
[edit routing-options traceoptions]
user@host# set file daemonlog size 10240 files 10
[edit routing-options traceoptions]
user@host# set flag general
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit routing-options traceoptions]
user@host# show
file daemonlog size 10k files 10;
flag general;
```

4. Commit the configuration:

```
user@host# commit
```

> **NOTE:** Some traceoptions flags generate an extensive amount of information. Tracing can also slow down the operation of routing protocols. Delete the traceoptions configuration if you no longer require it.

1. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
[edit routing-options traceoptions]
user@pro4-a# run show log daemonlog
Sep 17 14:17:31 trace_on: Tracing to "/var/log/daemonlog" started
Sep 17 14:17:31 Tracing flags enabled: general
Sep 17 14:17:31 inet_routerid_notify: Router ID: 10.255.245.44
Sep 17 14:17:31 inet_routerid_notify: No Router ID assigned
Sep 17 14:17:31 Initializing LSI globals
Sep 17 14:17:31 LSI initialization complete
Sep 17 14:17:31 Initializing OSPF instances
Sep 17 14:17:31 Reinitializing OSPFv2 instance master
Sep 17 14:17:31 OSPFv2 instance master running
[...Output truncated...]
```

Meaning

Table 5 on page 59 lists tracing flags and example output for Junos-supported routing protocol daemon tracing.

Table 5: Routing Protocol Daemon Tracing Flags

Tracing Flag	Description	Example Output
all	All operations	Not available.
general	Normal operations and routing table change	Not available.
normal	Normal operations	Not available.
policy	Policy operations and actions	**Nov 29 22:19:58 export: Dest 10.0.0.0 proto Static** **Nov 29 22:19:58 policy_match_qual_or: Qualifier proto Sense: 0** **Nov 29 22:19:58 policy_match_qual_or: Qualifier proto Sense: 0** **Nov 29 22:19:58 export: Dest 10.10.10.0 proto IS-IS**
route	Routing table changes	**Nov 29 22:23:59** **Nov 29 22:23:59 rtlist_walker_job: rt_list walk for RIB inet.0 started with 42 entries** **Nov 29 22:23:59 rt_flash_update_callback: flash KRT (inet.0) start** **Nov 29 22:23:59 rt_flash_update_callback: flash KRT (inet.0) done** **Nov 29 22:23:59 rtlist_walker_job: rt_list walk for inet.0 ended with 42 entries** **Nov 29 22:23:59** **Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 CHANGE route/user af 2 addr 172.16.0.0 nhop-type unicast nhop 10.10.10.33** **Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 ADD route/user af 2 addr 172.17.0.0 nhop-type unicast nhop 10.10.10.33** **Nov 29 22:23:59 KRT Request: send len 68 v14 seq 0 ADD route/user af 2 addr 10.149.3.0 nhop-type unicast nhop 10.10.10.33** **Nov 29 22:24:19 trace_on: Tracing to "/var/log/rpdlog" started** **Nov 29 22:24:19 KRT Request: send len 68 v14 seq 0 DELETE route/user af 2 addr 10.10.218.0 nhop-type unicast nhop 10.10.10.29** **Nov 29 22:24:19 RELEASE 10.10.218.0 255.255.255.0 gw 10.10.10.29,10.10.10.33 BGP pref 170/-101 metric so-1/1/0.0,so-1/1/1.0 <Release Delete Int Ext> as 65401** **Nov 29 22:24:19 KRT Request: send len 68 v14 seq 0 DELETE route/user af 2 addr 172.18.0.0 nhop-type unicast nhop 10.10.10.33**
state	State transitions	Not available.

Table 5: Routing Protocol Daemon Tracing Flags *(continued)*

Tracing Flag	Description	Example Output
task	Interface transactions and processing	**Nov 29 22:50:04 foreground dispatch running job task_collect for task Scheduler** **Nov 29 22:50:04 task_collect_job: freeing task MGMT_Listen (DELETED)** **Nov 29 22:50:04 foreground dispatch completed job task_collect for task Scheduler** **Nov 29 22:50:04 background dispatch running job rt_static_update for task RT** **Nov 29 22:50:04 task_job_delete: delete background job rt_static_update for task RT** **Nov 29 22:50:04 background dispatch completed job rt_static_update for task RT** **Nov 29 22:50:04 background dispatch running job Flash update for task RT** **Nov 29 22:50:04 background dispatch returned job Flash update for task RT** **Nov 29 22:50:04 background dispatch running job Flash update for task RT** **Nov 29 22:50:04 task_job_delete: delete background job Flash update for task RT** **Nov 29 22:50:04 background dispatch completed job Flash update for task RT** **Nov 29 22:50:04 background dispatch running job Flash update for task RT** **Nov 29 22:50:04 task_job_delete: delete background job Flash update for task RT**
timer	Timer usage	**Nov 29 22:52:07 task_timer_hiprio_dispatch: ran 1 timer** **Nov 29 22:52:07 main: running normal priority timer queue** **Nov 29 22:52:07 main: ran 1 timer** **Nov 29 22:52:07 task_timer_hiprio_dispatch: running high priority timer queue** **Nov 29 22:52:07 task_timer_hiprio_dispatch: ran 1 timer** **Nov 29 22:52:07 main: running normal priority timer queue** **Nov 29 22:52:07 main: ran 1 timer** **Nov 29 22:52:07 main: running normal priority timer queue** **Nov 29 22:52:07 main: ran 2 timers**

Configure Routing Protocol Tracing for a Specific Routing Protocol

Action

To configure routing protocol tracing for a specific routing protocol, follow these steps:

1. In configuration mode, go to the following hierarchy level:

```
[edit]
user@host# edit protocol protocol-name traceoptions
```

2. Configure the file, file size, number, and flags:

```
[edit protocols protocol name traceoptions]
```

```
user@host# set file filename size size files number
[edit protocols protocol name traceoptions]
user@host# set flag flag
```

For example:

```
[edit protocols ospf traceoptions]
user@host# set file ospflog size 10240 files 10
[edit protocols ospf traceoptions]
user@host# set flag general
```

3. Verify the configuration:

```
user@host# show
```

For example:

```
[edit protocols ospf traceoptions]
user@host# show
file ospflog size 10k files 10;
flag general;
```

4. Commit the configuration:

```
user@host# commit
```

5. View the contents of the file containing the detailed messages:

```
user@host# run show log filename
```

For example:

```
[edit protocols ospf traceoptions]
user@pro4-a# run show log ospflog
Sep 17 14:23:10 trace_on: Tracing to "/var/log/ospflog" started
Sep 17 14:23:10 rt_flash_update_callback: flash OSPF (inet.0) start
Sep 17 14:23:10 OSPF: multicast address 224.0.0.5/32, route ignored
Sep 17 14:23:10 rt_flash_update_callback: flash OSPF (inet.0) done
Sep 17 14:23:10 CHANGE 10.255.245.46/32 gw 10.10.208.67 OSPF pref 10/0 metric 1/0 fe-0/0/0.0 <Delete
   Int>
Sep 17 14:23:10 CHANGE 10.255.245.46/32 gw 10.10.208.67 OSPF pref 10/0 metric 1/0 fe-0/0/0.0 <Active
   Int>
```

```
Sep 17 14:23:10 ADD 10.255.245.46/32 gw 10.10.208.67 OSPF pref 10/0 metric 1/0 fe-0/0/0.0 <Active
    Int>
Sep 17 14:23:10 CHANGE 10.255.245.48/32 gw 10.10.208.69 OSPF pref 10/0 metric 1/0 fe-0/0/0.0 <Delete
    Int>
Sep 17 14:23:10 CHANGE 10.255.245.48/32 gw 10.10.208.69 OSPF pref 10/0 metric 1/0 fe-0/0/0.0 <Active
    Int>
Sep 17 14:23:10 ADD 10.255.245.48/32 gw 10.10.208.69 OSPF pref 10/0 metric 1/0 fe-0/0/0.0 <Active
    Int>
Sep 17 14:23:10 rt_close: 4/4 routes proto OSPF
[...Output truncated...]
```

Meaning

Table 6 on page 62 lists standard tracing options that are available globally or that can be applied to specific protocols. You can also configure tracing for a specific BGP peer or peer group. For more information, see the *Junos System Basics Configuration Guide*.

Table 6: Standard Trace Options for Routing Protocols

Tracing Flag	Description
all	All operations
general	Normal operations and routing table changes
normal	Normal operations
policy	Policy operations and actions
route	Routing table changes
state	State transitions
task	Interface transactions and processing
timer	Timer usage

Monitor Trace File Messages Written in Near-Real Time

Purpose

To monitor messages in near-real time as they are being written to a trace file.

Action

To monitor messages in near-real time as they are being written to a trace file, use the following Junos OS command-line interface (CLI) operational mode command:

```
user@host> monitor start filename
```

Sample Output

```
user@host> monitor start isis
```

```
user@host>
*** isis ***
Sep 15 18:32:21 Updating LSP isis5.02-00 in database
Sep 15 18:32:21 Updating L2 LSP isis5.02-00 in TED
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Scheduling L2 LSP isis5.02-00 sequence 0xd87 on interface fxp2.3
Sep 15 18:32:21 Updating LSP isis5.00-00 in database
Sep 15 18:32:21 Updating L1 LSP isis5.00-00 in TED
Sep 15 18:32:21 Sending L2 LSP isis5.02-00 on interface fxp2.3
Sep 15 18:32:21     sequence 0xd87, checksum 0xc1c8, lifetime 1200
```

Stop Trace File Monitoring

Action

To stop monitoring a trace file in near-real time, use the following Junos OS CLI operational mode command after you have started monitoring:

```
user@host monitor stop filename
```

Sample Output

```
user@host> monitor start isis
```

```
user@host>
*** isis ***
Sep 15 18:32:21 Updating LSP isis5.02-00 in database
Sep 15 18:32:21 Updating L2 LSP isis5.02-00 in TED
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis6.00
Sep 15 18:32:21 Adding a half link from isis5.02 to isis5.00
Sep 15 18:32:21 Scheduling L2 LSP isis5.02-00 sequence 0xd87 on interface fxp2.3
Sep 15 18:32:21 Updating LSP isis5.00-00 in database
Sep 15 18:32:21 Updating L1 LSP isis5.00-00 in TED
Sep 15 18:32:21 Sending L2 LSP isis5.02-00 on interface fxp2.3
Sep 15 18:32:21     sequence 0xd87, checksum 0xc1c8, lifetime 1200
```

monitor stop isis

```
user@host>
```

3

CHAPTER

Configuration Statements

autonomous | **67**

cryptographic-address | **68**

current-hop-limit | **69**

default-lifetime | **70**

interface (Protocols IPv6 Neighbor Discovery) | **71**

key-length | **73**

key-pair | **74**

link-mtu | **75**

managed-configuration | **76**

max-advertisement-interval (Protocols IPv6 Neighbor Discovery) | **77**

min-advertisement-interval (Protocols IPv6 Neighbor Discovery) | **78**

nd-retransmit-timer | **79**

nd-system-cache-limit | **80**

nd6-max-cache | **81**

nd6-new-hold-limit | **82**

neighbor-discovery | **84**

on-link | **86**

onlink-subnet-only | **87**

other-stateful-configuration | **88**

preference (IPv6 Router Advertisement) | **89**

preferred-lifetime | **90**

prefix (Protocols IPv6 Neighbor Discovery) | **91**

reachable-time | **92**

retransmit-timer | **93**

router-advertisement | **94**

secure | **95**

security-level | **96**

solicit-router-advertisement-unicast | **97**

timestamp | **98**

traceoptions (Protocols IPv6 Neighbor Discovery) | **100**

traceoptions (Protocols Secure Neighbor Discovery) | **102**

valid-lifetime | **104**

autonomous

Syntax

```
(autonomous | no-autonomous);
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols router-advertisement interface *interface-name* **prefix** *prefix*],
[edit protocols router-advertisement interface *interface-name* **prefix** *prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify whether prefixes in the router advertisement messages are used for stateless address autoconfiguration:

- **autonomous**—Use prefixes for address autoconfiguration.
- **no-autonomous**—Do not use prefixes for address autoconfiguration.

Default

autonomous

Required Privilege Level

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

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

cryptographic-address

Syntax

```
cryptographic-address {
    key-length number;
    key-pair pathname;
}
```

Hierarchy Level

[edit protocols neighbor-discovery secure]

Release Information
Statement introduced in Junos OS Release 9.3.

Description
Configure parameters for cryptographically generated addresses for Secure Neighbor Discovery.

The Secure Neighbor Discovery (SEND) Protocol uses cryptographically generated addresses (CGAs), as defined in RFC 3972, *Cryptographically Generated Addresses*, to ensure that the sender of a Neighbor Discovery Protocol (NDP) message is the “owner” of the claimed address. Each node must generate a public-private key pair before it can claim an address. The CGA is included in all outgoing neighbor solicitation and neighbor advertisement messages.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
routing level—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring Secure IPv6 Neighbor Discovery | 31
Understanding Secure IPv6 Neighbor Discovery | 31

current-hop-limit

Syntax

```
current-hop-limit number;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Set the default value placed in the hop count field of the IP header for outgoing packets.

Options

number—Hop limit. A value of 0 means the limit is unspecified by this router.

Range: 0 through 255

Default: 64

Required Privilege Level

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

RELATED DOCUMENTATION

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

default-lifetime

Syntax

```
default-lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure the lifetime associated with a default router.

Options

seconds—Default lifetime. A value of 0 means this router is not the default router.

Range: Maximum advertisement interval value through 9000 seconds

Default: Three times the maximum advertisement interval value

Required Privilege Level

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

RELATED DOCUMENTATION

max-advertisement-interval | 77

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

interface (Protocols IPv6 Neighbor Discovery)

Syntax

```
interface interface-name {
    current-hop-limit number;
    default-lifetime seconds;
    (link-mtu | no-link-mtu);
    (managed-configuration | no-managed-configuration);
    max-advertisement-interval seconds;
    min-advertisement-interval seconds;
    (other-stateful-configuration | no-other-stateful-configuration);
    prefix prefix {
        (autonomous | no-autonomous);
        (on-link | no-on-link);
        preferred-lifetime seconds;
        valid-lifetime seconds;
    }
    reachable-time milliseconds;
    retransmit-timer milliseconds;
    solicit-router-advertisement-unicast;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement],
[edit protocols router-advertisement]
```

Release Information

Statement introduced before Junos OS Release 7.4.
solicit-router-advertisement-unicast statement added from 15.1 Release onwards.

Description

Configure router advertisement properties on an interface. To configure more than one interface, include the **interface** statement multiple times.

The Junos OS enters the Neighbor Discovery Protocol (NDP) packets into the routing platform cache even if there is no known route to the source.

If you are using Virtual Router Redundancy Protocol (VRRP) for IPv6, you must include the **virtual-router-only** statement on both the master and backup VRRP on the IPv6 router.

Options

interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery | **30**

key-length

Syntax

```
key-length number;
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure cryptographic-address]
```

Release Information
Statement introduced in Junos OS Release 9.3.

Description
Specify the length of the RSA key used to generate the public-private key pair for the cryptographic address.

Default
1024

Options
number—RSA key length.
Range: 1024 through 2048

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery | 30

key-pair

Syntax

```
key-pair pathname;
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure cryptographic-address]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Specify the directory path of the public-private key file generated for the cryptographic address.

A cryptographic address is dynamically generated based on a public key and a subnet prefix.

Default

The default location of the file is the **/var/etc/rsa_key** directory.

Options

pathname—Directory path of the public-private key file.

Required Privilege Level

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

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery | 30

link-mtu

Syntax

```
(link-mtu | no-link-mtu);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced in Junos OS 10.3.

Description

Specify whether to include the maximum transmission unit (MTU) option in router advertisement messages:

- **link-mtu**–Includes the MTU option in router advertisements.
- **no-link-mtu**–Does not include the MTU option in router advertisements.

The MTU option included in router advertisement messages ensures that all nodes on a link use the same MTU value in situations where the link MTU is not well known.

Default

Router advertisement messages do not include the MTU option.

Required Privilege Level

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

RELATED DOCUMENTATION

> *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

managed-configuration

Syntax

```
(managed-configuration | no-managed-configuration);
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols router-advertisement **interface** *interface-name*],
[edit protocols router-advertisement **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify whether to enable the host to use a stateful autoconfiguration protocol for address autoconfiguration, along with any stateless autoconfiguration already configured:

- **managed-configuration**—Enable host to use stateful autoconfiguration.
- **no-managed-configuration**—Disable host from using stateful autoconfiguration.

You can set two fields in the router advertisement message to enable stateful autoconfiguration on a host: the managed configuration field and the other stateful configuration field. Setting the managed configuration field enables the host to use a stateful autoconfiguration protocol for address autoconfiguration, along with any stateless autoconfiguration already configured. Setting the other stateful configuration field enables autoconfiguration of other nonaddress-related information.

Default

Stateful autoconfiguration is disabled.

Required Privilege Level

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

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

other-stateful-configuration | 88

max-advertisement-interval (Protocols IPv6 Neighbor Discovery)

Syntax

```
max-advertisement-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Set the maximum interval between each router advertisement message.

The router sends router advertisements on each interface configured to transmit messages. The advertisements include route information and indicate to network hosts that the router is operational. The router sends these messages periodically, with a time range defined by minimum and maximum values.

Options

seconds—Maximum interval.

Range: 4 through 1800 seconds

Default: 600 seconds

Required Privilege Level

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

RELATED DOCUMENTATION

min-advertisement-interval | 78

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

min-advertisement-interval (Protocols IPv6 Neighbor Discovery)

Syntax

```
min-advertisement-interval seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Set the minimum interval between each router advertisement message.

The router sends router advertisements on each interface configured to transmit messages. The advertisements include route information and indicate to network hosts that the router is operational. The router sends these messages periodically, with a time range defined by minimum and maximum values.

Options

seconds—Minimum interval.

Range: 3 seconds through three-quarter times the maximum advertisement interval value

Default: One-third the maximum advertisement interval valueBy default, the maximum advertisement interval is 600 seconds and the minimum advertisement interval is one-third the maximum interval, or 200 seconds.

Required Privilege Level

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

RELATED DOCUMENTATION

max-advertisement-interval | 77

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

nd-retransmit-timer

Syntax

```
nd-retransmit-timer milliseconds;
```

Hierarchy Level

```
[edit system]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Set the retransmit timer for neighbor discovery messages. Whenever the state of a neighbor during the Neighbor Discovery (ND) process changes from stale to probe, the value of the retransmit timer controls the interval between the neighbor solicitation messages that are sent out. Also, the retransmit timer controls the time for which the neighbor is in the probe state. A device sends a neighbor solicitation message after the specified number of milliseconds in the nd-retransmit-timer statement, until a reachability confirmation is received. If a solicited neighbor advertisement (NA) message is not received from the neighbor in response to the soliciation message sent from the device, the neighbor remains in the probe state.

Options

milliseconds—Retransmission frequency.

Default: 0 milliseconds

Required Privilege Level

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

RELATED DOCUMENTATION

Using NDRA to Provide IPv6 WAN Link Addressing Overview

nd-system-cache-limit

Syntax

```
nd-system-cache-limit number;
```

Hierarchy Level

```
[edit system]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Specify the maximum system cache size for IPv6 next-hop addresses. This limit enforces a systemwide cap on the neighbor discovery cache entries for all interfaces, including the loopback interface for the internal routing instance, management interfaces, and the public interfaces.

Default

100,000

Options

number—Maximum system cache size for IPv6 next-hop addresses.

Range: 200 through 2,000,000

Required Privilege Level

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

RELATED DOCUMENTATION

- Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks | 40
- nd6-max-cache | 81
- nd6-new-hold-limit | 82

nd6-max-cache

Syntax

```
nd6-max-cache nd6-max-cache;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Specify the maximum number of entries that can be added to the Neighbor Discovery Protocol (NDP) IPv6 neighbor discovery cache for an interface. When this maximum is reached, no new entries are allowed.

Default

- 100,000 for M Series.
- 75,000 for MX Series and QFX Series.
- 20,000 for EX Series.

Options

nd6-max-cache—Maximum size of the neighbor discovery next-hop cache for an interface.

Range: 1 through 2,000,000 for MX Series or QFX Series.

Range: 1 through 700,000 for M Series.

Required Privilege Level

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

RELATED DOCUMENTATION

Configuring Neighbor Discovery Cache Protection | 39

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks | 40

IPv6 Neighbor Discovery Overview | 15

nd6-new-hold-limit | 82

nd-system-cache-limit | 80

nd6-new-hold-limit

Syntax

```
nd6-new-hold-limit nd6-new-hold-limit;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet6]
```

Release Information

Statement introduced in Junos OS Release 15.1.

Description

Specify the maximum number of entries for unresolved next-hop addresses that can be added to the Neighbor Discovery Protocol (NDP) IPv6 neighbor discovery cache for an interface.

Default

- 100,000 for M Series.
- 75,000 for MX Series and QFX Series.
- 20,000 for EX Series.

Options

nd6-new-hold-limit—Maximum number of new unresolved next-hop addresses that can be added to the IPv6 neighbor discovery cache.

Range: 1 through 2,000,000

Required Privilege Level

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

RELATED DOCUMENTATION

Configuring Neighbor Discovery Cache Protection | 39

Example: Configuring Neighbor Discovery Cache Protection to Prevent Denial-of-Service Attacks | 40

IPv6 Neighbor Discovery Overview | 15

nd-system-cache-limit | 80

nd6-max-cache | 81

nd6-stale-time

neighbor-discovery

Syntax

```
neighbor-discovery {
    onlink-subnet-only;
    secure {
        security-level {
            (default | secure-messages-only);
        }
        cryptographic-address {
            key-length number;
            key-pair pathname;
        }
        timestamp {
            clock-drift number;
            known-peer-window number;
            new-peer-window number;
        }
        traceoptions {
            file filename <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
            flag flag;
            no-remote-trace;
        }
    }
}
```

Hierarchy Level

```
[edit protocols]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Enable Secure Neighbor Discovery.

The remaining statements are explained separately. See CLI Explorer.

Default

Disabled

Required Privilege Level

routing—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery | 30

on-link

Syntax

```
(on-link | no-on-link);
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols router-advertisement interface *interface-name* **prefix** *prefix*],
[edit protocols router-advertisement interface *interface-name* **prefix** *prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify whether to enable prefixes to be used for onlink determination:

- **no-on-link**—Disable prefixes from being used for onlink determination.
- **on-link**—Enable prefixes to be used for onlink determination.

Router advertisement messages carry prefixes and information about them. A prefix is onlink when it is assigned to an interface on a specified link. The prefixes specify whether they are onlink or not onlink. A node considers a prefix to be onlink if it is represented by one of the link's prefixes, a neighboring router specifies the address as the target of a redirect message, a neighbor advertisement message is received for the (target) address, or any neighbor discovery message is received from the address. These prefixes are also used for address autoconfiguration. The information about the prefixes specifies the lifetime of the prefixes, whether the prefix is autonomous, and whether the prefix is onlink.

Default

Prefixes are onlink unless explicitly disabled.

Required Privilege Level

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

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

onlink-subnet-only

Syntax

```
onlink-subnet-only;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols neighbor-discovery],
[edit protocols neighbor-discovery]
```

Release Information

Statement introduced in Junos OS Release 10.0.
Statement introduced in Junos OS Release 11.3 for SRX Series devices.

Description

Enable this option to prevent the device from responding to a neighbor solicitation (NS) from a prefix that is not included as one of the device interface prefixes.

After configuring the **onlink-subnet-only** statement, the Routing Engine needs to be restarted using the **request system reboot both-routing-engines** command. If the attacker's IPv6 destination address is already in the forwarding-table, it is not removed after you configure the **onlink-subnet-only** statement, and therefore the device continues to respond to ping NSs. Restarting the Routing Engine removes the entry from the forwarding table.

Default

Disabled

Required Privilege Level

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

RELATED DOCUMENTATION

> *Understanding How to Control Inbound Traffic Based on Protocols*
>
> *IPv6 Neighbor Discovery User Guide*

other-stateful-configuration

Syntax

```
(other-stateful-configuration | no-other-stateful-configuration);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Specify whether to enable autoconfiguration of other nonaddress-related information:

- **no-other-stateful-configuration**—Disable autoconfiguration of other nonaddress-related information.
- **other-stateful-configuration**—Enable autoconfiguration of other nonaddress-related information.

Default
By default, stateful autoconfiguration is disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery | 30

preference (IPv6 Router Advertisement)

Syntax

```
preference (high | low | medium);
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information
Statement introduced in Junos OS Release 16.1 for the MX Series.

Description
Specify the router preference that is communicated to IPv6 hosts through router advertisements. The **preference** value in the router advertisements enables IPv6 hosts to select a default router to reach a remote destination.

The **preference** can be configured when there are multiple devices that route to distinct sets of prefixes and where one of the devices would lead to considerably fewer redirects. You can indicate a lower preference for a new device that is not completely configured yet, so that hosts do not adopt this new device as the default device and thus avoid traffic loss.

Options
You can specify different levels of preference depending on your requirements:

high—Specify a high preference for a device.

low—Specify a low preference for a device.

medium—Specify a medium preference for a device. This is the default preference.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

preferred-lifetime

Syntax

```
preferred-lifetime seconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name prefix prefix],
[edit protocols router-advertisement interface interface-name prefix prefix]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify how long the prefix generated by stateless autoconfiguration remains preferred.

Options

seconds—Preferred lifetime, in seconds. If you set the preferred lifetime to **0xffffffff**, the lifetime is infinite. The preferred lifetime is never greater than the valid lifetime.

Default: 604,800 seconds

Required Privilege Level

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

RELATED DOCUMENTATION

valid-lifetime | **104**

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery | **30**

prefix (Protocols IPv6 Neighbor Discovery)

Syntax

```
prefix prefix {
    (autonomous | no-autonomous);
    (on-link | no-on-link);
    preferred-lifetime seconds;
    valid-lifetime seconds;
}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure prefix properties in router advertisement messages.

Options

prefix—Prefix name.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

RELATED DOCUMENTATION

- *Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery*

reachable-time

Syntax

```
reachable-time milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Set the length of time that a node considers a neighbor reachable until another reachability confirmation is received from that neighbor.

After receiving a reachability confirmation from a neighbor, a node considers that neighbor reachable for a certain amount of time without receiving another confirmation. This mechanism is used for neighbor unreachability detection, a mechanism for finding link failures to a target node.

Options

milliseconds—Reachability time limit.

Range: 0 through 3,600,000 milliseconds

Default: 0 milliseconds

Required Privilege Level

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

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery | 30

retransmit-timer

Syntax

```
retransmit-timer milliseconds;
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols router-advertisement interface interface-name],
[edit protocols router-advertisement interface interface-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Set the retransmission frequency of neighbor solicitation messages. This timer is used to detect when a neighbor has become unreachable and to resolve addresses.

Options

milliseconds—Retransmission frequency.

Default: 0 milliseconds

Required Privilege Level

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

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery | 30

router-advertisement

Syntax

```
router-advertisement {...}
```

Hierarchy Level

```
[edit logical-systems logical-system-name protocols],
[edit protocols]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable router advertisement.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Secure IPv6 Neighbor Discovery | 30

secure

Syntax

```
secure {
    security-level {
        (default | secure-messages-only);
    }
    cryptographic-address {
        key-length number;
        key-pair pathname;
    }
    timestamp {
        clock-drift number;
        known-peer-window seconds;
        new-peer-window seconds;
    }
    traceoptions {
        file filename <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure parameters for Secure Neighbor Discovery.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

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

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery | 30

security-level

Syntax

```
security-level {
    (default | secure-messages-only);
}
```

Hierarchy Level

[edit protocols **neighbor-discovery secure**]

Release Information
Statement introduced in Junos OS Release 9.3.

Description
Configure the type of security mode for Secure Neighbor Discovery.

Options
default—Accept and transmit both secure and unsecured messages.

secure-messages-only—Accept secure messages only. Discard unsecured messages.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Secure IPv6 Neighbor Discovery | 30

solicit-router-advertisement-unicast

Syntax

```
solicit-router-advertisement-unicast;
```

Hierarchy Level

[edit protocols router-advertisement interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 15.1R1 onwards.

Description
Configure devices to send router advertisements as unicast in response to the router solicitation message sent by IPv6 routers.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

- *IPv6 Neighbor Discovery User Guide*

timestamp

Syntax

```
timestamp {
    clock-drift value;
    known-peer-window seconds;
    new-peer-window seconds;
}
```

Hierarchy Level

[edit protocols neighbor-discovery secure]

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure timestamp options, which are used to ensure that solicitation and redirect messages are not being replayed.

Options

clock-drift *value*—Specify the allowable drift in time between the synchronization of peers. For ***value***, specify a fractional value of 100.

Default: 0.01

known-peer-window *seconds*—Specify the expected interval in seconds between Secure Neighbor Discovery messages from an established peer. A message from a known peer that arrives after the specified interval is discarded.

Default: 1 second

new-peer-window *seconds*—Specify the maximum allowable time in seconds between the timestamp of a Secure Neighbor Discovery message from a new peer and when it can be accepted.

Default: 300 seconds

Required Privilege Level

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

RELATED DOCUMENTATION

traceoptions (Protocols IPv6 Neighbor Discovery)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **router-advertisement**],
[edit protocols **router-advertisement**]

Release Information
Statement introduced before Junos OS Release 7.4.

Description
For IPv6 neighbor discovery, specify router advertisement protocol-level tracing options.

Trace IPv6 Neighbor Discovery protocol traffic to help debug Neighbor Discovery protocol issues.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the IPv6 Neighbor Discovery protocol using the **traceoptions flag** statement included at the **[edit protocols router-advertisement]** hierarchy level:

Default
The default trace options are inherited from the global **traceoptions** statement.

Options
disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place router advertisement tracing output in the file **/var/log/router-advertisement-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

- **all**—All tracing operations

  NOTE: Use the trace flag **all** with caution as this may cause the CPU to become very busy.

- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—IPv6 interface transactions and processing
- **timer**—IPv6 neighbor discovery protocol timer processing

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

traceoptions (Protocols Secure Neighbor Discovery)

Syntax

```
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit protocols neighbor-discovery secure]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Configure tracing operations for Secure Neighbor Discovery events. To specify more than one tracing operation, include multiple **flag** statements.

Options

file *filename*—Name of the file to receive the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1** and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

Secure Neighbor Discovery Tracing Options

- **configuration**—All configuration events.
- **cryptographic-address**—Cryptographically generated address events.
- **protocol**—All protocol processing events.
- **rsa**—RSA events.

Global Tracing Options

- **all**—All tracing operations.

You can specify one or more of following flag modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

match *regular-expression*—(Optional) Specify a regular expression to match the output of the trace file you want to log.

no-remote-trace—Disable remote tracing globally or for a specific tracing operation.

no-world-readable—(Optional) Prevent any user from reading this log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1**, and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read this log file.

Required Privilege Level
routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery | 19

Understanding Secure IPv6 Neighbor Discovery | 31

Understanding IPv6 Neighbor Discovery

valid-lifetime

Syntax

valid-lifetime *seconds*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols router-advertisement interface *interface-name* **prefix** *prefix*],
[edit protocols router-advertisement interface *interface-name* **prefix** *prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify how long the prefix remains valid for onlink determination.

Options

seconds—Valid lifetime, in seconds. If you set the valid lifetime to **0xffffffff**, the lifetime is infinite.

Default: 2,592,000 seconds

Required Privilege Level

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

RELATED DOCUMENTATION

preferred-lifetime | **90**

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

4

CHAPTER

Operational Commands

clear ipv6 neighbors | **106**

clear ipv6 router-advertisement | **108**

monitor interface | **109**

monitor start | **125**

monitor stop | **127**

ping | **129**

show ipv6 neighbors | **137**

show ipv6 router-advertisement | **140**

show log | **144**

traceroute | **152**

clear ipv6 neighbors

Syntax

```
clear ipv6 neighbors
<all>
< host hostname>
< interface interface-name>
< logical-system logical-system-name>
< tenant name>
< vpn vpn-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Clear IPv6 neighbor cache information.

> **NOTE:** On Junos OS Evolved, issuing the **clear ipv6 neighbors** command clears the cache for IPv6 neighbors in a reachable state.

Options

none—Clear all IPv6 neighbor cache information.

all—(Optional) Clear all IPv6 neighbor cache information.

host *hostname*—(Optional) Clear the information for the specified IPv6 neighbors.

interface *interface-name*—(Optional) Clear the information about IPv6 for the specified logical interface.

logical-system *logical-system-name*—(Optional) Clear the IPv6 entries for the specified logical system; only available in the main router context.

tenant *name*—(Optional) Clear the IPv6 entries for the specified tenant system; only available in the main router context.

vpn *vpn-name*—(Optional) Clear entries in the IPv6 table for the specified virtual private network's (VPN) routing table.

Required Privilege Level
view

RELATED DOCUMENTATION

- show ipv6 neighbors | **137**

List of Sample Output
clear ipv6 neighbors on page 107

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

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

clear ipv6 router-advertisement

Syntax

```
clear ipv6 router-advertisement
<interface interface>
<logical-system (all | logical-system-name)>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear IPv6 router advertisement counters.

Options

none—Clear IPv6 router advertisement counters for all interfaces.

interface *interface*—(Optional) Clear IPv6 router advertisement counters for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

RELATED DOCUMENTATION

show ipv6 router-advertisement | 140

List of Sample Output

clear ipv6 router-advertisement on page 108

Output Fields

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

Sample Output

clear ipv6 router-advertisement

```
user@host> clear ipv6 router-advertisement
```

monitor interface

Syntax

```
monitor interface
<interface-name | traffic <detail>>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as SONET/SDH and T3 alarms, loopbacks detected, and increases in framing errors.

NOTE: On Junos OS Evolved, you can use the **monitor interface** command over SSH sessions, but console and Telnet sessions are not supported.

NOTE: This command is not supported on the QFX3000 QFabric switch.

Options

none—Display real-time statistics for all interfaces.

detail—(Optional) With traffic option only, display detailed output.

interface-name—(Optional) Display real-time statistics for the specified interface. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified line-card chassis (LCC) only.

traffic—(Optional) Display traffic data for all active interfaces. In a TX Matrix or TX Matrix Plus router, display real-time statistics for the physical interfaces on the specified LCC only.

Additional Information

The output of this command shows how much each field has changed since you started the command or since you cleared the counters by pressing the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type

in the CLI Explorer. To control the output of the **monitor interface** command while it is running, use the keys listed in Table 7 on page 110. The keys are not case-sensitive.

Table 7: Output Control Keys for the monitor interface interface-name Command

Key	Action
c	Clears (returns to zero) the delta counters since **monitor interface** was started. This does not clear the accumulative counter. To clear the accumulative counter, use the **clear interfaces interval** command.
f	Freezes the display, halting the display of updated statistics and delta counters.
i	Displays information about a different interface. The command prompts you for the name of a specific interface.
n	Displays information about the next interface. The **monitor interface** command displays the physical or logical interfaces in the same order as the **show interfaces terse** command.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in Table 8 on page 110. The keys are not case-sensitive.

Table 8: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bits per second (bps).
c	Clears (return to 0) the delta counters in the **Current Delta** column. The statistics counters are not cleared.
d	Displays the **Current Delta** column (instead of the rate column) in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or Esc	Quits the command and returns to the command prompt.
r	Displays the rate column (instead of the **Current Delta** column) in bps and pps.

Required Privilege Level

trace

List of Sample Output

monitor interface (Physical) on page 112
monitor interface (OTN Interface) on page 114
monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface) on page 115
monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface) on page 116
monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface) on page 117
monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface) on page 119
monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface) on page 119
monitor interface (Logical) on page 120
monitor interface (QFX3500 Switch) on page 121
monitor interface traffic on page 122
monitor interface traffic (QFX3500 Switch) on page 122
monitor interface traffic detail (QFX3500 Switch) on page 123

Output Fields

Table 9 on page 111 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 9: monitor interface Output Fields

Field Name	Field Description	Level of Output
router1	Hostname of the router.	All levels
Seconds	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
Time	Current time (UTC).	All levels
Delay *x*/*y*/*z*	Time difference between when the statistics were displayed and the actual clock time. • ***x***—Time taken for the last polling (in milliseconds). • ***y***—Minimum time taken across all pollings (in milliseconds). • ***z***—Maximum time taken across all pollings (in milliseconds).	All levels
Interface	Short description of the interface, including its name, status, and encapsulation.	All levels
Link	State of the link: **Up**, **Down**, or **Test**.	All levels
Current delta	Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.	All levels

Table 9: monitor interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local Statistics	(Logical interfaces only) Number and rate of bytes and packets destined to the router or switch through the specified interface. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels
Remote Statistics	(Logical interfaces only) Statistics for traffic transiting the router or switch. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels
Traffic statistics	Total number of bytes and packets received and transmitted on the interface. These statistics are the sum of the local and remote statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It usually takes less than 1 second for this counter to stabilize. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface. • **Output packets**—Number of packets transmitted on the interface.	All levels
Description	With the **traffic** option, displays the interface description configured at the **[edit interfaces *interface-name*]** hierarchy level.	**detail**

Sample Output

monitor interface (Physical)

```
user@host> monitor interface so-0/0/0
```

```
router1                          Seconds: 19                 Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC48
Traffic statistics:                                        Current Delta
  Input packets:                     6045 (0 pps)                  [11]
  Input bytes:                    6290065 (0 bps)                [13882]
  Output packets:                   10376 (0 pps)                  [10]
  Output bytes:                  10365540 (0 bps)                 [9418]
Encapsulation statistics:
  Input keepalives:                  1901                          [2]
  Output keepalives:                 1901                          [2]
  NCP state: Opened
  LCP state: Opened
Error statistics:
  Input errors:                         0                          [0]
  Input drops:                          0                          [0]
  Input framing errors:                 0                          [0]
  Policed discards:                     0                          [0]
  L3 incompletes:                       0                          [0]
  L2 channel errors:                    0                          [0]
  L2 mismatch timeouts:                 0                          [0]
  Carrier transitions:                  1                          [0]
  Output errors:                        0                          [0]
  Output drops:                         0                          [0]
  Aged packets:                         0                          [0]
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                             1                          [0]
  LOF count                             1                          [0]
  SEF count                             1                          [0]
  ES-S                                  0                          [0]
  SES-S                                 0                          [0]
SONET statistics:
  BIP-B1                           458871                          [0]
  BIP-B2                           460072                          [0]
  REI-L                            465610                          [0]
  BIP-B3                           458978                          [0]
  REI-P                            458773                          [0]
Received SONET overhead:
  F1      : 0x00  J0        : 0x00  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0x00
  C2(cmp) : 0x00  F2        : 0x00  Z3        : 0x00
```

```
  Z4      : 0x00  S1(cmp)   : 0x00
Transmitted SONET overhead:
  F1      : 0x00  J0        : 0x01  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0xcf
  F2      : 0x00  Z3        : 0x00  Z4        : 0x00

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (OTN Interface)

user@host> **monitor interface ge-7/0/0**

```
Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                          0 (0 bps)
  Output bytes:                         0 (0 bps)
  Input packets:                        0 (0 pps)
  Output packets:                       0 (0 pps)
Error statistics:
  Input errors:                         0
  Input drops:                          0
  Input framing errors:                 0
  Policed discards:                     0
  L3 incompletes:                       0
  L2 channel errors:                    0
  L2 mismatch timeouts:                 0
  Carrier transitions:                  5
  Output errors:                        0
  Output drops:                         0
  Aged packets:                         0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets                       0
  Broadcast packets                     0
  Multicast packets                     0
  Oversized frames                      0
  Packet reject count                   0
  DA rejects                            0
  SA rejects                            0
Output MAC/Filter Statistics:
```

```
  Unicast packets                      0
  Broadcast packets                    0
  Multicast packets                    0
  Packet pad count                     0
  Packet error count                   0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                                2
    LOF                                9
  OTN OTU - FEC Statistics
    Corr err ratio                   N/A
    Corr bytes                         0
    Uncorr words                       0
  OTN OTU - Counters
    BIP                                0
    BBE                                0
    ES                                 0
    SES                                0
    UAS                              422
  OTN ODU - Counters
    BIP                                0
    BBE                                0
    ES                                 0
    SES                                0
    UAS                              422
  OTN ODU - Received Overhead    APSPCC 0-3:             0
```

monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface)

`user@host>` **monitor interface xe-0/0/3**

```
Interface: xe-0/0/3, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                                          Current delta
  Input bytes:                       0 (0 bps)                       [0]
  Output bytes:                      0 (0 bps)                       [0]
  Input packets:                     0 (0 pps)                       [0]
  Output packets:                    0 (0 pps)                       [0]
Error statistics:
  Input errors:                      0                               [0]
  Input drops:                       0                               [0]
  Input framing errors:              0                               [0]
```

```
  Policed discards:                    0                                    [0]
  L3 incompletes:                      0                                    [0]
  L2 channel errors:                   0                                    [0]
  L2 mismatch timeouts:                0                                    [0]
  Carrier transitions:                 5                                    [0]
  Output errors:                       0                                    [0]
  Output drops:                        0                                    [0]
  Aged packets:                        0                                    [0]
Active alarms : None
Active defects: None
PCS statistics:                      Seconds
  Bit Errors                           0                                    [0]
  Errored blocks                       4                                    [0]
Input MAC/Filter statistics:
  Unicast packets                      0                                    [0]
  Broadcast packets                    0                                    [0]
  Multicast packets                    0                                    [0]
  Oversized frames                     0                                    [0]
  Packet reject count                  0                                    [0]
  DA rejects                           0                                    [0]
  SA rejects                           0                                    [0]
Output MAC/Filter Statistics:
  Unicast packets                      0                                    [0]
  Broadcast packets                    0                                    [0]
  Multicast packets                    0                                    [0]
  Packet pad count                     0                                    [0]
  Packet error count                   0                                    [0]



Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface)

`user@host>` **monitor interface et-2/1/0**

```
Interface: et-2/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:                                          Current delta
  Input bytes:                    0 (0 bps)                          [0]
  Output bytes:                   0 (0 bps)                          [0]
  Input packets:                  0 (0 pps)                          [0]
  Output packets:                 0 (0 pps)                          [0]
Error statistics:
```

```
  Input errors:                        0                                    [0]
  Input drops:                         0                                    [0]
  Input framing errors:                0                                    [0]
  Policed discards:                    0                                    [0]
  L3 incompletes:                      0                                    [0]
  L2 channel errors:                   0                                    [0]
  L2 mismatch timeouts:                0                                    [0]
  Carrier transitions:               263                                    [0]
  Output errors:                       0                                    [0]
  Output drops:                        0                                    [0]
  Aged packets:                        0                                    [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                              129                                    [0]
    LOF                                2                                    [0]
  OTN OTU - FEC Statistics
    Corr err ratio                 <8E-5
    Corr bytes               169828399453                                   [0]
    Uncorr words              28939961456                                   [0]
  OTN OTU - Counters                                                       [0]
    BIP                                0
    BBE                                0                                    [0]
    ES                                24                                    [0]
    SES                                0                                    [0]
    UAS                             1255                                    [0]
  OTN ODU - Counters                                                       [0]
    BIP                                0
    BBE                                0                                    [0]
    ES                                24                                    [0]
    SES                                0                                    [0]
    UAS                             1256                                    [0]
  OTN ODU - Received Overhead                                              [0]
    APSPCC 0-3:              00 00 00 00


Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

```
user@host> monitor interface xe-6/1/0
```

```
Interface: xe-6/1/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                        [0]
  Output bytes:                       0 (0 bps)                        [0]
  Input packets:                      0 (0 pps)                        [0]
  Output packets:                     0 (0 pps)                        [0]
Error statistics:
  Input errors:                       0                                [0]
  Input drops:                        0                                [0]
  Input framing errors:               0                                [0]
  Policed discards:                   0                                [0]
  L3 incompletes:                     0                                [0]
  L2 channel errors:                  0                                [0]
  L2 mismatch timeouts:               0                                [0]
  Carrier transitions:                1                                [0]
  Output errors:                      0                                [0]
  Output drops:                       0                                [0]
  Aged packets:                       0                                [0]
Active alarms : None
Active defects: None
PCS statistics:                     Seconds
  Bit Errors                          0                                [0]
  Errored blocks                      1                                [0]
Input MAC/Filter statistics:
  Unicast packets                     0                                [0]
  Broadcast packets                   0                                [0]
  Multicast packets                   0                                [0]
  Oversized frames                    0                                [0]
  Packet reject count                 0                                [0]
  DA rejects                          0                                [0]
  SA rejects                          0                                [0]
Output MAC/Filter Statistics:
  Unicast packets                     0                                [0]
  Broadcast packets                   0                                [0]
  Multicast packets                   0                                [0]
  Packet pad count                    0                                [0]
  Packet error count                  0                                [0]



Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface)

user@host> **monitor interface et-9/0/0**

```
Interface: et-9/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 100000mbps
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                        [0]
  Output bytes:                       0 (0 bps)                        [0]
  Input packets:                      0 (0 pps)                        [0]
  Output packets:                     0 (0 pps)                        [0]
Error statistics:
  Input errors:                       0                                [0]
  Input drops:                        0                                [0]
  Input framing errors:               0                                [0]
  Policed discards:                   0                                [0]
  L3 incompletes:                     0                                [0]
  L2 channel errors:                  0                                [0]
  L2 mismatch timeouts:               0                                [0]
  Carrier transitions:                1                                [0]
  Output errors:                      0                                [0]
  Output drops:                       0                                [0]
  Aged packets:                       0                                [0]




Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface)

user@host> **monitor interface xe-3/0/0**

```
host name                           Seconds: 67                 Time: 23:46:46
                                                              Delay: 0/0/13
Interface: xe-3/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:                                          Current delta
  Input bytes:                        0 (0 bps)                        [0]
  Output bytes:                       0 (0 bps)                        [0]
  Input packets:                      0 (0 pps)                        [0]
  Output packets:                     0 (0 pps)                        [0]
Error statistics:
  Input errors:                       0                                [0]
  Input drops:                        0                                [0]
```

```
  Input framing errors:                0                                   [0]
  Policed discards:                    0                                   [0]
  L3 incompletes:                      0                                   [0]
  L2 channel errors:                   0                                   [0]
  L2 mismatch timeouts:                0                                   [0]
  Carrier transitions:                 3                                   [0]
  Output errors:                       0                                   [0]
  Output drops:                        0                                   [0]
  Aged packets:                        0                                   [0]
OTN Link 0
  OTN Alarms:
  OTN Defects:
  OTN OC - Seconds
    LOS                                0                                   [0]
    LOF                                0                                   [0]
  OTN OTU - FEC Statistics
    Corr err ratio                   N/A
    Corr bytes                         0                                   [0]
    Uncorr words                       0                                   [0]
  OTN OTU - Counters                                                       [0]
    BIP                                0
    BBE                                0                                   [0]
    ES                                 0                                   [0]
    SES                                0                                   [0]
    UAS                                0                                   [0]
  OTN ODU - Counters                                                       [0]
    BIP                                0
    BBE                                0                                   [0]
    ES                                 0                                   [0]
    SES                                0                                   [0]
    UAS                                0                                   [0]
  OTN ODU - Received Overhead                                              [0]
    APSPCC 0-3:              00 00 00 00


Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (Logical)

`user@host>` **monitor interface so-1/0/0.0**

```
host name               Seconds: 16                  Time: 15:33:39
                                                             Delay: 0/0/1
Interface: so-1/0/0.0, Enabled, Link is Down
```

```
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:                                          Current delta
  Input bytes:                        0                              [0]
  Output bytes:                       0                              [0]
  Input packets:                      0                              [0]
  Output packets:                     0                              [0]
Remote statistics:
  Input bytes:                        0 (0 bps)                      [0]
  Output bytes:                       0 (0 bps)                      [0]
  Input packets:                      0 (0 pps)                      [0]
  Output packets:                     0 (0 pps)                      [0]
Traffic statistics:
  Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (QFX3500 Switch)

```
user@switch> monitor interface ge-0/0/0
```

```
Interface: ge-0/0/0, Enabled, Link is Down
Encapsulation: Ethernet, Speed: Unspecified
Traffic statistics:                                        Current delta
  Input bytes:                      0 (0 bps)                      [0]
  Output bytes:                     0 (0 bps)                      [0]
  Input packets:                    0 (0 pps)                      [0]
  Output packets:                   0 (0 pps)                      [0]
Error statistics:
  Input errors:                     0                              [0]
  Input drops:                      0                              [0]
  Input framing errors:             0                              [0]
  Policed discards:                 0                              [0]
  L3 incompletes:                   0                              [0]
  L2 channel errors:                0                              [0]
  L2 mismatch timeouts:             0                              [0]
  Carrier transitions:              0                              [0]
  Output errors:                    0                              [0]
  Output drops:                     0                              [0]
  Aged packets:                     0                              [0]
Active alarms : LINK
Active defects: LINK
Input MAC/Filter statistics:
  Unicast packets                   0                              [0]
```

```
  Broadcast packets                  0  Multicast packet             [0]

Interface warnings:
  o Outstanding LINK alarm
```

monitor interface traffic

user@host> **monitor interface traffic**

```
host name               Seconds: 15                Time: 12:31:09

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-1/0/0    Down              0          (0)                  0          (0)
 so-1/1/0    Down              0          (0)                  0          (0)
 so-1/1/1    Down              0          (0)                  0          (0)
 so-1/1/2    Down              0          (0)                  0          (0)
 so-1/1/3    Down              0          (0)                  0          (0)
 t3-1/2/0    Down              0          (0)                  0          (0)
 t3-1/2/1    Down              0          (0)                  0          (0)
 t3-1/2/2    Down              0          (0)                  0          (0)
 t3-1/2/3    Down              0          (0)                  0          (0)
 so-2/0/0      Up         211035          (1)              36778          (0)
 so-2/0/1      Up         192753          (1)              36782          (0)
 so-2/0/2      Up         211020          (1)              36779          (0)
 so-2/0/3      Up         211029          (1)              36776          (0)
 so-2/1/0      Up         189378          (1)              36349          (0)
 so-2/1/1    Down              0          (0)              18747          (0)
 so-2/1/2    Down              0          (0)              16078          (0)
 so-2/1/3      Up              0          (0)              80338          (0)
 at-2/3/0      Up              0          (0)                  0          (0)
 at-2/3/1    Down              0          (0)                  0          (0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

monitor interface traffic (QFX3500 Switch)

user@switch> **monitor interface traffic**

```
switch                           Seconds: 7                  Time: 16:04:37

Interface    Link  Input packets        (pps)     Output packets        (pps)
 ge-0/0/0    Down              0          (0)                  0          (0)
 ge-0/0/1      Up         392187          (0)             392170          (0)
```

```
 ge-0/0/2    Down              0          (0)                0          (0)
 ge-0/0/3    Down              0          (0)                0          (0)
 ge-0/0/4    Down              0          (0)                0          (0)
 ge-0/0/5    Down              0          (0)                0          (0)
 ge-0/0/6    Down              0          (0)                0          (0)
 ge-0/0/7    Down              0          (0)                0          (0)
 ge-0/0/8    Down              0          (0)                0          (0)
 ge-0/0/9      Up         392184          (0)           392171          (0)
 ge-0/0/10   Down              0          (0)                0          (0)
 ge-0/0/11   Down              0          (0)                0          (0)
 ge-0/0/12   Down              0          (0)                0          (0)
 ge-0/0/13   Down              0          (0)                0          (0)
 ge-0/0/14   Down              0          (0)                0          (0)
 ge-0/0/15   Down              0          (0)                0          (0)
 ge-0/0/16   Down              0          (0)                0          (0)
 ge-0/0/17   Down              0          (0)                0          (0)
 ge-0/0/18   Down              0          (0)                0          (0)
 ge-0/0/19   Down              0          (0)                0          (0)
 ge-0/0/20   Down              0          (0)                0          (0)
 ge-0/0/21   Down              0          (0)                0          (0)
 ge-0/0/22     Up         392172          (0)           392187          (0)
 ge-0/0/23     Up         392185          (0)           392173          (0)
 vcp-0       Down              0                             0
 vcp-1       Down              0                             0
 ae0         Down              0          (0)                0          (0)
 bme0          Up              0                       1568706
```

monitor interface traffic detail (QFX3500 Switch)

`user@switch>` **monitor interface traffic detail**

```
switch                                                  Seconds: 74
       Time: 16:03:02

Interface    Link  Input packets        (pps)     Output packets        (pps)
Description
 ge-0/0/0    Down              0          (0)                0          (0)
 ge-0/0/1      Up         392183          (0)           392166          (0)
 ge-0/0/2    Down              0          (0)                0          (0)
 ge-0/0/3    Down              0          (0)                0          (0)
 ge-0/0/4    Down              0          (0)                0          (0)
 ge-0/0/5    Down              0          (0)                0          (0)
 ge-0/0/6    Down              0          (0)                0          (0)
 ge-0/0/7    Down              0          (0)                0          (0)
```

```
ge-0/0/8    Down             0          (0)              0          (0)
ge-0/0/9      Up        392181          (0)         392168          (0)
ge-0/0/10   Down             0          (0)              0          (0)
ge-0/0/11   Down             0          (0)              0          (0)
ge-0/0/12   Down             0          (0)              0          (0)
ge-0/0/13   Down             0          (0)              0          (0)
ge-0/0/14   Down             0          (0)              0          (0)
ge-0/0/15   Down             0          (0)              0          (0)
ge-0/0/16   Down             0          (0)              0          (0)
ge-0/0/17   Down             0          (0)              0          (0)
ge-0/0/18   Down             0          (0)              0          (0)
ge-0/0/19   Down             0          (0)              0          (0)
ge-0/0/20   Down             0          (0)              0          (0)
ge-0/0/21   Down             0          (0)              0          (0)
ge-0/0/22     Up        392169          (0)         392184          (1)
ge-0/0/23     Up        392182          (0)         392170          (0)
vcp-0       Down             0                           0
vcp-1       Down             0                           0
ae0         Down             0          (0)              0          (0)
bme0          Up             0                     1568693
```

monitor start

Syntax

```
monitor start filename
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Start displaying the system log or trace file and additional entries being added to those files.

Options

filename—Specific log or trace file.

Additional Information

Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the **syslog** statement at the **[edit system]** hierarchy level and the **options** statement at the **[edit routing-options]** hierarchy level. The trace files generated by the routing protocol process are configured with **traceoptions** statements at the **[edit routing-options]**, **[edit interfaces]**, and **[edit protocols *protocol*]** hierarchy levels.

> **NOTE:** To monitor a log file within a logical system, issue the **monitor start *logical-system-name/filename*** command.

Required Privilege Level

trace

RELATED DOCUMENTATION

monitor list

monitor stop | **127**

List of Sample Output

monitor start on page 126

Output Fields

Table 10 on page 126 describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

Table 10: monitor start Output Fields

Field Name	Field Description
******filename******	Name of the file from which entries are being displayed. This line is displayed initially and when the command switches between log files.
Date and time	Timestamp for the log entry.

Sample Output

monitor start

```
user@host> monitor start system-log
```

```
*** system-log***
Jul 20 15:07:34 hang sshd[5845]: log: Generating 768 bit RSA key.
Jul 20 15:07:35 hang sshd[5845]: log: RSA key generation complete.
Jul 20 15:07:35 hang sshd[5845]: log: Connection from 204.69.248.180 port 912
Jul 20 15:07:37 hang sshd[5845]: log: RSA authentication for root accepted.
Jul 20 15:07:37 hang sshd[5845]: log: ROOT LOGIN as 'root' from host.example.com
Jul 20 15:07:37 hang sshd[5845]: log: Closing connection to 204.69.248.180
```

monitor stop

Syntax

```
monitor stop filename
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Stop displaying the system log or trace file.

Options

filename—Specific log or trace file.

Additional Information

Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are those configured with the **syslog** statement at the **[edit system]** hierarchy level and the **options** statement at the **[edit routing-options]** hierarchy level. The trace files generated by the routing protocol process are those configured with **traceoptions** statements at the **[edit routing-options]**, **[edit interfaces]**, and **[edit protocols *protocol*]** hierarchy levels.

Required Privilege Level

trace

RELATED DOCUMENTATION

monitor list

monitor start | 125

List of Sample Output

monitor stop on page 128

Output Fields

This command produces no output.

Sample Output

monitor stop

```
user@host> monitor stop
```

ping

List of Syntax

Syntax on page 129
Syntax (QFX Series) on page 129
Syntax (Junos OS Evolved) on page 130

Syntax

```
ping host
<bypass-routing>
<ce-ip destination-ip-address instance routing-instance-name source-ip source-ip-address>
<count requests>
<do-not-fragment>
<inet | inet6>
<interface source-interface>
<interval seconds>
<no-resolve>
<pattern string>
<rapid>
<record-route>
<routing-instance routing-instance-name>
<logical-system logical-system-name>
<tenant tenant-name>
<size bytes>
<source source-address>
<tos type-of-service>
<ttl value>
<verbose>
<wait seconds>
```

Syntax (QFX Series)

```
ping host
<bypass-routing>
<count requests>
<detail>
<do-not-fragment>
<inet>
<interface source-interface>
<interval seconds>
<logical-system logical-system-name>
<loose-source value>
```

<mac-address *mac-address*>
<no-resolve>
<pattern *string*>
<rapid>
<record-route>
<routing-instance *routing-instance-name*>
<size *bytes*>
<source *source-address*>
<strict>
< strict-source *value*>
<tos *type-of-service*>
<ttl *value*>
<verbose>
<wait *seconds*>

Syntax (Junos OS Evolved)

ping *host*
<bypass-routing>
<ce-ip *destination-ip-address* instance *routing-instance-name* source-ip *source-ip-address*>
<count *requests*>
<do-not-fragment>
<inet | inet6>
<interface *source-interface*>
<interval *seconds*>
<no-resolve>
<pattern *string*>
<rapid>
<record-route>
<routing-instance *routing-instance-name*>
<size *bytes*>
<source *source-address*>
<tos *type-of-service*>
<ttl *value*>
<verbose>
<wait *seconds*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
ce-ip option introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

The following options are deprecated for Junos OS Evolved Release 18.3R1: **detail**, **logical-system**, **loose-source**, **mac-address**, **strict**, **strict-source**, and **vpls**.
The command **tenant** option is introduced in Junos OS Release 19.2R1 for SRX Series.

Description

Check host reachability and network connectivity. The **ping** command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Press Ctrl+c to interrupt a ping command.

Options

host—IP address or hostname of the remote system to ping.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

ce-ip *destination-ip-address* instance *routing-instance-name* source-ip *source-ip-address*—(MX Series routers with MPC and MIC interfaces only) (Optional) Check the connectivity information of customer edge (CE) devices, such as reachability, attachment points, and MAC addresses, from a provider edge (PE) device in a virtual private LAN service (VPLS), hierarchical VPLS (H-VPLS), and Ethernet VPN (EVPN) network. The **ce-ip** option is based on the LSP ping infrastructure, where the **ping** utility is extended to use the CE device IP address as the target host and the PE device loopback address as the source for a specific VPLS or EVPN routing instance.

destination-ip-address—IPv4 address of the CE device to ping.

instance *routing-instance-name*—Name of the VPLS or EVPN routing instance. The command output displays the connectivity information of the CE device based on the configured routing instance type.

source-ip *source-ip-address*—Loopback address of the PE device.

count *requests*—(Optional) Number of ping requests to send. The range of values is **1** through **2,000,000,000**. The default value is an unlimited number of requests.

detail—(Optional) This option is not supported for Junos OS Evolved Release 18.3R1. Include in the output the interface on which the ping reply was received.

do-not-fragment—(Optional) Set the do-not-fragment (DF) flag in the IP header of the ping packets.

For Junos OS Evolved Release 18.3R1, IPv6 **ping** does not have **do-not-fragment** support. The **ping** command is identified as IPv6 Ping when destination is IPv6 address or **inet6** option is used.

For Junos OS IPv6 packets, this option disables fragmentation.

NOTE: In Junos OS Release 11.1 and later, when issuing the **ping** command for an IPv6 route with the **do-not-fragment** option, the maximum ping packet size is calculated by subtracting 48 bytes (40 bytes for the IPV6 header and 8 bytes for the ICMP header) from the MTU. Therefore, if the ping packet size (including the 48-byte header) is greater than the MTU, the ping operation might fail.

inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.

inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.

interface *source-interface*—(Optional) Interface to use to send the ping requests.

interval *seconds*—(Optional) How often to send ping requests. The range of values, in seconds, is **1** through infinity. The default value is **1**.

logical-system *logical-system-name*—(Optional) Name of logical system from which to send the ping requests.

Alternatively, enter the **set cli logical-system *logical-system-name*** command and then run the **ping** command. To return to the main router or switch, enter the **clear cli logical-system** command.

tenant *tenant-name*—(Optional) Name of tenant system from which to send the ping requests.

loose-source *value*—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

mac-address *mac-address*—(Optional) Ping the physical or hardware address of the remote system you are trying to reach.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern *string*—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the **count** option.

record-route—(Optional) Record and report the packet's path (IPv4).

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the ping attempt. For Junos OS Evolved, the **routing-instance** option supports only **mgmt_junos**.

size *bytes*—(Optional) Size of ping request packets. The range of values, in bytes, is **0** through **65,468**. The default value is **56**, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

strict—(Optional) Use the strict source route option (IPv4).

strict-source *value*—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos *type-of-service*—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is **0** through **255**.

If the device configuration includes the **dscp-code-point *value*** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level, the configured DSCP value overrides the value specified in this command option. In this case, the ToS field of ICMP echo request packets sent on behalf of this command carries the DSCP value specified in the **dscp-code-point** configuration statement instead of the value you specify in this command option.

ttl *value*—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is **0** through **255**.

verbose—(Optional) Display detailed output.

vpls *instance-name*—(Optional) Ping the instance to which this VPLS belongs.

wait *seconds*—(Optional) Maximum wait time, in seconds, after the final packet is sent. If this option is not specified, the default delay is **10** seconds. If this option is used without the count option, a default count of **5** packets is used.

Required Privilege Level
network

RELATED DOCUMENTATION

Rate Limit ICMPv4 and ICMPv6 Traffic

Pinging Customer Edge Device IP Address

List of Sample Output

ping ce-ip <destination-ip-address> instance <routing-instance-name> source-ip <source-ip-address> (EVPN) on page 134
ping ce-ip <destination-ip-address> instance <routing-instance-name> source-ip <source-ip-address> (VPLS) on page 134
ping hostname on page 134
ping hostname rapid on page 135
ping hostname size count on page 135

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. These packets are not counted in the received packets count. They are accounted for separately.

When pinging a nonexistant route, the display output of **ping** command does not print the number of packets sent or received or the packet loss.

Sample Output

ping ce-ip <destination-ip-address> instance <routing-instance-name> source-ip <source-ip-address> (EVPN)

```
user@host> ping ce-ip 10.0.0.4 instance foo source-ip 127.0.0.1
```

```
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
! -> PE5|foo|evpn|ge-0/0/2.100, 00:11:22:33:44:55:66:77:88:99|12:23:ab:98:34:05
--- ce-ip ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping ce-ip <destination-ip-address> instance <routing-instance-name> source-ip <source-ip-address> (VPLS)

```
user@host> ping ce-ip 10.0.0.4 instance foo source-ip 127.0.0.1
```

```
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
! -> PE2|foo|vpls|ge-0/0/2.100|12:23:ab:98:34:02
--- ce-ip ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping hostname

```
user@host> ping device1.example.com
```

```
PING device1.example.com (192.0.2.0): 56 data bytes
64 bytes from 192.0.2.0: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.0.2.0: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.0.2.0: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 1192.0.2.0: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.0.2.0: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.0.2.0: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname rapid

```
user@host> ping device1.example.com rapid
```

```
PING device1.example.com (192.0.2.0): 56 data bytes
!!!!!
--- device1.example.com ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

ping hostname size count

```
user@host> ping device1.example.com size 200 count 5
```

```
PING device1.example.com (192.0.2.0): 200 data bytes
208 bytes from 192.0.2.0: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.0.2.0: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.0.2.0: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.0.2.0: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.0.2.0: icmp_seq=4 ttl=253 time=17.898 ms

--- device1.example.com ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

Output for Junos OS Evolved:

```
user@host> ping 40.0.0.2 count 20 size 500
```

```
connect: No route to host
```

Output for Junos OS:

user@host> **ping 40.0.0.2 count 20 size 500**

```
Aug 02 12:56:56 [INFO ] Step 2: Host and Transit ping has to fail
Aug 02 12:56:56 [TRACE] [R0 evo-ptx-b] [cmd] run ping 40.0.0.2 rapid count 50 size
 500
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] PING 40.0.0.2 (40.0.0.2): 500 data bytes
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ping: sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] .ping: sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ping: sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] .ping: .sendto: No route to host
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ping: sendto: No route to host
ug 02 12:57:21 [TRACE] [R0 evo-ptx-b] ..
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] --- 40.0.0.2 ping statistics ---
Aug 02 12:57:21 [TRACE] [R0 evo-ptx-b] 50 packets transmitted, 0 packets received,
 100% packet loss
```

show ipv6 neighbors

Syntax

```
show ipv6 neighbors
<flags>
<hostname host-name>
<interface interface-name>
<logical-system logical-system-name>
<reference-count count>
<tenant name>
<vpn vpn-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
flags, **hostname**, **interface**, **logical-system**, **reference-count**, **tenant**, and **vpn** options added in Junos OS Release 18.3.

Description

Display information about the IPv6 neighbor cache.

> **NOTE:** Starting with Junos OS Release 16.1, **show ipv6 neighbors** 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 IPv6 table.

flags—(Optional) Display the flags set, if any.

hostname *host-name*—(Optional) Display the hostname.

interface *interface-name*—(Optional) Display information about IPv6 for the specified logical interface

logical-system *logical-system-name*—(Optional) Display the IPv6 entries for the specified logical system; only available on the main router context.

reference-count *count*—(Optional) Display the IPv6 next-hop reference count.

tenant *name*—(Optional) Displays the IPv6 entries for the specified tenant system; only available in the main router context.

vpn *vpn-name*—(Optional) Display entries in the IPv6 table for the specified virtual private network's (VPN) routing table.

Required Privilege Level
view

RELATED DOCUMENTATION

clear ipv6 neighbors | **106**

List of Sample Output
show ipv6 neighbors on page 139
show ipv6 neighbors on page 139

Output Fields
Table 11 on page 138 describes the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 11: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: **up**, **down**, **incomplete**, **reachable**, **stale**, or **unreachable**.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: **yes** or **no**.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: **yes** or **no**.
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
```

```
IPv6 Address                Linklayer Address  State       Exp Rtr Secure Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale       546 yes no
fe-1/2/0.1
fe80::2a0:a514:0:24c         00:05:85:8f:c8:bd  stale       258 yes no
fe-1/2/0.1
fe80::2a0:a514:0:64c         00:05:85:8f:c8:bd  stale       111 yes no
fe-1/2/1.5
fe80::2a0:a514:0:a4c         00:05:85:8f:c8:bd  stale       327 yes no
fe-1/2/2.9
```

show ipv6 neighbors

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

```
IPv6 Address                Linklayer Address  State       Exp Rtr Secure Interface
23::23:0:0:2                  00:00:23:00:00:02  reachable   0   no  no      irb.0
 [xe-2/2/0.0]
```

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

```
IPv6 Address                Linklayer Address  State       Exp Rtr Secure Interface
23::23:0:0:2                  00:00:23:00:00:02  reachable   0   no  no      irb.0
```

show ipv6 router-advertisement

Syntax

```
show ipv6 router-advertisement
<conflicts>
<interface interface>
<logical-system (all | logical-system-name)>
<prefix prefix/prefix length>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

Options

none—Display all IPv6 router advertisement information for all interfaces.

conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.

interface *interface*—(Optional) Display IPv6 router advertisement information for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix *prefix/prefix length*—(Optional) Display IPv6 router advertisement information for the specified prefix.

Additional Information

The display identifies conflicting information by enclosing the value the router is advertising in brackets.

Required Privilege Level

view

RELATED DOCUMENTATION

clear ipv6 router-advertisement | 108

List of Sample Output

show ipv6 router-advertisement on page 142
show ipv6 router-advertisement conflicts on page 142
show ipv6 router-advertisement prefix on page 142

Output Fields

Table 12 on page 141 describes the output fields for the **show ipv6 router-advertisement** command. Output fields are listed in the approximate order in which they appear.

Table 12: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.
Advertisements sent	Number of router advertisements sent and the elapsed time since they were sent.
Solicits received	Number of solicitation messages received.
Advertisements received	Number of router advertisements received.
Advertisements from	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
Managed	Managed address configuration flag: **0** (stateless) or **1** (stateful).
Other configuration	Other stateful configuration flag: **0** (stateless) or **1** (stateful).
Reachable time	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
Default lifetime	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
Retransmit timer	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
Current hop limit	Configured current hop limit.
Prefix	Name and length of the prefix.
Valid lifetime	How long the prefix remains valid for onlink determination.
Preferred lifetime	How long the prefix generated by stateless autoconfiguration remains preferred.
On link	Onlink flag: **0** (not onlink) or **1** (onlink).

Table 12: show ipv6 router-advertisement Output Fields *(continued)*

Field Name	Field Description
Autonomous	Autonomous address configuration flag: **0** (not autonomous) or **1** (autonomous).

Sample Output

show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
```

```
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
```

```
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 2001:db8:8040::/16
```

```
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
```

```
Solicits received: 0
Advertisements received: 3
Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
  Managed: 0
  Other configuration: 0
  Reachable time: 0 ms
  Default lifetime: 180 sec [1800 sec]
  Retransmit timer: 0 ms
  Current hop limit: 64
  Prefix: 2001:db8:8040:1::/64
    Valid lifetime: 2592000 sec
    Preferred lifetime: 604800 sec
    On link: 1
    Autonomous: 1
```

show log

List of Syntax

Syntax on page 144
Syntax (QFX Series and OCX Series) on page 144
Syntax (TX Matrix Router) on page 144

Syntax

```
show log
<filename | user <username>>
```

Syntax (QFX Series and OCX Series)

```
show log filename
<device-type (device-id | device-alias)>
```

Syntax (TX Matrix Router)

```
show log
<all-lcc | lcc number | scc>
<filename | user <username>>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Command introduced in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

> **NOTE:** On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options

none—List all log files.

<all-lcc | lcc *number* | scc>—(Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace ***number*** with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.

> **NOTE:** If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(*device-id* | *device-alias*)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of **messages**.

user <*username*>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include ***username***, display logging information about the specified user.

Required Privilege Level
trace

RELATED DOCUMENTATION

syslog (System)

List of Sample Output
show log on page 146
show log filename on page 147
show log filename (QFabric System) on page 149
show log user on page 150
show log accepted-traffic (SRX4600, SRX5400, SRX5600, and SRX5800) on page 150

Sample Output

show log

```
user@host> show log
```

```
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
```

```
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
```

show log filename

user@host> **show log rpd**

```
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr 192.0.2.21
 nhop type local nhop 192.0.2.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr 192.0.2.22
 nhop type unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

user@host:LSYS1> **show log flow_lsys1.log**

```
Nov  7 07:34:09 07:34:09.491800:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491809:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
table lock

Nov  7 07:34:09 07:34:09.491840:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491841:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
table lock

Nov  7 07:34:09 07:34:09.491854:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh
 0x0

Nov  7 07:34:09 07:34:09.491868:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491869:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
```

```
table lock

Nov  7 07:34:09 07:34:09.491881:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh
 0x0
```

user@host:TSYS1> **show log flow_tsys1.log**

```
Nov  7 13:21:47
13:21:47.217744:CID-0:THREAD_ID-05:LSYS_ID-32:RT:<192.0.2.0/0->198.51.100.0/9011;1,0x0>
 :

Nov  7 13:21:47 13:21:47.217747:CID-0:THREAD_ID-05:LSYS_ID-32:RT:packet [84] ipid
 = 39281, @0x7f490ae56d52

Nov  7 13:21:47 13:21:47.217749:CID-0:THREAD_ID-05:LSYS_ID-32:RT:----
flow_process_pkt: (thd 5): flow_ctxt type 0, common flag 0x0, mbuf 0x4882b600,
rtbl7

Nov  7 13:21:47 13:21:47.217752:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow process pak
 fast ifl 88 in_ifp lt-0/0/0.101

Nov  7 13:21:47 13:21:47.217753:CID-0:THREAD_ID-05:LSYS_ID-32:RT:
lt-0/0/0.101:192.0.2.0->198.51.100.0, icmp, (0/0)

Nov  7 13:21:47 13:21:47.217756:CID-0:THREAD_ID-05:LSYS_ID-32:RT: find flow: table
 0x11d0a2680, hash 20069(0xffff), sa 192.0.2.0, da 198.51.100.0, sp 0, d0

Nov  7 13:21:47 13:21:47.217760:CID-0:THREAD_ID-05:LSYS_ID-32:RT:Found: session
id 0x12. sess tok 28685

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow got session.


Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT:  flow session
id 18

Nov  7 13:21:47 13:21:47.217763:CID-0:THREAD_ID-05:LSYS_ID-32:RT: vector bits 0x200
 vector 0x84ae85f0

Nov  7 13:21:47 13:21:47.217764:CID-0:THREAD_ID-05:LSYS_ID-32:RT:set nat
0x11e463550(18) timeout const to 2

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT: set_nat_timeout
 2 on session 18
```

```
Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT:refresh nat
0x11e463550(18) timeout to 2

Nov  7 13:21:47 13:21:47.217767:CID-0:THREAD_ID-05:LSYS_ID-32:RT:insert usp tag
for apps

Nov  7 13:21:47 13:21:47.217768:CID-0:THREAD_ID-05:LSYS_ID-32:RT:mbuf 0x4882b600,
 exit nh 0xfffb0006
```

show log filename (QFabric System)

`user@qfabric>` **show log messages**

```
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486 file:
 UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486 file:
 UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default___NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default___NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
```

```
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492 file:
 UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492 file:
 UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

user@host> **show log user**

```
usera    mg2546                  Thu Oct  1 19:37   still logged in
usera    mg2529                  Thu Oct  1 19:08 - 19:36  (00:28)
usera    mg2518                  Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                  Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    aaa.bbbb.com   Wed Sep 30 18:39 - 18:41  (00:02)
userb    ttyp1    192.0.2.0      Wed Sep 30 01:03 - 01:22  (00:19)
```

show log accepted-traffic (SRX4600, SRX5400, SRX5600, and SRX5800)

user@host> **show log accepted-traffic**

```
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.5/2->4.4.4.2/63 0x0 None 3.3.3.5/2->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2617282058 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.4/4->4.4.4.2/63 0x0 None 3.3.3.4/4->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162754 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.4/1->4.4.4.2/63 0x0 None 3.3.3.4/1->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162755 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.3/0->4.4.4.2/63 0x0 None 3.3.3.3/0->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162752 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
```

```
3.3.3.5/5->4.4.4.2/63 0x0 None 3.3.3.5/5->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162751 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04  sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.3/3->4.4.4.2/63 0x0 None 3.3.3.3/3->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162753 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
```

traceroute

List of Syntax

Syntax on page 152
Syntax (QFX Series and OCX Series) on page 152

Syntax

```
traceroute
<host>
<as-number-lookup>
<bypass-routing>
<ce-ip ip address>
<clns>
<ethernet host>
<gateway address>
<inet>
<inet6>
<interface interface-name>
<logical-system logical-system-name>
<monitor host>
<mpls (ldp FEC address | rsvp label-switched-path-name | bgp FEC address)>
<next-hop next-hop address>
<no-resolve host>
<overlay>
<port>
<propagate-ttl host>
<routing-instance routing-instance-name>
<tenant tenant-name>
<source source-address>
<tos value>
<ttl value>
<wait seconds>
```

Syntax (QFX Series and OCX Series)

```
traceroute
<host>
<as-number-lookup>
<bypass-routing>
<gateway address>
<inet>
<inet6>
```

<interface *interface-name*>
<monitor *host*>
<mpls (ldp *FEC address* | rsvp *label-switched-path-name* | bgp *FEC address*)>
<next-hop *next-hop address*>
<no-resolve *host*>
<overlay>
<port>
<propagate-ttl *host*>
<routing-instance *routing-instance-name*>
<source *source-address*>
<tos *value*>
<ttl *value*>
<wait *seconds*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
mpls option introduced in Junos OS Release 9.2.
propagate-ttl option introduced in Junos OS Release 12.1.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Support for IPv6 traceroute with **as-number-lookup** introduced with Junos OS Release 18.3R1.
The command **tenant** option is introduced in Junos OS Release 19.2R1 for the SRX Series.
The following options are deprecated in Junos OS Evolved Release 18.3R1: **logical-system** and **propagate-ttl**.

Description

Display the route that packets take to a specified network host. Use **traceroute** as a debugging tool to locate points of failure in a network.

Options

host—IP address or name of remote host.

as-number-lookup—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

bypass-routing—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

ce-ip *IPv4 or IPv6 address*—(MX Series routers with MPC and MIC interfaces only) (Optional) Check the route to a customer edge (CE) IP address in a virtual private LAN service (VPLS) and Ethernet VPN (EVPN) network.

clns—(Optional) Trace the route belonging to Connectionless Network Service (CLNS).

ethernet *host*—(Optional) Trace the route to an Ethernet host using unicast MAC address.

gateway *address*—(Optional) Address of a router or switch through which the route transits.

inet | inet6—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

interface *interface-name*—(Optional) Name of the interface over which to send packets.

logical-system (all | *logical-system-name*)—(Optional) This option is not supported for Junos OS Evolved Release 18.3R1. Perform this operation on all logical systems or on a particular logical system.

tenant *tenant-name*—(Optional) Name of a particular tenant system for traceroute attempt.

monitor *host*—(Optional) Display real-time monitoring information for the specified host.

mpls (ldp *FEC address* | rsvp *label-switched-path name*)—(Optional) See **traceroute mpls ldp** and **traceroute mpls rsvp**.

next-hop—The next-hop through which to send packets to a destination.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

overlay—(Optional) Traceroute overlay path.

port—(Optional) Base port number to use in traceroute probes.

propagate-ttl—(Optional) On the PE routing device, use this option to view locally generated Routing Engine transit traffic. This is applicable for MPLS L3VPN traffic only.

Use for troubleshooting, when you want to view hop-by-hop information from the local provider router to the remote provider router, when TTL decrementing is disabled on the core network using the **no-proagate-ttl** configuration statement.

NOTE: Using **propagate-ttl** with **traceroute** on the CE router does not show hop-by-hop information.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the traceroute attempt.

source *source-address*—(Optional) Source address of the outgoing traceroute packets.

tos *value*—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is **0** through **255**.

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is **0** through **128**.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
network

RELATED DOCUMENTATION

traceroute monitor

List of Sample Output

traceroute on page 156
traceroute as-number-lookup host on page 156
traceroute no-resolve on page 156
traceroute propogate-ttl on page 157
traceroute (Between CE Routers, Layer 3 VPN) on page 157
traceroute (Through an MPLS LSP) on page 157
traceroute routing-instance no-resolve (Through an MPLS LSP) on page 157
traceroute (Junos OS Evolved, Through an MPLS LSP) on page 158

Output Fields

Table 13 on page 155 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 13: traceroute Output Fields

Field Name	Field Description
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
Round trip time	Average round-trip time, in milliseconds (ms).

Sample Output

traceroute

```
user@host> traceroute santacruz
```

```
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  blue23 (10.168.1.254)    2.370 ms  2.853 ms  0.367 ms
 2  red14  (10.168.255.250)  0.778 ms  2.937 ms  0.446 ms
 3  yellow (10.156.169.254)  7.737 ms  89.905 ms  0.834 ms
```

traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
```

```
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1  10.39.1.1 (10.39.1.1)  0.779 ms  0.728 ms  0.562 ms
 2  10.39.1.6 (10.39.1.6) [AS  32]  0.657 ms  0.611 ms  0.617 ms
 3  10.100.1.1 (10.100.1.1) [AS  10, 40, 50]  0.880 ms  0.808 ms  0.774 ms
```

```
user@host> traceroute as-number-lookup 1::1
```

```
traceroute6 to 1::1 (1::1) from 2001:b8::7, 64 hops max, 12 byte packets
```

```
user@host> traceroute 2001:b8::7 as-number-lookup
```

```
traceroute6 to 2001:b8::7 (2001:b8::7) from 2001:db8::9, 64 hops max, 12 byte
packets
 1  2001:db8::10 (2001:db8::10) [AS 18] 0.657 ms  17.319 ms  0.504 ms
 2  2001:b8::7 (2001:b8::7)  0.949 ms  0.930 ms  0.739 ms
```

traceroute no-resolve

```
user@host> traceroute santacruz no-resolve
```

```
traceroute to host1.example.com (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

traceroute propogate-ttl

```
user@host> traceroute propagate-ttl 100.200.2.2 routing-instance VPN-A
```

```
traceroute to 100.200.2.2 (100.200.2.2) from 1.1.0.2, 30 hops max, 40 byte packets

 1  1.2.0.2 (1.2.0.2)  2.456 ms  1.753 ms  1.672 ms
     MPLS Label=299776 CoS=0 TTL=1 S=0
     MPLS Label=299792 CoS=0 TTL=1 S=1
 2  1.3.0.2 (1.3.0.2)  1.213 ms  1.225 ms  1.166 ms
     MPLS Label=299792 CoS=0 TTL=1 S=1
 3  100.200.2.2 (100.200.2.2)  1.422 ms  1.521 ms  1.443 ms
```

traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
```

```
traceroute to host2.example.com (10.255.14.179), 30 hops max, 40
      byte packets
       1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
       2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
           MPLS Label=100006 CoS=0 TTL=1 S=1
       3  host2.example.com (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
```

```
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
     MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

traceroute routing-instance no-resolve (Through an MPLS LSP)

```
user@host> traceroute routing-instance VRF-1 198.51.100.1 no-resolve
```

```
traceroute to 198.51.100.1 (198.51.100.1), 30 hops max, 40 byte packets
 1  198.51.100.20  20.243 ms  13.256 ms  24.194 ms
     MPLS Label=299792 CoS=0 TTL=1 S=0
```

```
     MPLS Label=16 CoS=0 TTL=1 S=1
 2  198.51.100.21  14.126 ms  13.090 ms  29.082 ms
     MPLS Label=16 CoS=0 TTL=1 S=0
     MPLS Label=16 CoS=0 TTL=2 S=1
 3  198.51.100.22  16.419 ms  11.564 ms  17.068 ms
     MPLS Label=16 CoS=0 TTL=1 S=1
 4  198.51.100.1  12.794 ms  12.939 ms  17.123 ms
```

traceroute (Junos OS Evolved, Through an MPLS LSP)

The Junos OS Evolved **traceroute** command parses MPLS data in the same way as the Linux **traceroute** command: L=label, E=exp_use, S=stack_bottom, and T=TTL. In the example below, **T=1/L=16** indicates the TTL with label 16.

`user@host>` **traceroute 192.0.2.50 ttl 255**

```
traceroute to 192.0.2.50 (192.0.2.50), 255 hops max, 60 byte packets
 1  192.0.2.60 (192.0.2.60)  13.565 ms  11.696 ms  11.448 ms
 2  192.0.2.61 (192.0.2.61) <MPLS:L=17,E=0,S=0,T=1/L=16,E=0,S=1,T=1>  34.034 ms
31.538 ms  27.697 ms
 3  192.0.2.62 (192.0.2.62) <MPLS:L=299776,E=0,S=0,T=1/L=16,E=0,S=1,T=2>  23.174
ms  24.393 ms  21.009 ms
 4  192.0.2.63 (192.0.2.63)  24.553 ms  19.698 ms  25.648 ms
 5  192.0.2.50 (192.0.2.50)  33.322 ms  29.514 ms  24.706 ms
```