JUNIPER NETWORKS | Engineering Simplicity

Junos® OS

ICMP Router Discovery Protocol 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 ICMP Router Discovery Protocol 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 | vi

Documentation and Release Notes | **vi**

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

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

Documentation Conventions | **viii**

Documentation Feedback | **xi**

Requesting Technical Support | **xi**

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

1 **Overview**

ICMP Router Discovery Overview | 14

- Operation of a Router Discovery Server | **14**
- Router Advertisement Messages | **15**

Supported ICMP Router Discovery and IPv6 Neighbor Discovery Standards | 16

2 **Configuring the ICMP Protocol**

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

3 **Configuring Recursive DNS Servers**

Understanding Recursive DNS Servers for IPv6 | 27

Configuring a Recursive DNS Server Address for IPv6 Hosts | 28

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts | 29

4 Monitoring ICMP Router Discovery

Traceroute for Inactive Interface | 36

Example: Tracing Global Routing Protocol Operations | 36

How to Use the Probe Command | 42

Benefits of the Probe command | 42

What is the Probe Command? | 42

Enabling the Probe command | 43

Using the Probe command | 43

5 Troubleshooting Network Issues

Working with Problems on Your Network | 50

Isolating a Broken Network Connection | 50

Identifying the Symptoms of a Broken Network Connection | 52

Isolating the Causes of a Network Problem | 53

Taking Appropriate Action for Resolving the Network Problem | 54

Evaluating the Solution to Check Whether the Network Problem Is Resolved | 56

6 Configuration Statements

address (Protocols Router Discovery) | 60

advertise | 61

broadcast | 62

disable (Protocols Router Discovery) | 63

dns-server-address | 64

extended-echo | 65

ignore | 65

ineligible | 66

interface (Protocols Router Discovery) | 67

lifetime (Router Advertisement) | 68

lifetime (Router Discovery) | 69

max-advertisement-interval (Protocols Router Discovery) | 70

min-advertisement-interval (Protocols Router Discovery) | 71

multicast (Protocols Router Discovery) | 72

priority (Protocols Router Discovery) | 73

router-discovery | 74

traceoptions (Protocols Router Discovery) | 76

7

Operational Commands

monitor interface | 80

monitor start | 96

monitor stop | 98

ping | 100

probe | 108

show log | 113

traceroute | 121

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | vi
- Using the Examples in This Manual | vi
- Documentation Conventions | viii
- Documentation Feedback | xi
- Requesting Technical Support | xi

Use this guide to configure, monitor and troubleshoot the ICMP Router Discovery protocol on Juniper Networks 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 ix 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 ix 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:

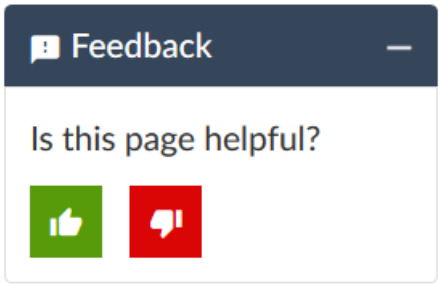

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. 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

Overview

ICMP Router Discovery Overview | **14**

Supported ICMP Router Discovery and IPv6 Neighbor Discovery Standards | **16**

ICMP Router Discovery Overview

IN THIS SECTION

- Operation of a Router Discovery Server | 14
- Router Advertisement Messages | 15

Router discovery uses Internet Control Message Protocol (ICMP) router advertisements and router solicitation messages to allow a host to discover the addresses of operational routers on the subnet. Hosts must discover routers before they can send IP datagrams outside their subnet. Router discovery allows a host to discover the addresses of operational routers on the subnet.

Router discovery uses Internet Control Message Protocol (ICMP) router advertisements and router solicitation messages to allow a host to discover the addresses of operational routers on the subnet. Hosts must discover routers before they can send IP datagrams outside their subnet.

Router discovery allows a host to discover the addresses of operational routers on the subnet. The Junos® operating system (Junos OS) implementation of router discovery supports server mode only.

Each router periodically multicasts a router advertisement from each of its multicast interfaces, announcing the IP address of that interface. Hosts listen for advertisements to discover the addresses of their neighboring routers. When a host starts, it can send a multicast 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 do not determine which router is best to reach a particular destination.

This section discusses the following topics:

Operation of a Router Discovery Server

The router discovery server distributes information about the addresses of all routers on directly connected networks and about their preferences for becoming the default router. (A host sends a packet to the default router if the host does not have a route to a destination in its routing table.) The server does this by periodically sending router advertisement packets out each interface on which router discovery is enabled. In addition to containing the router addresses, these packets also announce the existence of the server itself.

The server can either transmit broadcast or multicast router advertisement packets. Multicast packets are sent to 224.0.0.1, which is the all-hosts multicast address. When packets are sent to the all-hosts multicast address, or when an interface is configured for the limited-broadcast address 255.255.255.255, all IP addresses configured on the physical interface are included in the router advertisement. When the packets are being sent to a network or subnet broadcast address, only the address associated with that network or subnet is included in the router advertisement.

When the routing protocol process first starts on the server router, the server sends router advertisement packets every few seconds. Then, the server sends these packets less frequently, commonly every 10 minutes.

The server responds to router solicitation packets it receives from a client. The response is sent unicast unless a router advertisement packet is due to be sent out momentarily.

> **NOTE:** Junos OS does not support the ICMP router solicitation message with the source address as **0.0.0.0**.

Router Advertisement Messages

Router advertisement messages include a preference level and a lifetime field for each advertised router address.

The preference level specifies the router's preference to become the default router. When a host chooses a default router address, it chooses the address with the highest preference. You can configure the preference level by including the **priority** statement.

The lifetime field indicates the maximum length of time that the advertised addresses are to be considered valid by hosts in the absence of further advertisements. You can configure the advertising rate by including the **max-advertisement-interval** and **min-advertisement-interval** statements, and you can configure the lifetime by including the **lifetime** statement. .

RELATED DOCUMENTATION

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

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*

2

CHAPTER

Configuring the ICMP Protocol

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | **18**

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | **19**

Understanding the ICMP Protocol for Discovering Gateways to Other Networks

The ICMP Router Discovery Protocol (IRDP) enables hosts to locate routers on the local subnet and use them as a gateway to reach other networks. Junos OS supports running IRDP in server mode, meaning that router discovery packets are generated. Junos OS does not support IRDP in client mode running as a host sending router solicitation messages. IRDP is specified in RFC 1256, *ICMP Router Discovery Messages*.

For a host to participate on an internetwork, it needs connectivity to at least one router on the local network. One way to ensure that this is the case is to manually configure each host with the address of a local router as its default router (also called a *gateway*). This method is time-consuming to set up, difficult to maintain, and inflexible.

When you enable the Dynamic Host Configuration Protocol (DHCP) on a host, you do not need to configure the default router. DHCP uses a method called router discovery to automatically discover local routers, and learn other information about them.

The information provided includes the router's address (or addresses, if it has more than one) and how long the host should retain information about the router. Router advertisement messages are sent periodically. Hosts listen for these messages. When an advertisement is received, the host processes it and adds the information about the router to its routing table. A host that has no manually configured routing information has no connectivity to routers when it first powers on. Instead of waiting for the next Router Advertisement message, the host sends a router solicitation message on its local network. This prompts any router that receives this message to immediately send an extra router advertisement message directly to that host.

By default, router discovery is disabled on Junos OS routing devices. When router discovery is enabled, the default behavior is to advertise all interfaces. If the router supports multicast, all the IPv4 Layer 3 interfaces are advertised through multicast. Otherwise, all the IPv4 Layer 3 interfaces are advertised through broadcast.

RELATED DOCUMENTATION

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks

IN THIS SECTION

- Requirements | 19
- Overview | 19
- Configuration | 21
- Verification | 24

This example shows how to configure Internet Control Message Protocol (ICMP) router advertisements to allow IPv4 hosts to discover the addresses of operational routers on the subnet. Hosts must discover routers before they can send IP datagrams outside their subnet.

Requirements

This example assumes that a server or a client computer on the local network supports RFC 1256, *ICMP Router Discovery Messages*.

Overview

Before a host is able to send a message to a host outside its own subnet, it must be able to identify the address of the immediate router. This is typically done through reading a configuration file upon startup, and on some multicast networks by listening to routing protocol traffic. When a server or a client computer on the local network that supports RFC 1256 needs to locate a default gateway (router), the server or client computer uses ICMP to send a router solicitation. Hosts that support RFC 1256 send an ICMP router discovery message on the multicast address 224.0.0.2. Routers on the local network that support RFC 1256 immediately respond with a router advertisement.

The all-routers IP multicast address, 224.0.0.2, is the local IP broadcast address that IPv4 reserved. IPv4 multicast addresses in the range 224.0.0.0/24 (from 224.0.0.0 to 224.0.0.255) are reserved for the local subnet.

The ICMP Router Discovery Protocol (IRDP) uses router advertisements as well as router solicitation messages to allow hosts to learn the IP addresses of the router that is attached to the immediate network. When a host is started, it sends router solicitation messages to check for the address of the immediate router.

> **NOTE:** Not all hosts perform router discovery using the method specified in RFC 1256. If the host has DHCP enabled, it might not use ICMP router discovery. The performance of router discovery is one of the DHCP options that is defined in RFC 1541, *Dynamic Host Configuration Protocol*. This option specifies whether the client solicits routers using the ICMP router discovery method specified in RFC 1256. A value of 1 indicates that the client performs router discovery. A value of 0 indicates that the client does not.

To configure the router to be a router discovery server, you must include at least the following statement in the configuration. All other router discovery configuration statements are optional.

```
[edit]
protocols {
   router-discovery;
}
```

To configure a router as a server for ICMP router discovery, you can include the following statements in the configuration:

```
[edit]
protocols {
   router-discovery {
      disable;
      address address {
         (advertise | ignore);
         (broadcast | multicast);
         (ineligible | priority number);
      }
      interface interface-name {
         lifetime seconds;
         max-advertisement-interval seconds;
         min-advertisement-interval seconds;
      }
      traceoptions {
         file filename <files number> <size size> <world-readable | no-world-readable>;
```

```
            flag flag <detail> <disable>;
        }
    }
}
```

Figure 1 on page 21 shows a simplified sample topology.

Figure 1: ICMP Router Discovery Topology

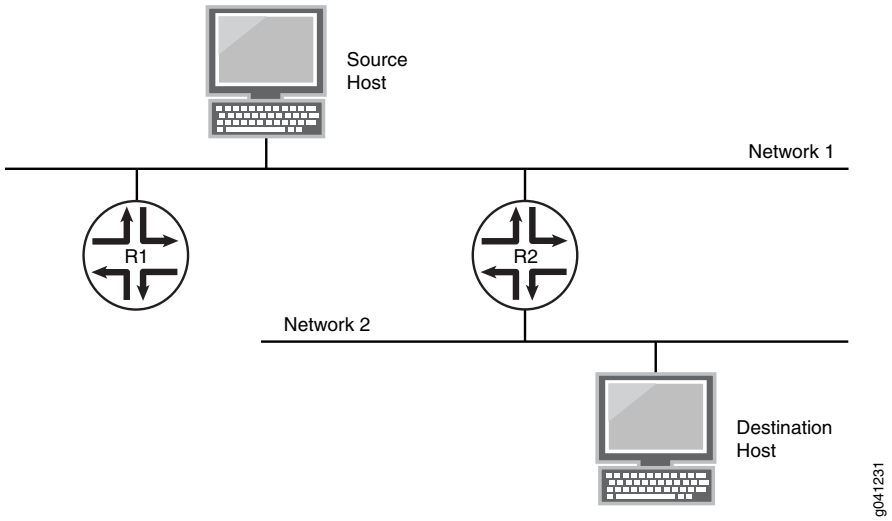

Configuration

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 interfaces ge-1/2/0 unit 6 description to-R2
set interfaces ge-1/2/0 unit 6 family inet address 10.0.0.6/24
set protocols router-discovery traceoptions file icmp-log
set protocols router-discovery traceoptions flag all
set protocols router-discovery interface ge-1/2/0.6 max-advertisement-interval 60
set protocols router-discovery interface ge-1/2/0.6 min-advertisement-interval 10
set protocols router-discovery interface ge-1/2/0.6 lifetime 120
set protocols router-discovery address 10.0.0.6 multicast
set protocols router-discovery address 10.0.0.6 priority 900
```

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 ICMP router discovery:

1. Configure the network interfaces.

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

   ```
   [edit interfaces]
   user@R1# set ge-1/2/0 unit 6 description to-R2
   user@R1# set ge-1/2/0 unit 6 family inet address 10.0.0.6/24
   ```

2. Enable router discovery.

   ```
   [edit protocols]
   user@R1# set router-discovery
   ```

3. (Optional) Enable trace operations for router discovery.

   ```
   [edit protocols router-discovery]
   user@R1# set traceoptions file icmp-log
   user@R1# set traceoptions flag all
   ```

4. (Optional) Set the IRDP maximum interval between advertisements.

   ```
   [edit protocols router-discovery]
   user@R1# set interface ge-1/2/0.6 max-advertisement-interval 60
   ```

5. (Optional) Set the IRDP minimum interval between advertisements.

   ```
   [edit protocols router-discovery]
   user@R1# set interface ge-1/2/0.6 min-advertisement-interval 10
   ```

6. (Optional) Set the IRDP period for which advertisements are valid.

   ```
   [edit protocols router-discovery]
   user@R1# set interface ge-1/2/0.6 lifetime 120
   ```

7. (Optional) Configure the router to include the 10.0.0.6 IP address in IRDP advertisements to the all-hosts multicast address (224.0.0.1).

   If the router supports IP multicast, and if the interface supports IP multicast, **multicast** is the default. Otherwise, the addresses are included in broadcast router advertisement packets.

   ```
   [edit protocols router-discovery]
   user@R1# set address 10.0.0.6 multicast
   ```

8. (Optional) Set the preference of the address to become a default router.

   This preference is set relative to the preferences of other router addresses on the same subnet.

   ```
   [edit protocols router-discovery]
   user@R1# set address 10.0.0.6 priority 900
   ```

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
ge-1/2/0 {
    unit 6 {
        description to-R2;
        family inet {
            address 10.0.0.6/24;
        }
    }
}
```

```
user@R1# show protocols
router-discovery {
    traceoptions {
        file icmp-log;
        flag all;
    }
    interface ge-1/2/0.6 {
        max-advertisement-interval 60;
        min-advertisement-interval 10;
        lifetime 120;
    }
```

```
    address 10.0.0.6 {
      multicast;
      priority 900;
    }
  }
```

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

Verification

Confirm that the configuration is working properly.

Checking the Trace Log

Purpose

Verify that the expected interfaces are sending messages.

Action

From operational mode, enter the **show log icmp-log** command.

```
user@R1> show log icmp-log
```

```
Mar 21 14:42:54 trace_on: Tracing to "/var/log/icmp-log" started
Mar 21 14:42:54.409027 rdisc_ifa_change: Preference for address 10.0.0.6(ge-1/2/0.6)
 set to 900
Mar 21 14:43:33.983695 task_timer_uset: timer RouterDiscoveryServer_Group <Touched
 Processing> set to offset 22 at 14:43:16
Mar 21 14:43:33.984263 rdisc_server_timer: group ge-1/2/0.6 timer set to 22
Mar 21 14:43:55.985225 task_timer_uset: timer RouterDiscoveryServer_Group <Touched
 Processing> set to offset 37 at 14:44:10
Mar 21 14:43:55.985520 rdisc_server_timer: group ge-1/2/0.6 timer set to 37
Mar 21 14:44:32.986407 task_timer_uset: timer RouterDiscoveryServer_Group <Touched
 Processing> set to offset 39 at 14:44:44
Mar 21 14:44:32.986961 rdisc_server_timer: group ge-1/2/0.6 timer set to 39
Mar 21 14:45:11.987331 task_timer_uset: timer RouterDiscoveryServer_Group <Touched
 Processing> set to offset 10 at 14:44:42
Mar 21 14:45:11.987888 rdisc_server_timer: group ge-1/2/0.6 timer set to 10
Mar 21 14:45:21.990974 task_timer_uset: timer RouterDiscoveryServer_Group <Touched
 Processing> set to offset 23 at 14:45:34
Mar 21 14:45:21.991548 rdisc_server_timer: group ge-1/2/0.6 timer set to 23
Mar 21 14:45:44.992150 task_timer_uset: timer RouterDiscoveryServer_Group <Touched
```

```
 Processing> set to offset 45 at 14:46:06
Mar 21 14:45:44.992710 rdisc_server_timer: group ge-1/2/0.6 timer set to 45
```

Meaning

The log output shows that the preference was set to 900 for IP address 10.0.0.6 and that messages are being sent on the ge-1/2/0.6 interface.

RELATED DOCUMENTATION

ICMP Router Discovery Overview | 14

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

3

CHAPTER

Configuring Recursive DNS Servers

Understanding Recursive DNS Servers for IPv6 | **27**

Configuring a Recursive DNS Server Address for IPv6 Hosts | **28**

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts | **29**

Understanding Recursive DNS Servers for IPv6

To access any location on the Internet, the domain name system (DNS) server plays a pivotal role in resolving the domain name into its associated IP address. The DNS resolution service can also be provided by the DHCP server. The routing protocol process (rpd) of routers generates router advertisements to facilitate IPv6 hosts in autoconfiguration and in learning network information. For IPv6 stateless autoconfiguration, DNS configuration is provided by router advertisements. The router advertisement-based DNS configuration is useful in networks where an IPv6 host's address is autoconfigured through an IPv6 stateless address and where there is no existing DHCPv6 infrastructure.

Depending on their configuration, DNS servers can be classified into the following types:

- Recursive domain name system
- Nonrecursive domain name system

DNS servers can resolve either recursive or nonrecursive queries. For a recursive query by a DNS client, the DNS server returns either the IP address associated with the domain name or an error. A recursive query does not return a referral. For a nonrecursive query, the DNS server returns the IP address of the domain name or an error or a referral to another DNS server which might have the resolution of the query.

For IPv6 hosts, a maximum of three recursive DNS server addresses can be configured along with their respective lifetimes. The default value of the lifetime of the configured recursive DNS server addresses is 1800 seconds. The configured IPv6 host uses the specified recursive DNS server address for DNS resolution where the IPv6 host's address is autoconfigured through an IPv6 stateless address and where there is no DHCPv6 infrastructure available.

CAUTION: The recursive DNS server configuration is included in the router advertisement packet, which is a part of the Neighbor Discovery Protocol (NDP). In general, in an unsecured deployment scenario, an attacker could send a router advertisement with a fraudulent recursive DNS server address, misleading the IPv6 host into contacting an unintended DNS server for DNS name resolution. These attacks are similar to neighbor discovery attacks and attacks against unauthenticated DHCP. We recommend using the Secure Neighbor Discovery (SEND) protocol as a security mechanism for neighbor discovery to allow all the neighbor discovery options including the recursive DNS server options to be automatically included in the signatures.

For more information about configuring the SEND protocol, see www.juniper.net/documentation/en_US/junos14.1/topics/topic-map/ipv6-secure-neighbor.html

RELATED DOCUMENTATION

dns-server-address | 64

lifetime | 68

Configuring a Recursive DNS Server Address for IPv6 Hosts | 28

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts | 29

Configuring a Recursive DNS Server Address for IPv6 Hosts

For IPv6 hosts, a maximum of three recursive DNS server addresses can be configured along with their respective lifetimes. The configured IPv6 host uses the specified recursive DNS server address for DNS resolution where the IPv6 host's address is autoconfigured through an IPv6 stateless address and where there is no DHCPv6 infrastructure available.

To configure a recursive DNS server address on IPv6 hosts, follow these steps:

1. Configure the recursive DNS server address for the IPv6 host.

   ```
   [edit protocols router-advertisement]
   user@host# set interfaces interface name dns-server-address address
   ```

   For example, to assign IPv6 address abcd:1::1 as the recursive dns server address to interface fe-1/0/1:

   ```
   [edit protocols router-advertisement]
   user@host# set interfaces fe-1/0/1 dns-server-address abcd:1::1
   ```

2. Configure the lifetime to specify the time in seconds for which the recursive DNS server address remains valid.

   ```
   [edit protocols router-advertisement interfaces interface name dns-server-address address]
   user@host# set lifetime seconds
   ```

   For example, to specify a lifetime of 60 seconds for the recursive DNS server address:

   ```
   [edit protocols router-advertisement interfaces interface name dns-server-address address]
   user@host# set lifetime 60
   ```

   The default value of the lifetime of the configured recursive DNS server address is 1800 seconds.

RELATED DOCUMENTATION

dns-server-address | 64
lifetime | 68
Understanding Recursive DNS Servers for IPv6 | 27
Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts | 29

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts

IN THIS SECTION

Requirements | 29
Overview | 30
Configuration | 30
Verification | 33

This example shows how to configure the recursive DNS server address of an IPv6 host. The recursive DNS server address is included in the router advertisement that is sent to the neighboring devices.

Requirements

This example uses the following hardware and software components:

- Two MX Series routers with IPv6 enabled on the connected interfaces.
- Junos OS Release 14.1 or later running on all devices.

Overview

The example includes two routers that are directly connected. Configure IPv6 on the directly connected interfaces. Enable router advertisement on the interfaces and configure the recursive DNS server addresses and their lifetimes on the interfaces. This example verifies that the router advertisement sent to the neighboring device includes the configured recursive DNS server addresses.

Topology

Figure 2 on page 30 shows the sample topology.

Figure 2: Configuring Recursive DNS Server Address for IPv6 Hosts

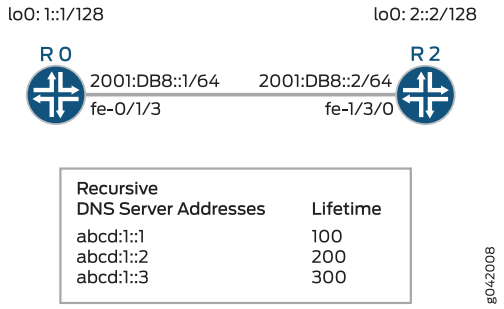

Configuration

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.

Router R0

```
set interfaces fe-0/1/3 unit 0 family inet6 address 2001:DB8::1/64
set interfaces lo0 unit 0 family inet6 address 1::1/128
set protocols router-advertisement interface fe-0/1/3 max-advertisement-interval 4
set protocols router-advertisement interface fe-0/1/3 min-advertisement-interval 3
set protocols router-advertisement interface fe-0/1/3 dns-server-address abcd:1::1 lifetime 100
set protocols router-advertisement interface fe-0/1/3 dns-server-address abcd:1::2 lifetime 200
set protocols router-advertisement interface fe-0/1/3 dns-server-address abcd:1::3 lifetime 300
```

Router R1

```
set interfaces fe-1/3/0 unit 0 family inet6 address 2001:DB8::2/64
set interfaces lo0 unit 0 family inet6 address 1::2/128
set protocols router-advertisement interface fe-1/3/0 max-advertisement-interval 4
set protocols router-advertisement interface fe-1/3/0 min-advertisement-interval 3
set protocols router-advertisement interface fe-1/3/0 dns-server-address abcd:1::4 lifetime 100
```

Configuring the Recursive DNS Server Address

Step-by-Step Procedure

The following example requires that you 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*.

NOTE: Repeat this procedure for Router R1, modifying the appropriate interface names, addresses, and any other parameters for the router.

To configure the recursive DNS server address on Router R0:

1. Enable IPv6 on the physical interface.

   ```
   [edit interfaces]
   user@R0# set fe-0/1/3 unit 0 family inet6 address 2001:DB8::1/64
   ```

2. Configure the loopback address.

   ```
   [edit interfaces]
   user@R0# set lo0 unit 0 family inet6 address 1::1/128
   ```

3. Specify the time interval between router advertisements on the interface.

   The router sends advertisements to neighbors after the specified time interval. In this example, Router R0 sends router advertisements to Router R1 after a minimum interval of 3 seconds and a maximum interval of 4 seconds.

   ```
   [edit protocols router-advertisement]
   user@R0# set interface fe-0/1/3 max-advertisement-interval 4
   ```

```
user@R0# set interface fe-0/1/3 min-advertisement-interval 3
```

4. Configure the recursive DNS addresses and their lifetimes on the interface.

```
[edit protocols router-advertisement]
user@R0# set interface fe-0/1/3 dns-server-address abcd:1::1 lifetime 100
user@R0# set interface fe-0/1/3 dns-server-address abcd:1::2 lifetime 200
user@R0# set interface fe-0/1/3 dns-server-address abcd:1::3 lifetime 300
```

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 instructions in this example to correct the configuration.

```
user@R0# show interfaces
fe-0/1/3 {
    unit 0 {
        family inet6 {
            address 2001:DB8::1/64;
        }
    }
}
lo0 {
    unit 0 {
        family inet6 {
            address ::1/128;
        }
    }
}
user@R0# show protocols
router-advertisement {
    interface fe-0/1/3.0 {
        max-advertisement-interval 4;
        min-advertisement-interval 3;
        dns-server-address abcd:1::1 {
            lifetime 100;
        }
        dns-server-address abcd:1::2 {
            lifetime 200;
        }
```

```
        dns-server-address abcd:1::3 {
            lifetime 300;
        }
    }
}
```

If you are done configuring the device, commit the configuration.

```
user@R0# commit
```

Verification

Verifying That the Router Advertisement Includes the Recursive DNS Server Address

Purpose

Verify that the router advertisement on Router R1 includes the recursive DNS server address configured on Router R0.

Action

From operational mode on Router R1, enter the **show ipv6 router-advertisement** command.

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

```
Interface: fe-1/3/0.0
  Advertisements sent: 18, last sent 00:00:02 ago
  Solicits received: 0
  Advertisements received: 18
  Advertisement from fe80::214:f6ff:fe22:5422, heard 00:00:02 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 12 sec
    Retransmit timer: 0 ms
    Current hop limit: 64
    RDNSS address: abcd:1::1
      Lifetime: 100 sec
    RDNSS address: abcd:1::2
      Lifetime: 200 sec
```

```
RDNSS address: abcd:1::3
    Lifetime: 300 sec
```

Meaning

The recursive DNS server address and the configured lifetime are included in the router advertisements on Router R1.

RELATED DOCUMENTATION

- Configuring a Recursive DNS Server Address for IPv6 Hosts | 28
- Understanding Recursive DNS Servers for IPv6 | 27
- dns-server-address | 64
- lifetime | 68

4

CHAPTER

Monitoring ICMP Router Discovery

Traceroute for Inactive Interface | 36

Example: Tracing Global Routing Protocol Operations | 36

How to Use the Probe Command | 42

Traceroute for Inactive Interface

Traceroute is a tool for displaying the route taken by a packet from an IP network on their way to a given host. When a traceroute is performed the packets are always sent out of the interface that is the NH for the active route and there is no option to bypass it.

When a traceroute is performed, packets are sent out of active interface even if we specify an inactive interface. From Junos OS Release 17.4R1 onwards, you can configure traceroute to send out packets through an inactive next-hop by specifying the **traceroute *next-hop address*** to a destination through an inactive next hop.

RELATED DOCUMENTATION

traceroute | **121**

Example: Tracing Global Routing Protocol Operations

IN THIS SECTION

- Requirements | **36**
- Overview | **37**
- Configuration | **37**
- Verification | **41**

This example shows how to list and view files that are created when you enable global routing trace operations.

Requirements

You must have the **view** privilege.

Overview

To configure global routing protocol tracing, include the **traceoptions** statement at the **[edit routing-options]** hierarchy level:

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

The flags in a **traceoptions flag** statement are identifiers. When you use the **set** command to configure a flag, any flags that might already be set are not modified. In the following example, setting the **timer** tracing flag has no effect on the already configured **task** flag. Use the **delete** command to delete a particular flag.

```
[edit routing-options traceoptions]
user@host# show
flag task;
user@host# set traceoptions flag timer
user@host# show
flag task;
flag timer;
user@host# delete traceoptions flag task
user@host# show
flag timer;
```

This example shows how to configure and view a trace file that tracks changes in the routing table. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.

TIP: To view a list of hierarchy levels that support tracing operations, enter the **help apropos traceoptions** command in configuration mode.

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 routing-options traceoptions file routing-table-changes
set routing-options traceoptions file size 10m
set routing-options traceoptions file files 10
set routing-options traceoptions flag route
set routing-options static route 1.1.1.2/32 next-hop 10.0.45.6
```

Configuring Trace Operations

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 the trace operations:

1. Configure trace operations.

```
[edit routing-options traceoptions]
user@host# set file routing-table-changes
user@host# set file size 10m
user@host# set file files 10
user@host# set flag route
```

2. Configure a static route to cause a change in the routing table.

```
[edit routing-options static]
user@host# set route 1.1.1.2/32 next-hop 10.0.45.6
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Viewing the Trace File

Step-by-Step Procedure

To view the trace file:

1. In operational mode, list the log files on the system.

   user@host> **file list /var/log**

   ```
   /var/log:
   ...
   routing-table-changes
   ...
   ```

2. View the contents of the **routing-table-changes** file.

   user@host> **file show /var/log/routing-table-changes**

   ```
   Dec 15 11:09:29 trace_on: Tracing to "/var/log/routing-table-changes" started
   Dec 15 11:09:29.496507
   Dec 15 11:09:29.496507 Tracing flags enabled: route
   Dec 15 11:09:29.496507
   Dec 15 11:09:29.533203 inet_routerid_notify: Router ID: 192.168.4.1
   Dec 15 11:09:29.533334 inet_routerid_notify: No Router ID assigned
   Dec 15 11:09:29.533381 inet_routerid_notify: No Router ID assigned
   Dec 15 11:09:29.533420 inet_routerid_notify: No Router ID assigned
   Dec 15 11:09:29.534915 inet_routerid_notify: Router ID: 192.168.4.1
   Dec 15 11:09:29.542934 inet_routerid_notify: No Router ID assigned
   Dec 15 11:09:29.549253 inet_routerid_notify: No Router ID assigned
   Dec 15 11:09:29.556878 inet_routerid_notify: No Router ID assigned
   Dec 15 11:09:29.582990 rt_static_reinit: examined 3 static nexthops, 0
   unreferenced
   Dec 15 11:09:29.589920
   Dec 15 11:09:29.589920 task_reconfigure reinitializing done
   ...
   ```

3. Filter the output of the log file.

   user@host> **file show /var/log/routing-table-changes | match 1.1.1.2**

   ```
   Dec 15 11:15:30.780314 ADD      1.1.1.2/32          nhid 0 gw 10.0.45.6
   Static   pref 5/0 metric  at-0/2/0.0 <ctive Int Ext>
   Dec 15 11:15:30.782276 KRT Request: send len 216 v104 seq 0 ADD route/user af
   2 table 0 infot 0 addr 1.1.1.2 nhop-type unicast nhindex 663
   ```

4. View the tracing operations in real time by running the **monitor start** command with an optional **match** condition.

```
user@host> monitor start routing-table-changes | match 1.1.1.2
```

```
Aug 10 19:21:40.773467 BGP RECV         0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlri: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlri: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlri: 0.0.0.0/0 qualified bnp->ribact 0x0 l2afcb
 0x0
```

5. Deactivate the static route.

```
user@host# deactivate routing-options static route 1.1.1.2/32
user@host# commit
```

```
*** routing-table-changes ***
Dec 15 11:42:59.355557 CHANGE   1.1.1.2/32          nhid 663 gw 10.0.45.6
  Static   pref 5/0 metric  at-0/2/0.0 <Delete Int Ext>
Dec 15 11:42:59.426887 KRT Request: send len 216 v104 seq 0 DELETE route/user
af 2 table 0 infot 0 addr 1.1.1.2 nhop-type discard filtidx 0
Dec 15 11:42:59.427366 RELEASE  1.1.1.2/32          nhid 663 gw 10.0.45.6
  Static   pref 5/0 metric  at-0/2/0.0 <Release Delete Int Ext>
```

6. Halt the **monitor** command by pressing Enter and typing **monitor stop**.

```
[Enter]
user@host> monitor stop
```

7. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

   When configuration is deactivated, it appears in the configuration with the **inactive** tag.

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

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

```
inactive: traceoptions {
    file routing-table-changes size 10m files 10;
    flag route;
```

```
}
static {
    inactive: route 1.1.1.2/32 next-hop 10.0.45.6;
}
```

8. To reactivate trace operations, use the **activate** configuration-mode statement.

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

Results

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

```
user@host# show routing-options
traceoptions {
    file routing-table-changes size 10m files 10;
    flag route;
}
static {
    route 1.1.1.2/32 next-hop 10.0.45.6;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the Trace Log File Is Operating

Purpose

Make sure that events are being written to the log file.

Action

```
user@host> show log routing-table-changes
```

```
Dec 15 11:09:29 trace_on: Tracing to "/var/log/routing-table-changes" started
```

How to Use the Probe Command

SUMMARY

Learn how to configure and use the Probe command.

IN THIS SECTION

- Benefits of the Probe command | 42
- What is the Probe Command? | 42

Benefits of the Probe command

The Probe command helps to capture the interface details such as probe packet statistics, and interface state (active/inactive), irrespective of whether the network family address configured is IPv4 or IPv6 on the probed interfaces.

What is the Probe Command?

RFC 8335, *PROBE: A Utility for Probing Interfaces*, describes the Probe utility. Probe is a network diagnostic tool similar to Ping that can be used to query the status of a probed interface on a node. The Probe command does not require bidirectional connectivity between the probing and the probed node. You can also have a proxy node between the probing and the probed node. Alternatively, the proxy node can reside on the same node as the probed node to which the probed interface is directly connected. Probe uses ICMP Extended Echo/Reply messages for communication between the probing interface and the proxy interface. The Probe command is used to probe the details of the interfaces on the probed node that is directly connected to the proxy, where the proxy learns the details from the ARP or NDP entries of those interfaces. This utility helps in scenarios where bidirectional connectivity between the probing and probed interfaces is lacking. For example, if the probed interface is an unnumbered interface or if the probed interface is in a different network family.

The following figure shows R0 as the probing node, R1 as the proxy node, and R2 as the probed node:

Figure 3: Probing, Proxy, and Probed Nodes

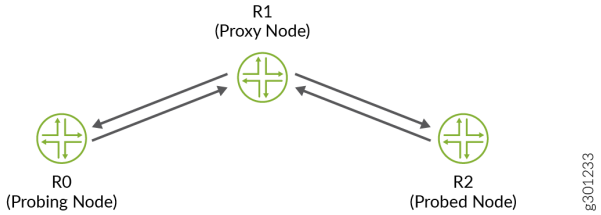

Enabling the Probe command

To enable the **probe** command, configure the **extended-echo** configuration statement at the **[edit system]** hierarchy level on the proxy node and the probed node.

```
[edit]
user@host# set system extended-echo
```

NOTE: You do not need to enable the **extended-echo** configuration statement on the probing node.

Using the Probe command

You can probe using three different options:

- By using IP address (the IP addresses that are configured in the probed node, the proxy node, and the remote node).
- By using interface names on the proxy node or the probed node.
- By using the interface index.

The following examples shows how to use the **probe** command. Consider R0 as the probing node, R1 as the proxy node, and R2 as the probed node and change the details of the interfaces and IP addresses to match your network configuration.

Purpose

Use case 1: Probing when proxy node (R1) and the probed node (R2) are two different nodes.

Query for the status of the probed IP address through the proxy IP address.

Action

From operational mode, probe for the status of the IP address 198.52.102.36 of the probed node (R2) using the **by-address** option with the proxy IP address 198.52.100.1 of the proxy node (R1). Count 1 indicates that the number of probe requests sent is 1.

user@R0>**probe 198.52.100.1 by-address 198.52.102.36 count 1**

```
PROBE 198.52.100.1 (198.52.100.1):
35 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=2 state=0 active=0 IPv4=0
IPv6=0 time=5.195 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 5.195/5.195/5.195/0.000 ms
```

Meaning

The proxy IP address, 198.52.100.1 displays the status of the probed IP address, 198.52.102.36. The output is verified with the following probe packet statistics:

- **code=2** indicates that the sent packet is not received by the probed node.
- **state=0** indicates that at the time of probing, the state of the IPv6 or IPv4 node is reachable.
- **active=0** indicates that the probed node is inactive as the IPv4 or IPv6 address is not available or configured on the probed node. That is, the probed IP address, 198.52.102.36 is not configured on the probed node (R2).
- **IPv4=0** and **IPv6=0** indicates that the IPv4 and IPv6 addresses are not configured.
- **time=*milliseconds* ms** indicates the time taken to transmit the packets.
- **ttl** is the number of hops/devices to cross.

Purpose

Use Case 2: Probing when proxy node and the probed node are the same.

Query for the status of the probed IP address through the proxy IP address (here, the probed IP address and the proxy IP address are configured on the same node, R1)

Action

From operational mode, probe for the status of the probed IP address 198.52.102.2 using the **by-address** option with the proxy IP address 198.52.100.1 at the proxy node (R1). Count 1 indicates that the number of probe requests sent is 1.

user@R0>**probe 198.52.100.1 by-address 198.52.102.2 count 1**

```
PROBE 198.52.100.1 (198.52.100.1):
32 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=1 time=4.908 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 5.195/5.195/5.195/0.000 ms
```

Meaning

The proxy IP address, 198.52.100.1 displays the status of the probed IP address, 198.52.102.2. The output is verified with the following probe packet statistics:

- **code=0** indicates that the sent packet is received by the probed node. The code value will be 0 as proxy and probed nodes are the same.
- **state=0** indicates that at the time of probing the state of the IPv6 or IPv4 node is reachable.
- **active=1** indicates that the probed node is active and either the IPv4 or IPv6 addresses or both are configured and available on the probed node.
- **IPv4=1** and **IPv6=1** indicates that the IPv4 and IPv6 addresses are configured and available on the probed node.
- **time=*milliseconds* ms** indicates the time taken to transmit the packets.
- **ttl** is the number of hops/devices to cross.

Purpose

Query for the status of the IPv4 or IPv6 address using the interface index of the probed interface.

Action

From operational mode, probe for the status of the IPv4 or IPv6 address using the **by-index** option to specify the interface index of the probed interface.

user@R0>**probe 198.52.100.1 by-index 1037**

```
PROBE 198.52.100.1 (198.52.100.1):
28 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.871 ms
28 bytes from 198.52.100.1: icmp_seq=1 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.864 ms


--- 198.52.100.1 probe statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.864/0.867/0.871/0.004 ms
```

Meaning

The proxy IP address, 198.52.100.1 displays the status of the probed interface using the interface index-id, 1037. The output is verified with the following probe packet statistics:

- **code=0** indicates that the sent packet is received by the probed node.
- **state=0** indicates that at the time of probing the state of the IPv6 or IPv4 node is reachable.
- **active=1** indicates that the probed node is active and either the IPv4 or the IPv6 address is available.
- **IPv4=1** indicates that the IPv4 address is configured and available. **IPv6=0** indicates that the IPv6 address is not configured.
- **time=*milliseconds* ms** indicates the time taken to transmit the packets.
- **ttl** is the number of hops/devices to cross.

Purpose

Query the proxy node (R1) for information about the remote node using the remote address.

Action

From operational mode, probe for the status of the probed node and verify using the **by-remote-address** option with the IP address of a remote device. Count 1 indicates that the number of probe requests sent is 1.

user@R0>**probe 198.52.100.1 by-remote-address 2001:65::2 count 1**

```
PROBE 198.52.100.1 (198.52.100.1):
44 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=2 active=0 IPv4=0
IPv6=0 time=6.691 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 6.691/6.691/6.691/0.000 ms
```

Meaning

The proxy IP address, 198.52.100.1 displays the status of the probed node using the IP address, 2001:65::2 of the remote device. The **by-remote-address** option is used to determine the status of the remote node only. The output is verified with the following probe packet statistics:

- **code=0** indicates that the sent packet is received by the probed node.
- **state=2** indicates that the connection is established with a remote node where the IPv4 and the IPv6 address is not available.
- **active=0** indicates that the probed node is inactive.
- **IPv4=0** and **IPv6=0** indicates that the IPv4 and the IPv6 addresses are not configured or available. Note that when you probe using the **by-remote-address** option, the IPv4 or the IPv6 address is not available.
- **time=*milliseconds* ms** indicates the time taken to transmit the packets.
- **ttl** is the number of hops/devices to cross.

Purpose

Query for the status of the probed node using the interface name.

Action

From operational mode, probe for the status of the probed node using the interface name ge-1/2/1.1 and verify the output through the IP address of the proxy node 198.52.100.1. Count 1 indicates that the number of probe requests sent is 1.

user@R0>**probe 198.52.100.1 by-name ge-1/2/1.1 count 1**

```
PROBE 198.52.100.1 (198.52.100.1):
34 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.807 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.807/0.807/0.807/0.000 ms
```

Meaning

The proxy IP address, 198.52.100.1 displays the status of the device with reference to the interface ge-1/2/1.1 of the device on the probed node. The output is verified with the following probe packet statistics:

- **code=0** indicates that the sent packet is received by the probed node.
- **state=0** indicates that the proxy is able to establish a connection with the probed node.
- **active=1** indicates that the probed node is active and either the IPv4 or the IPv6 address configured is available.
- **IPv4=1** indicates that the IPv4 address is configured and available. **IPv6=0** indicates that the IPv6 address is not configured.
- **time=*milliseconds* ms** indicates the time taken to transmit the packets.
- **ttl** is the number of hops/devices to cross.

SEE ALSO

extended-echo | **65**

probe | **108**

5
CHAPTER

Troubleshooting Network Issues

Working with Problems on Your Network | **50**

Isolating a Broken Network Connection | **50**

Identifying the Symptoms of a Broken Network Connection | **52**

Isolating the Causes of a Network Problem | **53**

Taking Appropriate Action for Resolving the Network Problem | **54**

Evaluating the Solution to Check Whether the Network Problem Is Resolved | **56**

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 50	
1. Identifying the Symptoms of a Broken Network Connection on page 52	**ping (*ip-address* | *hostname*)** **show route (*ip-address* | *hostname*)** **traceroute (*ip-address* | *hostname*)**
2. Isolating the Causes of a Network Problem on page 53	**show < configuration | interfaces | protocols | route >**
3. Taking Appropriate Action for Resolving the Network Problem on page 54	**[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 56	**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 4 on page 50, 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 4: Process for Diagnosing Problems in Your Network

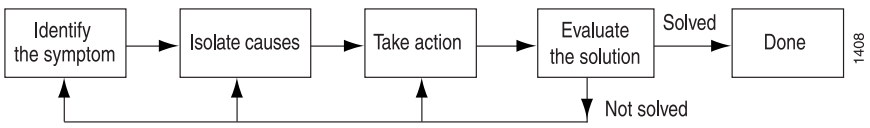

Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic 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 5 on page 51 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 5: Network with a Problem

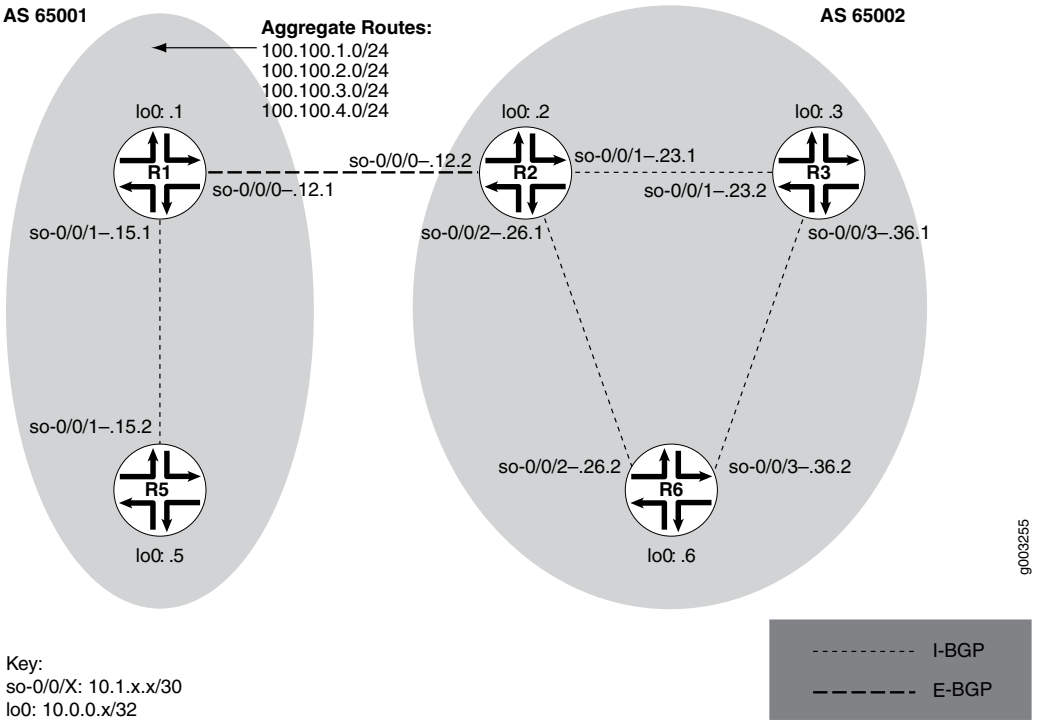

The network in Figure 5 on page 51 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (**R1**) in AS 65001 announces aggregated prefixes **100.100/24** to the AS 65002 network. The problem in this network is that **R6** does 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 53
- Taking Appropriate Action for Resolving the Network Problem on page 54

- Taking Appropriate Action for Resolving the Network Problem on page 54
- Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 56

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 50, 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**).

CHAPTER 6

Configuration Statements

address (Protocols Router Discovery) | **60**

advertise | **61**

broadcast | **62**

disable (Protocols Router Discovery) | **63**

dns-server-address | **64**

extended-echo | **65**

ignore | **65**

ineligible | **66**

interface (Protocols Router Discovery) | **67**

lifetime (Router Advertisement) | **68**

lifetime (Router Discovery) | **69**

max-advertisement-interval (Protocols Router Discovery) | **70**

min-advertisement-interval (Protocols Router Discovery) | **71**

multicast (Protocols Router Discovery) | **72**

priority (Protocols Router Discovery) | **73**

router-discovery | 74

traceoptions (Protocols Router Discovery) | 76

address (Protocols Router Discovery)

Syntax

```
address (Protocols Router Discovery) address {
    (advertise | ignore);
    (broadcast | multicast);
    (ineligible | priority number);
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the IP addresses to include in router advertisement packets.

Options

address—IP address. To specify more than one address, specify multiple addresses or include multiple **address** statements.

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

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

advertise

Syntax

```
(advertise | ignore);
```

Hierarchy Level

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Specify whether the server should advertise the IP address in its router advertisement packets:

- **advertise**—Advertise the IP address in its router advertisement packets.
- **ignore**—Do not advertise the IP addresses in router advertisement packets.

Default
advertise

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

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18
Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

broadcast

Syntax

```
(broadcast | multicast);
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify when the server should include the IP addresses in router advertisement packets. On the same physical interfaces, some addresses might be included only in multicast packets, while others might be included only in broadcast packets.

If you specify **broadcast**, the server includes the addresses in router advertisement packets only if the packets are broadcast.

Default

multicast if the router supports IP multicast; **broadcast** if the router does not support IP multicast.

Required Privilege Level

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

RELATED DOCUMENTATION

- Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18
- Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19
- multicast | 72

disable (Protocols Router Discovery)

Syntax

```
disable;
```

Hierarchy Level

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Disable router discovery.

Default
The configured object is enabled (operational) 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

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | **18**
Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | **19**

dns-server-address

Syntax

```
dns-server-address address {
    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 in Junos OS Release 14.1.

Description

Specify the recursive DNS server address that the device must use to resolve DNS names. The recursive DNS server address is the 128-bit IPv6 address of the recursive DNS server. You can configure a maximum of three recursive DNS server addresses at the interface level.

The remaining statement is 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 Recursive DNS Server Addresses for IPv6 Hosts | 29
- Understanding Recursive DNS Servers for IPv6 | 27
- lifetime (Router Advertisement) | 68

extended-echo

Syntax

```
extended-echo;
```

Hierarchy Level

```
[edit system]
```

Release Information
Statement introduced in Junos OS Release 20.3R1.

Description
Configure ICMP extended echo to enable the **probe** command. On enabling this statement on the proxy node and the probed node, you can query for the status of the probed node. The ICMP extended echo and ICMP extended reply messages are used for the communication between the probing interface and the proxy interface.

Default
Extended echo is disabled unless explicitly enabled.

Required Privilege Level
system-To view this statement in the configuration.

RELATED DOCUMENTATION

What is the Probe Command? | 42

probe | 108

ignore

See
advertise

ineligible

Syntax

```
ineligible;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify that the address can never become the default router.

Required Privilege Level

routing—To view this statement in the configuration.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

priority | 73

interface (Protocols Router Discovery)

Syntax

```
interface interface-name {
    lifetime seconds;
    max-advertisement-interval seconds;
    min-advertisement-interval seconds;
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify physical interfaces on which to configure timers for router advertisement messages.

Options

interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, specify **all**.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

lifetime (Router Advertisement)

Syntax

```
lifetime seconds;
```

Hierarchy Level

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

Release Information

Statement introduced in Junos OS Release 14.1.

Description

Specify the maximum time in seconds for which the recursive DNS server address remains valid. The device can use the specified recursive DNS server address for DNS name resolution until the time specified by this statement.

seconds– Maximum time for which the recursive DNS server address remains valid.

Options

Range: 0 through 4294967295 seconds

Default: 1800 seconds

Values: 0 indicates that the advertised recursive DNS server address is no longer valid and that this recursive DNS server address entry can be deleted. 4294967295 seconds indicates an infinite lifetime and a persistent entry in the device for this recursive DNS server address.

Required Privilege Level

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

RELATED DOCUMENTATION

Understanding Recursive DNS Servers for IPv6 | 27

Example: Configuring Recursive DNS Server Addresses for IPv6 Hosts | 29

dns-server-address | 64

lifetime (Router Discovery)

Syntax

```
lifetime seconds;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify how long the addresses sent by the server in its router advertisement packets are valid. This time must be long enough so that another router advertisement packet is sent before the lifetime has expired. The lifetime value is placed in the advertisement lifetime field of the router advertisement packet. If this amount of time passes and the host has not received a router advertisement from the server, the router marks the advertised addresses as invalid.

Options

seconds—Lifetime value. A value of 0 indicates that one or more addresses are no longer valid.

Range: Three times the value set by the **max-advertisement-interval** statement through 2 hours, 30 minutes (9000 seconds)

Default: 1800 seconds (30 minutes, which is three times the default value for the **max-advertisement-interval** statement)

Required Privilege Level

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

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | **18**

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | **19**

max-advertisement-interval | **70**

max-advertisement-interval (Protocols Router Discovery)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the maximum time the router waits before sending periodic router advertisement packets out the interface. These packets are broadcast or multicast, depending on how the address corresponding to this physical interface is configured.

Options

seconds—Maximum time between router advertisement packets.

Range: 4 through 1800 seconds

Default: 600 seconds (10 minutes)

Required Privilege Level

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

RELATED DOCUMENTATION

- Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18
- Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19
- broadcast | 62
- lifetime | 69
- min-advertisement-interval | 71
- multicast | 72

min-advertisement-interval (Protocols Router Discovery)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the minimum time the router waits before sending router advertisement packets out the interface in response to router solicitation packets it receives from a client. These packets are broadcast or multicast, depending on how the address corresponding to this physical interface is configured.

Options

seconds—Minimum time between router advertisement packets.

Range: 3 seconds through 1800 seconds

Default: 400 seconds (0.75 times the default value for the **max-advertisement-interval** statement)

Required Privilege Level

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

RELATED DOCUMENTATION

- Understanding the ICMP Protocol for Discovering Gateways to Other Networks | **18**
- Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | **19**
- broadcast | **62**
- max-advertisement-interval | **70**
- multicast | **72**

multicast (Protocols Router Discovery)

Syntax

```
(multicast | broadcast);
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify when the server should include the IP addresses in router advertisement packets. On the same physical interfaces, some addresses might be included only in multicast packets, while others might be included only in broadcast packets.

If you specify **multicast**, the server includes the addresses in router advertisement packets only if the packets are multicast. If the router supports IP multicast, and if the interface supports IP multicast, **multicast** is the default. Otherwise, the addresses are included in broadcast router advertisement packets. If the router does not support IP multicast, the addresses are not included.

Default

multicast if the router supports IP multicast; **broadcast** if the router does not support IP multicast.

Required Privilege Level

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

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

broadcast | 62

priority (Protocols Router Discovery)

Syntax

```
priority number;
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the preference of the address to become a default router. This preference is set relative to the preferences of other router addresses on the same subnet.

Options

number—Preference of the addresses for becoming the default router. A higher value indicates that the address has a greater preference for becoming the default router.

Range: 0 through 0x80000000

Default: 0 (This address has the least chance of becoming the default router.)

Required Privilege Level

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

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

router-discovery

Syntax

```
router-discovery {
    disable;
    address address {
        (advertise | ignore);
        (broadcast | multicast);
        (ineligible | priority number);
    }
    interface interface-name {
        lifetime seconds;
        max-advertisement-interval seconds;
        min-advertisement-interval seconds;
    }
    traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier> <disable>;
    }
}
```

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Enable ICMP router discovery (server mode) on the router.

The remaining statements are explained separately. See CLI Explorer.

Default

Router discovery is disabled on the router. When router discovery is enabled, the default behavior is to advertise all interfaces. If the router supports multicast, all the IPv4 Layer 3 interfaces are advertised through multicast. Otherwise, all the IPv4 Layer 3 interfaces are advertised through broadcast.

Required Privilege Level

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

RELATED DOCUMENTATION

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

traceoptions (Protocols Router Discovery)

Syntax

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

Hierarchy Level

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

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure ICMP protocol-level tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default ICMP protocol-level tracing options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

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 within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place ICMP tracing output in the file **icmp-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 also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 2 files

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

These are the global tracing options:

- **all**—All tracing operations
- **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**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

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

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

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). 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 also must 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: 1 MB

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

Understanding the ICMP Protocol for Discovering Gateways to Other Networks | 18

Example: Configuring the ICMP Protocol for Discovering Gateways to Other Networks | 19

CHAPTER 7

Operational Commands

monitor interface | **80**

monitor start | **96**

monitor stop | **98**

ping | **100**

probe | **108**

show log | **113**

traceroute | **121**

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 4 on page 81. The keys are not case-sensitive.

Table 4: 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 5 on page 81. The keys are not case-sensitive.

Table 5: 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 83
monitor interface (OTN Interface) on page 85
monitor interface (MX480 Router with MPC5E and 10-Gigabit Ethernet OTN Interface) on page 86
monitor interface (MX480 Router with MPC5E and 100-Gigabit Ethernet Interface) on page 87
monitor interface (MX2010 Router with MPC6E and 10-Gigabit Ethernet OTN Interface) on page 88
monitor interface (MX2010 Router with MPC6E and 100-Gigabit Ethernet OTN Interface) on page 90
monitor interface (MX2020 Router with MPC6E and 10-Gigabit Ethernet OTN Interface) on page 90
monitor interface (Logical) on page 91
monitor interface (QFX3500 Switch) on page 92
monitor interface traffic on page 93
monitor interface traffic (QFX3500 Switch) on page 93
monitor interface traffic detail (QFX3500 Switch) on page 94

Output Fields

Table 6 on page 82 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 6: 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 6: 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 | 98

List of Sample Output

monitor start on page 97

Output Fields

Table 7 on page 97 describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

Table 7: 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 | 96

List of Sample Output

monitor stop on page 99

Output Fields

This command produces no output.

Sample Output

monitor stop

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

ping

List of Syntax

Syntax on page 100
Syntax (QFX Series) on page 100
Syntax (Junos OS Evolved) on page 101

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 105
ping ce-ip <destination-ip-address> instance <routing-instance-name> source-ip <source-ip-address> (VPLS) on page 105
ping hostname on page 105
ping hostname rapid on page 106
ping hostname size count on page 106

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
```

probe

Syntax

```
probe <proxy-ip-address-or-hostname>
<by-address>
<by-index>
<by-name>
<by-remote-address>
<count>
<detail>
<inet>
<inet6>
<interface>
<interval>
<logical-system>
<routing-instance>
<source>
<tos>
<ttl>
<wait>
```

Release Information

Command introduced in Junos OS Release 20.3R1.

Description

Query interface details on a remote node.

Options

by-address—(Optional) Probe a local interface by address.

by-index—(Optional) Probe a local interface by index.

by-name—(Optional) Probe a local interface by name.

by-remote-address—(Optional) Probe a remote interface by remote address.

count—(Optional) Specify the number of probe requests or packets sent out.

Range: 1 through 255

detail—(Optional) Display details of the incoming interface of the received packet.

inet—(Optional) Force probe to IPv4 network family destination.

inet6—(Optional) Force probe to IPv6 network family destination.

interface *interface-name*—(Optional) Specify the source interface to send the probe requests.

interval *seconds*—(Optional) Specify the time interval between the probe requests.

Default: 1 second

Range: 1 through infinity

logical-system *logical-system-name*—(Optional) Specify the name of the logical system used to send the probe requests.

routing-instance *routing-instance-name*—(Optional) Specify the routing instance for the probe attempt.

source *source-address*—(Optional) Specify the source IP address of the probing interface.

tos *type-of-service*—(Optional) Specify the type-of-service (ToS) field in the IP header of the probe packets.

Range: 0 through 255

ttl *value*—(Optional) Specify the time-to-live (TTL) value for IPv4 packets and hop-limit value for IPv6 packets for the number of hops/devices to cross.

Range: 0 through 255

wait *seconds*—(Optional) Specify the maximum wait time in seconds after the final packet is sent. If this option is not specified, the default delay is 10 seconds.

Default: 10 seconds

Required Privilege Level

system-To view this statement in the configuration.

RELATED DOCUMENTATION

What is the Probe Command? | 42

extended-echo | 65

List of Sample Output

Probing Using Interface Names of the Probed Interface through Proxy-IP Address on page 110
Probing Using Interface-index of the Probed Interface Resides in the Proxy Node on page 111
Probing Using IPv4 or IPv6 Address in the Proxy Node on page 111
Probing Proxy Node for Information about the Remote Device on page 111
Probing the Proxy IP Address for Information About the Probed Interface When Both the Addresses are Configured on the Same Node on page 112
Probing the Proxy IP Address for Information About the Probed Interface Details Using the Proxy Interface Name on page 112

Output Fields

Table 1 lists the output fields for the **probe** command. The output fields are listed in the appropriate order in which they appear.

Table 8: probe Output Fields

Field Name	Field Description	Level of Output
icmp-seq	ICMP sequence number of the proxy IP address.	All levels
ttl	Time to live of the packets sent—0 to 255.	none
code	ICMP code to specify the state of the unreachable packet—0 to 15.	All levels
state	Reachable or unreachable state of the IPv6 or IPv4 address.	All levels
active	Inactive/active state (0 indicates inactive state and 1 indicates active state) of the IPv4 or IPv6 address configured.	All levels
IPv4	Unconfigured/configured state (0 indicates unconfigured and 1 indicates configured state) of IPv4 address.	All levels
IPv6	Unconfigured/configured state (0 indicates unconfigured and 1 indicates configured state) of IPv6 address.	All levels
time	Time taken by the packet to reach the destination in milliseconds (ms).	All levels
hlim	Hop limit to the destination.	detail

Sample Output

Probing Using Interface Names of the Probed Interface through Proxy-IP Address

`user@host>` **probe 198.52.100.1 count 1 by-name ge-0/0/0.0**

```
PROBE 198.52.100.1 (198.52.100.1): 56 data bytes
120 bytes from 198.52.100.1: icmp_seq=0 ttl=64 code=0 state=0 active=1 IPv4=1
IPv6=1 time=1.019 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.019/1.019/1.019/0.000 ms
```

Probing Using Interface-index of the Probed Interface Resides in the Proxy Node

`user@host>` **probe 198.52.100.1 count 2 by-index 322**

```
PROBE 198.52.100.1 (198.52.100.1):

28 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.871 ms
28 bytes from 198.52.100.1: icmp_seq=1 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.864 ms


--- 198.52.100.1 probe statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.864/0.867/0.871/0.004 ms
```

Probing Using IPv4 or IPv6 Address in the Proxy Node

`user@host>` **probe 198.52.102.1 count 2 by-address 2001::5**

```
PROBE 198.52.100.1 (198.52.100.1):
28 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.861 ms
28 bytes from 198.52.100.1: icmp_seq=1 ttl=255 code=0 state=0 active=1 IPv4=1
IPv6=0 time=0.854 ms


--- 198.52.100.1 probe statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.861/0.854/0.847/0.006 ms
```

Probing Proxy Node for Information about the Remote Device

`user@host>` **probe 198.52.100.1 count 1 by-remote-address 198.52.102.36**

```
PROBE 198.52.100.1 (198.52.100.1):
44 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=0 state=2 active=0 IPv4=0
IPv6=0 time=5.691 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 5.691/5.691/5.691/0.000 ms
```

Probing the Proxy IP Address for Information About the Probed Interface When Both the Addresses are Configured on the Same Node

```
user@host> probe 198.52.100.1 count 1 by-address 198.21.1.14
```

```
PROBE 198.52.100.1 (198.52.100.1):
35 bytes from 198.52.100.1: icmp_seq=0 ttl=255 code=2 state=0 active=0 IPv4=0
IPv6=0 time=5.185 ms

--- 198.52.100.1 probe statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
round-trip min/avg/max/stddev = 5.185/5.185/5.185/0.000 ms
```

Probing the Proxy IP Address for Information About the Probed Interface Details Using the Proxy Interface Name

```
user@host>probe 2001:15::2 by-name ge-0/0/2.0 detail
```

```
PROBE6(112=40+8+64 bytes) 2001:15::1 --> 2001:15::2
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=0 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=8.372 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=1 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=5.646 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=2 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=240.701 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=3 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=5.100 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=4 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=39.146 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=5 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=3.231 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=6 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=4.469 ms
72 bytes from 2001:15::2 via ge-0/0/0.0: icmp_seq=7 hlim=64 code=0 state=0 active=1
 IPv4=1 IPv6=1 time=3.360 ms
^C
--- 2001:15::2 probe6 statistics ---
8 packets transmitted, 8 packets received, 0% packet loss
round-trip min/avg/max/std-dev = 3.231/38.753/240.701/77.156 ms
```

show log

List of Syntax

Syntax on page 113
Syntax (QFX Series and OCX Series) on page 113
Syntax (TX Matrix Router) on page 113

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 115
show log filename on page 116
show log filename (QFabric System) on page 118
show log user on page 119
show log accepted-traffic (SRX4600, SRX5400, SRX5600, and SRX5800) on page 119

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 121
Syntax (QFX Series and OCX Series) on page 121

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 125
traceroute as-number-lookup host on page 125
traceroute no-resolve on page 125
traceroute propogate-ttl on page 126
traceroute (Between CE Routers, Layer 3 VPN) on page 126
traceroute (Through an MPLS LSP) on page 126
traceroute routing-instance no-resolve (Through an MPLS LSP) on page 126
traceroute (Junos OS Evolved, Through an MPLS LSP) on page 127

Output Fields

Table 9 on page 124 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 9: 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
```