JUNIPER NETWORKS

Junos® OS

IPv6 Neighbor Discovery Configuration Guide

Release
12.1

Published: 2012-08-10

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Junos® OS IPv6 Neighbor Discovery Configuration Guide
12.1

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 http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

List of Figures

List of Tables

About the Documentation

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 http://www.juniper.net/techpubs/.

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 http://www.juniper.net/books .

Supported Platforms

For the features described in this document, the following platforms are supported:

- T Series
- MX Series
- M Series

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 the *Junos OS CLI User Guide*.

Documentation Conventions

Table 1 on page xi 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.

Table 2 on page xi 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`

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

Convention	Description	Examples
Italic text like this	• Introduces or emphasizes important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
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)	Enclose 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)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify 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.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web 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 J-Web selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at https://www.juniper.net/cgi-bin/docbugreport/ . If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need 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 http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf .
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/ .
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://www.juniper.net/alerts/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/ .
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

PART 1

Overview

CHAPTER 1

Introduction to Neighbor Discovery

Neighbor Discovery Overview

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

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

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

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

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

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

This section discusses the following topics:

Router Discovery

Router advertisements can contain a list of prefixes. These prefixes are used for address autoconfiguration, to maintain a database of onlink (on the same data link) prefixes, and

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

Address Resolution

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

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

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

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

Redirect

Redirect messages are sent to inform a host of a better next-hop router to a particular destination or an onlink neighbor. This is similar to ICMPv4 redirect.

CHAPTER 2

IPv6 Neighbor Discovery Standards

Supported ICMP and Neighbor Discovery Standards

The Junos OS substantially supports the following RFCs, which define standards for 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 4862, *IPv6 Stateless Address Autoconfiguration*
- RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*

Related Documentation

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

PART 2

Configuration

CHAPTER 3

Concept and Example

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

Understanding IPv6 Neighbor Discovery

IPv6 Neighbor Discovery has many improvements when compared to the corresponding IPv4 protocols.

For instance, Neighbor Discovery moves address resolution to the ICMP layer, which makes it much less media dependent than ARP, as well as adding the ability to use IP layer security when needed.

Additionally, Neighbor Discovery uses link-local addresses. This allows all nodes to maintain their router associations even when the site is renumbered to a new global prefix.

Another improvement worth noting is that Neighbor Discovery messages carry link-layer address information, so a single message (or pair of messages) is all that is needed for nodes to resolve the others' addresses. No additional address resolution is needed.

Neighbor unreachability detection is built in, making packet delivery much more robust in a changing network. Using neighbor unreachability detection, Neighbor Discovery detects router failures, link failures, and partial link failures such as one-way communication.

And finally, IPv6 router advertisements carry prefixes (including network masks) and support multiple prefixes on the same link. Hosts can learn on-link prefixes from router advertisements or, when the router is configured to withhold them, from redirects as needed.

SLAAC

In addition to all the other improvements it brings to the networking world, Neighbor Discovery also enables address autoconfiguration, namely Stateless Address Autoconfiguration (SLAAC). IPv6 maintains the capability for stateful address assignment

through DHCPv6 (and static assignment), but SLAAC provides a lightweight address configuration method that might be desirable in many circumstances.

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

- Phase 1—Steps for local connectivity:

  1. Link-Local Address Generation: Any time that a multicast-capable IPv6-enabled interface is turned up, the node generates a link-local address for that interface. This is done by appending an interface identifier to the link-local prefix (FE80::/10).
  2. Duplicate Detection: Before assigning the new link-local address to its interface, the node verifies that the address is unique. This is accomplished by sending a Neighbor Solicitation message destined to the new address. If there is a reply, then the address is a duplicate and the process stops, requiring operator intervention.
  3. Link-Local Address Assignment: If the address is unique, the node assigns it to the interface for which it was generated.

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

- Phase 2—Steps for global connectivity:

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

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

Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery

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

Requirements

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

Overview

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

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

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

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

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

```
            flag flag;
        }
    }
}
```

Figure 1 on page 12 shows a simplified sample topology.

Figure 1: ICMP Router Discovery Topology

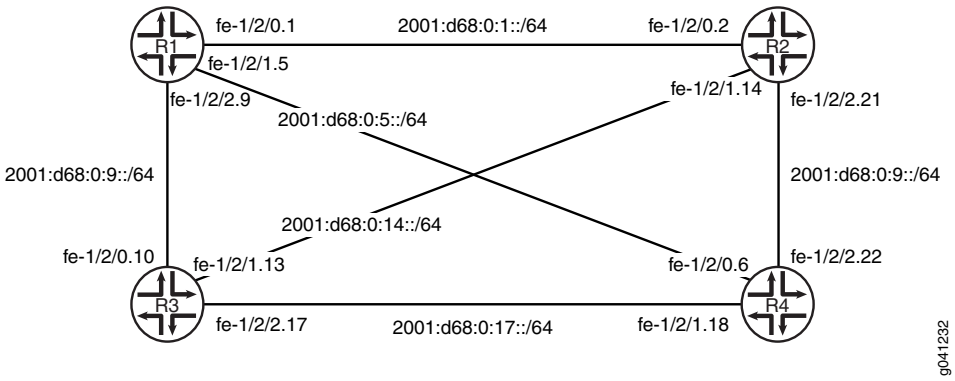

This example shows how to make sure that all of the IPv6 hosts attached to the subnets in the sample topology can auto-configure a local EUI-64 address.

"CLI Quick Configuration" on page 12 shows the configuration for all of the devices in Figure 1 on page 12. "Step-by-Step Procedure" on page 13 describes the steps on Device R1.

Configuration

CLI Quick Configuration

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

Device R1

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

Device R2

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

```
set protocols router-advertisement interface fe-1/2/2.21 prefix 2001:db8:0:21::/64
```

Device R3

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

Device R4

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

Step-by-Step Procedure

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

To configure a IPv6 neighbor discovery:

1. Configure the network interfaces.

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

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

   user@R1# set fe-1/2/1 unit 5 description to-P4
   user@R1# set fe-1/2/1 unit 5 family inet6 address 2001:db8:0:5::/64 eui-64

   user@R1# set fe-1/2/2 unit 9 description to-P3
   user@R1# set fe-1/2/2 unit 9 family inet6 address 2001:db8:0:9::/64 eui-64

   user@R1# set lo0 unit 1 family inet6 address 2001:db8::1/128
   ```

2. Enable neighbor discovery.

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

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

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-P2;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
fe-1/2/1 {
  unit 5 {
    description to-P4;
    family inet6 {
      address 2001:db8:0:5::/64 {
        eui-64;
      }
    }
  }
}
fe-1/2/2 {
  unit 9 {
    description to-P3;
    family inet6 {
      address 2001:db8:0:9::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 1 {
    family inet6 {
      address 2001:db8::1/128;
    }
  }
}
```

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


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

Verification

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

Checking the Interfaces

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

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

```
user@R1> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
fe-1/2/0
fe-1/2/0.1              up    up   inet6    2001:db8:0:1:2a0:a514:0:14c/64
                                            fe80::2a0:a514:0:14c/64
fe-1/2/1.5              up    up   inet6    2001:db8:0:5:2a0:a514:0:54c/64
                                            fe80::2a0:a514:0:54c/64
fe-1/2/2.9              up    up   inet6    2001:db8:0:9:2a0:a514:0:94c/64
                                            fe80::2a0:a514:0:94c/64
lo0
lo0.1                   up    up   inet6    2001:db8::1
                                            fe80::2a0:a50f:fc56:14c
```

Meaning The output shows that all interfaces are configured with the IPv6 (inet6) address family. Each IPv6-enabled interface has two IPv6 addresses; one link-local address, and one global address. The global addresses match those shown in Figure 1 on page 12. Junos OS automatically creates a link-local address for any interface that is enabled for IPv6 operation. All link-local addresses begin with the fe80::/64 prefix. The host portion of the address is a full 64 bits long and matches the link-local interface identifier. When an interface address is configured using the **eui-64** statement, its interface identifier matches the interface identifier of the link-local address. This is because link-local addresses are coded according to the EUI-64 specification.

Pinging the Interfaces

Purpose Verify connectivity between the directly connected interfaces.

Action

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

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

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

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

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

```
user@R1> ping 2001:db8:0:1:2a0:a514:0:24c
PING6(56=40+8+8 bytes) 2001:db8:0:1:2a0:a514:0:14c -->
2001:db8:0:1:2a0:a514:0:24c
16 bytes from 2001:db8:0:1:2a0:a514:0:24c, icmp_seq=0 hlim=64 time=20.412 ms
16 bytes from 2001:db8:0:1:2a0:a514:0:24c, icmp_seq=1 hlim=64 time=18.897 ms
16 bytes from 2001:db8:0:1:2a0:a514:0:24c, icmp_seq=2 hlim=64 time=1.389 ms
```

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

Checking the IPv6 Neighbor Cache

Purpose Display information about the IPv6 neighbors.

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

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

```
user@R1> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale       546 yes no
fe-1/2/0.1
fe80::2a0:a514:0:24c         00:05:85:8f:c8:bd  stale       258 yes no
fe-1/2/0.1
fe80::2a0:a514:0:64c         00:05:85:8f:c8:bd  stale       111 yes no
fe-1/2/1.5
fe80::2a0:a514:0:a4c         00:05:85:8f:c8:bd  stale       327 yes no
fe-1/2/2.9
```

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

Verifying IPv6 Router Advertisements

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

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

```
user@R1> show ipv6 router-advertisement
Interface: fe-1/2/0.1
  Advertisements sent: 37, last sent 00:01:41 ago
  Solicits received: 0
  Advertisements received: 38
  Advertisement from fe80::2a0:a514:0:24c, heard 00:05:46 ago
    Managed: 0
```

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

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

Tracing Neighbor Discovery Events

Purpose Perform additional validation by tracing router advertisements.

Action 1. Configure trace operations.

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

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


```
user@R1> show log ipv6-nd-trace
Mar 29 14:07:16 trace_on: Tracing to "/var/log/P1/ipv6-nd-trace" started
Mar 29 14:07:16.287229 background dispatch running job
ipv6_ra_delete_interface_config_job for task Router-Advertisement
Mar 29 14:07:16.287452 task_job_delete: delete background job
ipv6_ra_delete_interface_config_job for task Router-Advertisement
Mar 29 14:07:16.287505 background dispatch completed job
ipv6_ra_delete_interface_config_job for task Router-Advertisement
Mar 29 14:07:16.288288 ipv6_ra_iflchange(Router-Advertisement): ifl 0xb904378
 ifl fe-1/2/2.9 104 change 0, intf 0xba140d8
Mar 29 14:07:16.288450 ipv6_ra_iflchange(Router-Advertisement): ifl 0xb904250
 ifl fe-1/2/0.1 85 change 0, intf 0xba14000
Mar 29 14:07:16.288656 ipv6_ra_iflchange(Router-Advertisement): ifl 0xb9044a0
 ifl fe-1/2/1.5 80 change 0, intf 0xba1406c
Mar 29 14:07:16.289293 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba002bc
 fe80::2a0:a514:0:54c ifl fe-1/2/1.5 80 change 0, intf 0xba1406c
Mar 29 14:07:16.289358 -- nochange/add
Mar 29 14:07:16.289624 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba00230
 2001:db8:0:5:2a0:a514:0:54c ifl fe-1/2/1.5 80 change 0, intf 0xba1406c
Mar 29 14:07:16.289682 -- nochange/add
Mar 29 14:07:16.289950 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba001a4
 fe80::2a0:a514:0:14c ifl fe-1/2/0.1 85 change 0, intf 0xba14000
Mar 29 14:07:16.290009 -- nochange/add
Mar 29 14:07:16.290302 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba00118
 2001:db8:0:1:2a0:a514:0:14c ifl fe-1/2/0.1 85 change 0, intf 0xba14000
Mar 29 14:07:16.290365 -- nochange/add
Mar 29 14:07:16.290634 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba003d4
 fe80::2a0:a514:0:94c ifl fe-1/2/2.9 104 change 0, intf 0xba140d8
Mar 29 14:07:16.290694 -- nochange/add
Mar 29 14:07:16.290958 ipv6_ra_ifachange(Router-Advertisement): ifa 0xba00348
 2001:db8:0:9:2a0:a514:0:94c ifl fe-1/2/2.9 104 change 0, intf 0xba140d8
Mar 29 14:07:16.291017 -- nochange/add
Mar 29 14:07:20.808516 task_job_create_foreground: create job ipv6 ra for task
 Router-Advertisement
Mar 29 14:07:20.808921 foreground dispatch running job ipv6 ra for task
Router-Advertisement
Mar 29 14:07:20.809027 ipv6_ra_send_advertisement: sending advertisement for
ifl 104 to ff02::1
Mar 29 14:07:20.809087 (4810916) sending advertisement for ifl 104
Mar 29 14:07:20.809170  ifa 0xba00348 2001:db8:0:9:2a0:a514:0:94c/64
Mar 29 14:07:20.809539  --> sent 56 bytes
Mar 29 14:07:20.809660 task_timer_reset: reset Router-Advertisement_ipv6ra
Mar 29 14:07:20.809725 task_timer_set_oneshot_latest: timer
Router-Advertisement_ipv6ra interval set to 7:07
Mar 29 14:07:20.809772 foreground dispatch completed job ipv6 ra for task
Router-Advertisement
```

Related Documentation

- Example: Configuring ICMP Router Discovery

CHAPTER 4

Configuration Statements

[edit protocols router-advertisement] Hierarchy Level

The following statement hierarchy can also be included at the **[edit logical-systems *logical-system-name*]** hierarchy level.

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

Related Documentation

- Notational Conventions Used in Junos OS Configuration Hierarchies
- [edit protocols] Hierarchy Level

autonomous

Syntax (autonomous | no-autonomous);

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

Default autonomous

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

Related Documentation
- OBSOLETE - Configuring the Prefix Information Included in Neighbor Discovery Advertisements

current-hop-limit

Syntax current-hop-limit *number*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Default value placed in the hop count field of the IP header for outgoing packets.

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

Range: 0 through 255

Default: 64

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

Related Documentation
- OBSOLETE - Configuring the Hop Count in Outgoing Neighbor Discovery Packets

default-lifetime

Syntax default-lifetime *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Lifetime associated with a default router.

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

Range: Maximum advertisement interval value through 9000 seconds

Default: Three times the maximum advertisement interval value

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

Related Documentation

- max-advertisement-interval on page 24
- OBSOLETE - Configuring the Lifetime for the Default Neighbor Discovery Router

interface

Syntax

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

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

The remaining statements are explained separately.

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

Related Documentation

- OBSOLETE - Configuring an Interface to Send Neighbor Discovery Advertisements

link-mtu

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

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

Release Information Statement introduced in Junos OS 10.3.

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

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

Default Router advertisement messages do not include the MTU option.

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

Related Documentation
- OBSOLETE - Configuring the MTU Option for Neighbor Discovery Advertisements

managed-configuration

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

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

Default The configured object is disabled unless explicitly enabled.

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

Related Documentation
- OBSOLETE - Enabling Stateful Autoconfiguration with Neighbor Discovery

max-advertisement-interval

Syntax max-advertisement-interval *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Maximum interval between each router advertisement message.

Options *seconds*—Maximum interval.

Range: 4 through 1800 seconds

Default: 600 seconds

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

Related Documentation

- min-advertisement-interval on page 24
- OBSOLETE - Configuring the Frequency of Neighbor Discovery Advertisements

min-advertisement-interval

Syntax min-advertisement-interval *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Minimum interval between each router advertisement message.

Options *seconds*—Minimum interval.

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

Default: One-third the maximum advertisement interval value

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

Related Documentation

- max-advertisement-interval on page 24
- OBSOLETE - Configuring the Frequency of Neighbor Discovery Advertisements

on-link

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

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

Default The configured object is enabled 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
- OBSOLETE - Configuring the Prefix Information Included in Neighbor Discovery Advertisements

other-stateful-configuration

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

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

Default The configured object is disabled unless explicitly enabled.

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

Related Documentation
- OBSOLETE - Enabling Stateful Autoconfiguration with Neighbor Discovery

preferred-lifetime

Syntax preferred-lifetime *seconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

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

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

Default: 604,800 seconds

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

Related Documentation
- valid-lifetime on page 30
- OBSOLETE - Configuring the Prefix Information Included in Neighbor Discovery Advertisements

prefix

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

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure prefix properties in router advertisement messages.

Options *prefix*—Prefix name.

The remaining statements are explained separately.

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

Related Documentation
- OBSOLETE - Configuring the Prefix Information Included in Neighbor Discovery Advertisements

reachable-time

Syntax reachable-time *milliseconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

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

Options ***milliseconds***—Reachability time limit.

Range: 0 through 3,600,000 milliseconds

Default: 0 milliseconds

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

Related Documentation

- OBSOLETE - Configuring the Delay Before Neighbor-Discovery Neighbors Mark the Router as Down

retransmit-timer

Syntax retransmit-timer *milliseconds*;

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Set the retransmission frequency of neighbor solicitation messages.

Options ***milliseconds***—Retransmission frequency.

Default: 0 milliseconds

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

Related Documentation

- OBSOLETE - Configuring the Frequency of Neighbor Solicitation Messages

router-advertisement

Syntax router-advertisement {...}

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit protocols]

Release Information Statement introduced before Junos OS Release 7.4.

Description Enable router advertisement.

The remaining statements are explained separately.

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

Related Documentation

- OBSOLETE - Configuring an Interface to Send Neighbor Discovery Advertisements

traceoptions

Syntax

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

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

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify router advertisement protocol-level tracing options.

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

Options

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

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

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

Range: 2 through 1000 files

Default: 10 files

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

- **all**—All tracing operations
- **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

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

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

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

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

Default: 128 KB

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

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

Related Documentation
- OBSOLETE - Tracing Neighbor Discovery Protocol Traffic

valid-lifetime

Syntax
valid-lifetime *seconds*;

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

Release Information
Statement introduced before Junos OS Release 7.4.

Description
Specify how long the prefix remains valid for onlink determination.

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

Default: 2,592,000 seconds

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

Related Documentation
- preferred-lifetime on page 26
- OBSOLETE - Configuring the Prefix Information Included in Neighbor Discovery Advertisements

PART 3

Administration

CHAPTER 5

Operational Commands

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.

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

Options **none**—Display real-time statistics for all interfaces.

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.

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

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 using 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 *Junos OS Interfaces Command Reference*. To control the output of the **monitor interface** ***interface-name*** command while it is running, use the keys listed in Table 3 on page 34. The keys are not case-sensitive.

Table 3: 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.

Table 3: Output Control Keys for the monitor interface interface-name Command *(continued)*

Key	Action
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 4 on page 35. The keys are not case-sensitive.

Table 4: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bytes 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 37
monitor interface (OTN Interface) on page 39
monitor interface (Logical) on page 40
monitor interface traffic on page 40
monitor interface (QFX3500 Switch) on page 41
monitor interface traffic on page 41
monitor interface traffic (QFX3500 Switch) on page 41
monitor interface traffic detail (QFX3500 Switch) on page 42

Output Fields Table 5 on page 36 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 5: 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
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 takes awhile (generally, 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 takes awhile (generally, 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

Table 5: monitor interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
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 takes awhile (generally, 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 (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 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 (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 on page 46

List of Sample Output monitor start on page 44

Output Fields Table 6 on page 44 describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

Table 6: 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 trip.jcmax.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 on page 44

List of Sample Output monitor stop on page 46

Output Fields This command produces no output.

Sample Output

monitor stop

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


ping

Syntax

ping *host*
<bypass-routing>
<count *requests*>
<detail>
<do-not-fragment>
<inet | inet6>
<interface *source-interface*>
<interval *seconds*>
<logical-system (all | *logical-system-name*)>
<loose-source *value*>
<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 (QFX Series)

ping *host*
<bypass-routing>
<count *requests*>
<detail>
<do-not-fragment>
<inet>
<interface *source-interface*>
<interval *seconds*>
<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*>

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.

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. Type 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.

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) 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 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**.

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.

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**.

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

- Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages

List of Sample Output

ping hostname on page 50
ping hostname size count on page 50
ping hostname rapid on page 50

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.

Sample Output

ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

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

ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!!
--- skye.net 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
```

show ipv6 neighbors

Syntax show ipv6 neighbors

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.

Description Display information about the IPv6 neighbor cache.

Options This command has no options.

Required Privilege Level view

Related Documentation
- clear ipv6 neighbors

List of Sample Output
show ipv6 neighbors on page 51
show ipv6 neighbors on page 51

Output Fields Table 7 on page 51 describes the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 7: show ipv6 neighbors Output Fields

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

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp  Rtr  Interface
fe80::2a0:c9ff:fe5b:4c1e     00:a0:c9:5b:4c:1e  reachable   15   yes  fxp0.0
```

show ipv6 neighbors

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

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
fe80::14fb:5dcf:54bd:ff76    00:90:69:a0:a8:bc  stale       1113 yes yes
ge-3/2/0.0
```

show ipv6 router-advertisement

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

Release Information
Command introduced before Junos OS Release 7.4.

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

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

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

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

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

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

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

Required Privilege Level
view

Related Documentation
- clear ipv6 router-advertisement

List of Sample Output
show ipv6 router-advertisement on page 54
show ipv6 router-advertisement conflicts on page 55
show ipv6 router-advertisement prefix on page 55

Output Fields
Table 8 on page 53 describes the output fields for the **show ipv6 router-advertisement** command. Output fields are listed in the approximate order in which they appear.

Table 8: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.
Advertisements sent	Number of router advertisements sent and elapsed time since they were sent.

Table 8: show ipv6 router-advertisement Output Fields *(continued)*

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

Sample Output

show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
```

```
    Retransmit timer: 0 ms
    Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 180 sec [1800 sec]
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 8040:1::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
```

show log

Syntax
show log
<*filename* | user <*username*>>

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.

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

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).

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

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

List of Sample Output
show log on page 56
show log filename on page 57
show log user on page 57

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
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
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
...
```

show log user

```
user@host> show log user
darius   mg2546                   Thu Oct  1 19:37   still logged in
darius   mg2529                   Thu Oct  1 19:08 - 19:36  (00:28)
darius   mg2518                   Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                   Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    jun.site.per    Wed Sep 30 18:39 - 18:41  (00:02)
alex     ttyp1    192.168.1.2      Wed Sep 30 01:03 - 01:22  (00:19)
```

traceroute

Syntax

traceroute *host*
<as-number-lookup>
<bypass-routing>
<clns>
<gateway *address*>
<inet | inet6>
<interface *interface-name*>
<logical system (all | *logical-system-name*)>
<mpls (ldp *FEC address* | rsvp *label-switched-path-name*)>
<no-resolve>
<propagate-ttl>
<routing-instance *routing-instance-name*>
<source *source-address*>
<tos *value*>
<ttl *value*>
<wait *seconds*>

Syntax (QFX Series)

traceroute *host*
<as-number-lookup>
<bypass-routing>
<gateway *address*>
<inet>
<interface *interface-name*>
<monitor *host*>
<no-resolve>
<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 11.1 for the QFX Series.

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.

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

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) Perform this operation on all logical systems or on a particular logical system.

monitor ***host***—(Optional) Display real-time monitoring information for the specified host.

monitor*host*—(Optional) Perform this operation to display real-time monitoring information.

monitor*host*—(Optional) Perform this operation to display real-time monitoring information.

mpls (ldp ***FEC address*** **| rsvp** ***label-switched-path name*****)**—(Optional) See traceroute mpls ldp and traceroute mpls rsvp.

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

propagate-ttl—(Optional) On the PE router, 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.

Required Privilege Level network

Related Documentation

- traceroute monitor

List of Sample Output

Output Fields

Table 9 on page 60 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 green.company.net (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
```

traceroute no-resolve

```
user@host> traceroute santacruz no-resolve
traceroute to green.company.net (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 vpn09.skybank.net (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  vpn09.skybank.net (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
```

PART 4

Troubleshooting

CHAPTER 6

Routing Protocol Process Memory FAQ

Routing Protocol Process Memory FAQ Overview

The Juniper Networks Junos operating system (Junos OS) is based on the FreeBSD Unix operating system. The open source software is modified and hardened to operate in the device's specialized environment. For example, some executables have been deleted while other utilities have been de-emphasized. Additionally, certain software processes have been added to enhance the routing functionality. The result of this transformation is the kernel, the heart of the Junos OS software.

The kernel is responsible for generating multiple processes that perform the actual functions of the device. Each process operates in its own protected memory space, providing isolation between the processes and resiliency in the event of a process failure. This is important in a core routing platform because a single process failure does not cause the entire device to cease functioning.

Some of the common software processes include the routing protocol process (rpd) that controls the device's protocols, the device control process (dcd) that controls the device's interfaces, the management process (mgd) that controls user access to the device, the chassis process (chassisd) that controls the device's properties itself, and the Packet Forwarding Engine process (pfed) that controls the communication between the device's Packet Forwarding Engine and the Routing Engine. Besides the above processes, there are other specialized processes that support additional functionality, such as the Simple Network Management Protocol (SNMP), Virtual Router Redundancy Protocol (VRRP), and Class of Service (CoS).

The routing protocol process is a software process within the Routing Engine software that controls the routing protocols that run on the device. Its functionality includes all protocol messages, routing table updates, and implementation of routing policies.

The routing protocol process starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements the routing policy, which allows you to control the routing information that is transferred between the routing

protocols and the routing table. Using the routing policy, you can filter and limit the transfer of information as well as set properties associated with specific routes.

Related Documentation

- Routing Protocol Process Memory FAQs on page 66

Routing Protocol Process Memory FAQs

The following sections present the most frequently asked questions and answers related to the routing protocol process memory utilization, operation, interpretation of related command outputs, and troubleshooting the software process.

Routing Protocol Process Memory Utilization FAQs

This section presents frequently asked questions and answers related to the memory usage of the routing protocol process.

Why does the routing protocol process use excessive memory?

The routing protocol process uses hundreds of megabytes of RAM in the Routing Engine to store information needed for the operation of routing and related protocols, such as BGP, OSPF, ISIS, RSVP, LDP, and MPLS. Such huge consumption of memory is common for the process, as the information it stores includes routes, next hops, interfaces, routing policies, labels, and label-switched paths (LSPs). Because access to the RAM memory is much faster than access to the hard disk, most of the routing protocol process information is stored in the RAM memory instead of using the hard disk space. This ensures that the performance of the routing protocol process is maximized.

How can I check the amount of memory the routing protocol process is using?

You can check the routing protocol process memory usage by entering the **show system processes** and the **show task memory** Junos OS command-line interface (CLI) operational mode commands.

The **show system processes** command displays information about software processes that are running on the device. You can check the routing protocol process memory usage by using the **show system processes** command with the **extensive** option.

The **show task memory** command displays a report generated by the routing protocol process on the memory utilization for routing protocol tasks on the Routing Engine. Although the report generated by the routing protocol process is on its own memory usage, it does not display all the memory used by the process. The value reported by the routing protocol process does not account for the memory used for the **TEXT** and **STACK** segments, or the memory used by the process's internal memory manager. The **show task memory** command also does not include the memory which has been deactivated by the routing protocol process, although some or all of that deactivated memory has not actually been freed by the kernel.

For more information about checking the routing protocol process memory usage, see Check Routing Protocol Process (rpd) Memory Usage in the *Junos OS Baseline Network Operations Guide*.

For more information about the show system processes command and the show task memory command, see the *Junos OS System Basics and Services Command Reference*.

I just deleted many routes from the routing protocol process. Why is the routing protocol process still using so much memory?

The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of process memory resident in the physical memory. This is also known as RSS or Resident Set Size. Any amount of memory deactivated by the process might still be considered part of the **RES** value. Generally, the kernel defers the actual freeing of deactivated memory until there is a memory shortage. This can lead to large discrepancies between the values reported by the routing protocol process and the kernel, even after the routing protocol process has deactivated a large amount of memory.

Interpreting Routing Protocol Process-Related Command Outputs FAQs

This section presents frequently asked questions and answers about the routing protocol process-related Junos OS CLI command outputs that are used to display the memory usage of the routing protocol process.

How do I interpret memory numbers displayed in the show system processes extensive command output?

The **show system processes extensive** command displays exhaustive system process information about software processes that are running on the device. This command is equivalent to the UNIX **top** command. However, the UNIX **top** command shows real-time memory usage, with the memory values constantly changing, while the **show system processes extensive** command provides a snapshot of memory usage in a given moment.

To check overall CPU and memory usage, enter the **show system processes extensive** command. Refer to Table 10 on page 69 for information about the **show system processes extensive** command output fields.

```
user@host> show system processes extensive
last pid:   544;  load averages:  0.00,  0.00,  0.00    18:30:33
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 184K Cache, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  544 root      30   0   604K   768K RUN      0:00  0.00%  0.00% top
    3 root      28   0     0K    12K psleep   0:00  0.00%  0.00% vmdaemon
    4 root      28   0     0K    12K update   0:03  0.00%  0.00% update
  528 aviva     18   0   660K   948K pause    0:00  0.00%  0.00% tcsh
  204 root      18   0   300K   544K pause    0:00  0.00%  0.00% csh
  131 root      18   0   332K   532K pause    0:00  0.00%  0.00% cron
  186 root      18   0   196K    68K pause    0:00  0.00%  0.00% watchdog
   27 root      10   0   512M 16288K mfsidl   0:00  0.00%  0.00% mount_mfs
    1 root      10   0   620K   344K wait     0:00  0.00%  0.00% init
  304 root       3   0   884K   900K ttyin    0:00  0.00%  0.00% bash
  200 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  203 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  202 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  201 root       3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  194 root       2   0  2248K  1640K select   0:11  0.00%  0.00% rpd
  205 root       2   0   964K   800K select   0:12  0.00%  0.00% tnp.chassisd
  189 root       2 -12   352K   740K select   0:03  0.00%  0.00% xntpd
  114 root       2   0   296K   612K select   0:00  0.00%  0.00% amd
```

```
  188 root      2    0   780K   600K select   0:00  0.00%  0.00% dcd
  527 root      2    0   176K   580K select   0:00  0.00%  0.00% rlogind
  195 root      2    0   212K   552K select   0:00  0.00%  0.00% inetd
  187 root      2    0   192K   532K select   0:00  0.00%  0.00% tnetd
   83 root      2    0   188K   520K select   0:00  0.00%  0.00% syslogd
  538 root      2    0  1324K   516K select   0:00  0.00%  0.00% mgd
   99 daemon    2    0   176K   492K select   0:00  0.00%  0.00% portmap
  163 root      2    0   572K   420K select   0:00  0.00%  0.00% nsrexecd
  192 root      2    0   560K   400K select   0:10  0.00%  0.00% snmpd
  191 root      2    0  1284K   376K select   0:00  0.00%  0.00% mgd
  537 aviva     2    0   636K   364K select   0:00  0.00%  0.00% cli
  193 root      2    0   312K   204K select   0:07  0.00%  0.00% mib2d
    5 root      2    0     0K    12K pfesel   0:00  0.00%  0.00% if_pfe
    2 root    -18    0     0K    12K psleep   0:00  0.00%  0.00% pagedaemon
    0 root    -18    0     0K     0K sched    0:00  0.00%  0.00% swapper
```

Table 10 on page 69 describes the output fields that represent the memory values for the **show system processes extensive** command. Output fields are listed in the approximate order in which they appear.

Table 10: show system processes extensive Output Fields

Field Name	Field Description
Mem	Information about physical and virtual memory allocation.
Active	Memory allocated and actively used by the process.
Inact	Memory allocated but not recently used, or memory deactivated by the processes. Inactive memory remains mapped in the address space of one or more processes and, therefore, counts toward the RSS value of those processes.
Wired	Memory that is not eligible to be swapped, usually used for in-kernel memory structure, memory physically locked by a process, or both.
Cache	Freed memory that is no longer associated with any process but still has valid contents that correspond to some file system blocks. Cache pages can be reclaimed as is when the corresponding file system blocks are accessed again. However, when the system is under memory pressure, the contents of Cache pages could be erased by the kernel and the pages reused to service any memory allocation requests.
Buf	Size of the virtual memory buffer used to hold data recently called from the disk.
Free	Free memory that is neither associated with any process nor contains any valid contents.
Swap	Information about swap memory. • Total—Total space on the swap device. • Used—Memory swapped to disk. • Free—Unused space available on the swap device.

The rest of the command output displays information about the memory usage of each process. The **SIZE** field indicates the size of the virtual address space, and the **RES** field indicates the amount of the process in physical memory, which is also known as RSS or Resident Set Size. For more information, see the show system processes command in the *Junos OS System Basics and Services Command Reference*.

What is the difference between Active and Inact memory that is displayed by the show system processes extensive command?

When the system is under memory pressure, the pageout process can free up memory from the **Inact** and, if necessary, **Active** pools after first preserving the contents of those pages on the swap device or backing file systems if necessary. When the pageout process runs, it scans memory to see which pages are good candidates to be unmapped and freed up. Thus, the distinction between **Active** and **Inact** memory is only used by the pageout process to determine which pool of pages to free first at the time of a memory shortage.

The pageout process first scans the **Inact** list and checks whether the pages on this list have been accessed since the time they have been listed here. The pages that have been accessed are moved from the **Inact** list to the **Active** list. On the other hand, pages that have not been accessed become prime candidates to be freed by the pageout process. If the pageout process cannot produce enough free pages from the **Inact** list, pages from the **Active** list are freed up.

Because the pageout process runs only when the system is under memory pressure, the pages on the **Inact** list remain untouched – even if they have not been accessed recently – when the amount of **Free** memory is adequate.

How do I interpret memory numbers displayed in the show task memory command output?

The **show task memory** command provides a comprehensive picture of the memory utilization for routing protocol tasks on the Routing Engine. The routing protocol process is the main task that uses Routing Engine memory.

To check routing process memory usage, enter the **show task memory** command.

```
user@host> show task memory
Memory                 Size (kB)  %Available  When
  Currently In Use:        29417          3%  now
  Maximum Ever Used:       33882          4%  00/02/11 22:07:03
  Available:              756281        100%  now
```

Table 11 on page 70 describes the output fields for the **show task memory** command. Output fields are listed in the approximate order in which they appear.

Table 11: show task memory Output Fields

Field Name	Field Description
Memory Currently In Use	Memory currently in use. Dynamically allocated memory plus the **DATA** segment memory in kilobytes.
Memory Maximum Ever Used	Maximum memory ever used.
Memory Available	Memory currently available.

The **show task memory** command does not display all the memory used by the routing protocol process. This value does not account for the memory used for the **TEXT** and

STACK segments, or the memory used by the routing protocol process's internal memory manager. The **show task memory** command also does not include the memory which has been deactivated by the routing protocol process, although some or all of that deactivated memory has not actually been freed by the kernel.

Why is the Memory Currently In Use value less than the RES value?

The **show task memory** command displays a **Memory Currently In Use** value measured in kilobytes. This value is the dynamically allocated memory plus the **DATA** segment memory. The **show system processes extensive** command displays a **RES** value measured in kilobytes. This value represents the amount of process memory resident in the physical memory. This is also known as RSS or Resident Set Size.

The **Memory Currently In Use** value does not account for all of the memory that the routing protocol process uses. This value does not include the memory used for the **TEXT** and the **STACK** segments, and a small percentage of memory used by the routing protocol process's internal memory manager. The **show task memory** command also does not include the memory which has been deactivated by the routing protocol process, although some or all of that deactivated memory has not actually been freed by the kernel.

Any amount of memory deactivated by the routing protocol process might still be considered part of the **RES** value. Generally, the kernel defers the actual freeing of deactivated memory until there is a memory shortage This can lead to large discrepancies between the **Memory Currently In Use** value and the **RES** value.

Routing Protocol Process Memory Swapping FAQs

This section presents frequently asked questions and answers related to the memory swapping of the routing protocol process from the Routing Engine memory to the hard disk memory.

Why does the system start swapping when I try to perform a core dump using the request system core-dumps command?

The **request system core-dumps** command displays a list of system core files created when the device has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, path, and filename. You can use the **core-filename** option and the **core-file-info**, **brief**, and **detail** options to display more information about the specified core dump files.

You can use the **request system core-dumps** command to perform a non-fatal core dump without aborting the routing protocol process. To do this, the routing protocol process is forked, generating a second copy, and then aborted. This process can double the memory consumed by the two copies of the routing protocol process, pushing the system into swap.

Why does the show system processes extensive command show that memory is swapped to disk even though there is plenty of free memory?

Memory can remain swapped out indefinitely if it is not accessed again. Therefore, the **show system processes extensive** command shows that memory is swapped to disk even though there is plenty of free memory. Such a situation is not unusual.

Troubleshooting the Routing Protocol Process FAQs

This section presents frequently asked questions and answers related to a shortage of memory and memory leakage by the routing protocol process.

What does the RPD_OS_MEMHIGH message mean?

The **RPD_OS_MEMHIGH** message is written into the system message file if the routing protocol process is running out of memory. This message alerts you that the routing protocol process is using the indicated amount and percentage of Routing Engine memory, which is considered excessive. This message is generated either because the routing protocol process is leaking memory or the use of system resources is excessive, perhaps because routing filters are not configured properly or the configured network topology is very complex.

When the memory utilization for the routing protocol process is using all available Routing Engine DRAM memory or reaches the maximum memory limit, a message of the following form is written every minute in the syslog message file:

```
RPD_OS_MEMHIGH: Using 188830 KB of memory, 100 percent of available
```

This message includes the amount (in kilobytes), the percentage, or both of the available memory in use.

This message should not appear under normal conditions, as any further memory allocations usually require a portion of existing memory to be written to swap. As a recommended solution, increase the amount of RAM in the Routing Engine. For more information, see http://kb.juniper.net/InfoCenter/index?page=content&id=KB14186 .

What can I do when there is a memory shortage even after a swap?

We do not recommend that the system operate in this state, notwithstanding the existence of swap. The protocols that run in the routing protocol process usually have a real-time requirement that cannot reliably withstand the latency of being swapped to hard disk. If the memory shortage has not resulted from a memory leak, then either a reduction in the memory usage or an upgrade to a higher memory-capacity Routing Engine is required.

What is the task_timer?

The source of a routing protocol process memory leak can usually be identified by dumping the timers for each task. You can use the **show task *task-name*** command to display routing protocol tasks on the Routing Engine. Tasks can be baseline tasks performed regardless of the device's configuration, and other tasks that depend on the device configuration.

For more information, see the show task command in the *Junos OS System Basics and Services Command Reference*.

Related Documentation

- Routing Protocol Process Memory FAQ Overview on page 65

PART 5

Index

Index

U

V