



---

Junos<sup>®</sup> OS

# Tunnel and Encryption Services Interfaces Feature Guide for Routing Devices

Release

15.1



---

Modified: 2016-07-26

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://www.juniper.net)

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

*Junos<sup>®</sup> OS Tunnel and Encryption Services Interfaces Feature Guide for Routing Devices*

15.1

Copyright © 2016, Juniper Networks, Inc.  
All rights reserved.

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

#### YEAR 2000 NOTICE

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

#### END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <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

	About the Documentation . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Supported Platforms . . . . .	xi
	Using the Examples in This Manual . . . . .	xi
	Merging a Full Example . . . . .	xii
	Merging a Snippet . . . . .	xii
	Documentation Conventions . . . . .	xiii
	Documentation Feedback . . . . .	xv
	Requesting Technical Support . . . . .	xv
	Self-Help Online Tools and Resources . . . . .	xv
	Opening a Case with JTAC . . . . .	xvi
<b>Part 1</b>	<b>Tunnel Services</b>	
<b>Chapter 1</b>	<b>Overview . . . . .</b>	<b>3</b>
	Tunnel Services Overview . . . . .	3
	Configuring Tunnel Interfaces on MX Series Routers . . . . .	6
	Configuring Tunnel Interfaces on T4000 Routers . . . . .	7
<b>Chapter 2</b>	<b>Encapsulating One Protocol over Another Using GRE Interfaces . . . . .</b>	<b>9</b>
	GRE Keepalive Time Overview . . . . .	9
	Configuring GRE Keepalive Time . . . . .	9
	Configuring Keepalive Time and Hold time for a GRE Tunnel Interface . . . . .	10
	Display GRE Keepalive Time Configuration . . . . .	10
	Display Keepalive Time Information on a GRE Tunnel Interface . . . . .	11
	Enabling Fragmentation on GRE Tunnels . . . . .	12
<b>Chapter 3</b>	<b>Encapsulating One IP Packet over Another Using IP-IP Interfaces . . . . .</b>	<b>15</b>
	Configuring IPv6-over-IPv4 Tunnels . . . . .	15
	Example: Configuring an IPv6-over-IPv4 Tunnel . . . . .	15
<b>Chapter 4</b>	<b>Filtering Unicast Packets Through Multicast Tunnel Interfaces . . . . .</b>	<b>17</b>
	Configuring Unicast Tunnels . . . . .	17
	Configuring a Key Number on GRE Tunnels . . . . .	19
	Enabling Fragmentation on GRE Tunnels . . . . .	20
	Specifying an MTU Setting for the Tunnel . . . . .	20
	Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header . . . . .	21
	Configuring Packet Reassembly . . . . .	21
	Examples: Configuring Unicast Tunnels . . . . .	22
	Restricting Tunnels to Multicast Traffic . . . . .	23

<b>Chapter 5</b>	<b>Connecting Logical Systems Using Logical Tunnel Interfaces . . . . .</b>	<b>25</b>
	Configuring Logical Tunnel Interfaces . . . . .	25
	Connecting Logical Systems . . . . .	25
	Example: Configuring Logical Tunnels . . . . .	26
	Redundant Logical Tunnels Overview . . . . .	28
	Redundant Logical Tunnel Configuration . . . . .	29
	Redundant Logical Tunnel Failure Detection and Failover . . . . .	30
	Configuring Redundant Logical Tunnels . . . . .	31
	Example: Configuring Redundant Logical Tunnels . . . . .	32
<b>Chapter 6</b>	<b>Configuring Layer 2 Ethernet Services over GRE Tunnel Interfaces . . . . .</b>	<b>43</b>
	Layer 2 Services Over GRE Tunnel Interfaces on MX Series with MS-MPCs and MS-MICs . . . . .	43
	Format of GRE Frames and Processing of GRE Interfaces for Layer 2 Ethernet Packets . . . . .	44
	Guidelines for Configuring Layer 2 Ethernet Traffic Over GRE Tunnels . . . . .	45
	Sample Scenarios of Configuring Layer 2 Ethernet Traffic Over GRE Tunnels . . . . .	46
	Configuring Layer 2 Services over GRE Logical Interfaces in Bridge Domains . . . . .	47
	Example: Configuring Layer 2 Services Over GRE Logical Interfaces in Bridge Domains . . . . .	48
<b>Chapter 7</b>	<b>Understanding Default PIM Tunnel Configurations . . . . .</b>	<b>55</b>
	Configuring PIM Tunnels . . . . .	55
<b>Chapter 8</b>	<b>Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces . . . . .</b>	<b>57</b>
	Configuring Virtual Loopback Tunnels for VRF Table Lookup . . . . .	57
	Configuring Tunnel Interfaces for Routing Table Lookup . . . . .	59
	Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup . . . . .	59
	Example: Virtual Routing and Forwarding (VRF) and Service Configuration . . . . .	60
<b>Chapter 9</b>	<b>Enabling a VPN to Travel Through a Non-MPLS Network Using Dynamic Tunnels . . . . .</b>	<b>63</b>
	Configuring Dynamic Tunnels . . . . .	63
<b>Part 2</b>	<b>Encryption Services</b>	
<b>Chapter 10</b>	<b>Overview . . . . .</b>	<b>67</b>
	Encryption Overview . . . . .	67
	Configuring an ES Tunnel Interface for a Layer 3 VPN . . . . .	67
<b>Chapter 11</b>	<b>Sending Encrypted Traffic Through Tunnels . . . . .</b>	<b>69</b>
	Configuring Encryption Interfaces . . . . .	69
	Specifying the Security Association Name for Encryption Interfaces . . . . .	70
	Configuring the MTU for Encryption Interfaces . . . . .	70
	Example: Configuring an Encryption Interface . . . . .	70
	Configuring Filters for Traffic Transiting the ES PIC . . . . .	71
	Traffic Overview . . . . .	71
	Configuring the Security Association . . . . .	72
	Configuring an Outbound Traffic Filter . . . . .	73
	Example: Configuring an Outbound Traffic Filter . . . . .	73

	Applying the Outbound Traffic Filter . . . . .	74
	Example: Applying the Outbound Traffic Filter . . . . .	74
	Configuring an Inbound Traffic Filter . . . . .	74
	Example: Configuring an Inbound Traffic Filter . . . . .	75
	Applying the Inbound Traffic Filter to the Encryption Interface . . . . .	75
	Example: Applying the Inbound Traffic Filter to the Encryption Interface . . . . .	75
<b>Chapter 12</b>	<b>Configuring Redundancy in Case of Service Failure . . . . .</b>	<b>77</b>
	Configuring ES PIC Redundancy . . . . .	77
	Example: Configuring ES PIC Redundancy . . . . .	77
	Configuring IPsec Tunnel Redundancy . . . . .	78
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 13</b>	<b>Configuration Statements . . . . .</b>	<b>83</b>
	address (Interfaces) . . . . .	84
	allow-fragmentation . . . . .	85
	apply-groups-except . . . . .	85
	backup-destination . . . . .	86
	backup-interface . . . . .	86
	copy-tos-to-outer-ip-header . . . . .	87
	core-facing . . . . .	87
	destination (Interfaces) . . . . .	88
	destination (Routing Instance) . . . . .	89
	destination (Tunnel Remote End) . . . . .	89
	destination-networks . . . . .	90
	do-not-fragment . . . . .	91
	dynamic-tunnels . . . . .	92
	es-options . . . . .	93
	family . . . . .	94
	family bridge . . . . .	95
	family bridge (GRE Interfaces) . . . . .	96
	filter . . . . .	97
	hold-time (OAM) . . . . .	97
	interfaces . . . . .	98
	ipsec-sa . . . . .	98
	keepalive-time . . . . .	99
	key . . . . .	100
	multicast-only . . . . .	100
	peer-unit . . . . .	101
	peer-certificate-type . . . . .	101
	reassemble-packets . . . . .	102
	redundancy-group (Interfaces) . . . . .	102
	redundancy-group (Logical Tunnels) . . . . .	103
	routing-instance . . . . .	104
	routing-instances . . . . .	104
	routing-options . . . . .	105
	source . . . . .	105
	source . . . . .	106

	source-address . . . . .	106
	ttl . . . . .	107
	tunnel . . . . .	108
	tunnel . . . . .	109
	unit (Interfaces) . . . . .	110
	unit (Interfaces) . . . . .	111
<b>Chapter 14</b>	<b>Operational Commands . . . . .</b>	<b>113</b>
	clear ike security-associations . . . . .	114
	clear ipsec security-associations . . . . .	115
	request ipsec switch . . . . .	117
	request security certificate enroll (Signed) . . . . .	118
	request security certificate enroll (Unsigned) . . . . .	120
	request security key-pair . . . . .	121
	request system certificate add . . . . .	122
	show ike security-associations . . . . .	123
	show interfaces (Encryption) . . . . .	127
	show interfaces (GRE) . . . . .	133
	show interfaces (IP-over-IP) . . . . .	142
	show interfaces (Logical Tunnel) . . . . .	146
	show interfaces (Multicast Tunnel) . . . . .	151
	show interfaces (PIM) . . . . .	156
	show interfaces (Virtual Loopback Tunnel) . . . . .	160
	show ipsec certificates . . . . .	165
	show ipsec redundancy . . . . .	168
	show ipsec security-associations . . . . .	170
	show system certificate . . . . .	173
<b>Part 4</b>	<b>Index</b>	
	Index . . . . .	177

# List of Figures

<b>Part 1</b>	<b>Tunnel Services</b>	
<b>Chapter 5</b>	<b>Connecting Logical Systems Using Logical Tunnel Interfaces . . . . .</b>	<b>25</b>
	Figure 1: Redundant Logical Tunnels . . . . .	29
	Figure 2: Redundant Logical Tunnels . . . . .	33
<b>Part 2</b>	<b>Encryption Services</b>	
<b>Chapter 11</b>	<b>Sending Encrypted Traffic Through Tunnels . . . . .</b>	<b>69</b>
	Figure 3: Example: IPsec Tunnel Connecting Security Gateways . . . . .	71
<b>Chapter 12</b>	<b>Configuring Redundancy in Case of Service Failure . . . . .</b>	<b>77</b>
	Figure 4: IPsec Tunnel Redundancy . . . . .	78



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xi</b>
	Table 1: Notice Icons . . . . .	xiii
	Table 2: Text and Syntax Conventions . . . . .	xiv
<b>Part 1</b>	<b>Tunnel Services</b>	
<b>Chapter 1</b>	<b>Overview</b> . . . . .	<b>3</b>
	Table 3: Tunnel Interface Types . . . . .	3
<b>Chapter 8</b>	<b>Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces</b> . . . . .	<b>57</b>
	Table 4: Methods for Configuring Egress Filtering . . . . .	57
<b>Part 3</b>	<b>Configuration Statements and Operational Commands</b>	
<b>Chapter 14</b>	<b>Operational Commands</b> . . . . .	<b>113</b>
	Table 5: show ike security-associations Output Fields . . . . .	123
	Table 6: Encryption show interfaces Output Fields . . . . .	127
	Table 7: GRE show interfaces Output Fields . . . . .	134
	Table 8: IP-over-IP show interfaces Output Fields . . . . .	142
	Table 9: Logical Tunnel show interfaces Output Fields . . . . .	146
	Table 10: Multicast Tunnel show interfaces Output Fields . . . . .	152
	Table 11: PIM show interfaces Output Fields . . . . .	156
	Table 12: Virtual Loopback Tunnel show interfaces Output Fields . . . . .	160
	Table 13: show ipsec certificates Output Fields . . . . .	165
	Table 14: show ipsec redundancy Output Fields . . . . .	168
	Table 15: show ipsec security-associations Output Fields . . . . .	170
	Table 16: show system certificate Output Fields . . . . .	173



# About the Documentation

- Documentation and Release Notes on page xi
- Supported Platforms on page xi
- Using the Examples in This Manual on page xi
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

## 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:

- M Series
- MX Series
- T 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.xml;
    }
  }
}
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.xml; }
```

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 *CLI User Guide*.

## Documentation Conventions

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

Table 1: Notice Icons







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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> <li>To configure a stub area, include the <b>stub</b> statement at the [edit protocols ospf area area-id] hierarchy level.</li> <li>The console port is labeled <b>CONSOLE</b>.</li> </ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub</b> <default-metric <i>metric</i> >;
(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.	<b>broadcast</b>   <b>multicast</b>  ( <i>string1</i>   <i>string2</i>   <i>string3</i> )
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members</b> [ <i>community-ids</i> ]
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

---

#### GUI Conventions

---

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

Convention	Description	Examples
<b>Bold text like this</b>	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> <li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li> <li>To cancel the configuration, click <b>Cancel</b>.</li> </ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software version (if applicable).

## Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service 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:  
<http://kb.juniper.net/InfoCenter/>
- 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

# Tunnel Services

- [Overview on page 3](#)
- [Encapsulating One Protocol over Another Using GRE Interfaces on page 9](#)
- [Encapsulating One IP Packet over Another Using IP-IP Interfaces on page 15](#)
- [Filtering Unicast Packets Through Multicast Tunnel Interfaces on page 17](#)
- [Connecting Logical Systems Using Logical Tunnel Interfaces on page 25](#)
- [Configuring Layer 2 Ethernet Services over GRE Tunnel Interfaces on page 43](#)
- [Understanding Default PIM Tunnel Configurations on page 55](#)
- [Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces on page 57](#)
- [Enabling a VPN to Travel Through a Non-MPLS Network Using Dynamic Tunnels on page 63](#)



## CHAPTER 1

# Overview

- [Tunnel Services Overview on page 3](#)
- [Configuring Tunnel Interfaces on MX Series Routers on page 6](#)
- [Configuring Tunnel Interfaces on T4000 Routers on page 7](#)

## Tunnel Services Overview

---

By encapsulating arbitrary packets inside a transport protocol, tunneling provides a private, secure path through an otherwise public network. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and MPLS. If you have a Tunnel Physical Interface Card (PIC) installed in your M Series or T Series router, you can configure unicast, multicast, and logical tunnels.

You can configure two types of tunnels for VPNs: one to facilitate routing table lookups and another to facilitate VPN routing and forwarding instance (VRF) table lookups.

For information about encryption interfaces, see [“Configuring Encryption Interfaces” on page 69](#) and the *Junos OS Administration Library for Routing Devices*. For information about VPNs, see the *Junos OS VPNs Library for Routing Devices*. For information about MPLS, see the *Junos OS MPLS Applications Library for Routing Devices*.

On SRX Series devices, Generic Routing Encapsulation (GRE) and IP-IP tunnels use internal interfaces, `gr-0/0/0` and `ip-0/0/0`, respectively. The Junos OS creates these interfaces at system bootup; they are not associated with physical interfaces.

The Juniper Networks Junos OS supports the tunnel types shown in the following table.

**Table 3: Tunnel Interface Types**

Interface	Description
<code>gr-0/0/0</code>	<p>Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol over another routing protocol.</p> <p>Within a router, packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then re-encapsulated with another protocol packet to complete the GRE. The GRE interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform GRE.</p>

Table 3: Tunnel Interface Types (*continued*)

Interface	Description
<b>gre</b>	Internally generated GRE interface. This interface is generated by the Junos OS to handle GRE. You cannot configure this interface.
<b>ip-0/0/0</b>	<p>Configurable IP-over-IP encapsulation (also called IP tunneling) interface. IP tunneling allows the encapsulation of one IP packet over another IP packet.</p> <p>Packets are routed to an internal interface where they are encapsulated with an IP packet and then forwarded to the encapsulating packet's destination address. The IP-IP interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform IP tunneling.</p>
<b>ipip</b>	Internally generated IP-over-IP interface. This interface is generated by the Junos OS to handle IP-over-IP encapsulation. It is not a configurable interface.
<b>lt-0/0/0</b>	<p>The <b>lt</b> interface on M Series and T Series routers supports configuration of logical systems—the capability to partition a single physical router into multiple logical devices that perform independent routing tasks.</p> <p>On SRX Series devices, the <b>lt</b> interface is a configurable logical tunnel interface that interconnects logical systems. See the <i>Junos OS Logical Systems Configuration Guide for Security Devices</i>.</p>
<b>mt-0/0/0</b>	<p>Internally generated multicast tunnel interface. Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a 224/8-or-greater prefix, the packet is dropped and a counter is incremented.</p> <p>Within a router, packets are routed to this internal interface for multicast filtering. The multicast tunnel interface is an internal interface only and is not associated with a physical interface. If your router has a Tunnel Services PIC, the Junos OS automatically configures one multicast tunnel interface (<b>mt-</b>) for each virtual private network (VPN) you configure. You do not need to configure multicast tunnel interfaces. However, you can configure properties on <b>mt-</b> interfaces, such as the <b>multicast-only</b> statement.</p>
<b>mtun</b>	Internally generated multicast tunnel interface. This interface is generated by the Junos OS to handle multicast tunnel services. It is not a configurable interface.
<b>pd-0/0/0</b>	<p>Configurable Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a router, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform PIM de-encapsulation.</p> <p><b>NOTE:</b> On SRX Series devices, this interface type is <b>ppd0</b>.</p>

Table 3: Tunnel Interface Types (*continued*)

Interface	Description
<b>pe-0/0/0</b>	<p>Configurable PIM encapsulation interface. In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a router, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical interface. You must configure the interface for it to perform PIM encapsulation.</p> <p><b>NOTE:</b> On SRX Series devices, this interface type is <b>ppe0</b>.</p>
<b>pimd</b>	Internally generated PIM de-encapsulation interface. This interface is generated by the Junos OS to handle PIM de-encapsulation. It is not a configurable interface.
<b>pime</b>	Internally generated PIM encapsulation interface. This interface is generated by the Junos OS to handle PIM encapsulation. It is not a configurable interface.
<b>vt-0/0/0</b>	<p>Configurable virtual loopback tunnel interface. Facilitates VRF table lookup based on MPLS labels. This interface type is supported on M Series and T Series routers, but not on SRX Series devices.</p> <p>To configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels, you specify a virtual loopback tunnel interface name and associate it with a routing instance that belongs to a particular routing table. The packet loops back through the virtual loopback tunnel for route lookup.</p>

Starting in Junos OS Release 15.1, you can configure Layer 2 Ethernet services over GRE interfaces (**gr-fpc/pic/port** to use GRE encapsulation). To enable Layer 2 Ethernet packets to be terminated on GRE tunnels, you must configure the bridge domain protocol family on the **gr-** interfaces and associate the **gr-** interfaces with the bridge domain. You must configure the GRE interfaces as core-facing interfaces, and they must be access or trunk interfaces. To configure the bridge domain family on **gr-** interfaces, include the **family bridge** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. To associate the **gr-** interface with a bridge domain, include the **interface gr-fpc/pic/port** statement at the **[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]** hierarchy level. You can associate GRE interfaces in a bridge domain with the corresponding VLAN ID or list of VLAN IDs in a bridge domain by including the **vlan-id (all | none | number)** statement or the **vlan-id-list [vlan-id-numbers]** statement at the **[edit bridge-domains bridge-domain-name]** hierarchy level. The VLAN IDs configured for the bridge domain must match with the VLAN IDs that you configure for GRE interfaces by using the **vlan-id (all | none | number)** statement or the **vlan-id-list [vlan-id-numbers]** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. You can also configure GRE interfaces within a bridge domain associated with a virtual switch instance. Layer 2 Ethernet packets over GRE tunnels are also supported with the GRE key option. The gre-key match condition allows a user to match against the GRE key field, which is an optional field in GRE encapsulated packets. The key can be matched as a single key value, a range of key values, or both.

**Related Documentation**

- [GRE Keepalive Time Overview on page 9](#)
- [Configuring Unicast Tunnels on page 17](#)
- [Restricting Tunnels to Multicast Traffic on page 23](#)
- [Configuring Tunnel Interfaces on MX Series Routers on page 6](#)
- [Configuring Tunnel Interfaces on T4000 Routers on page 7](#)

---

## Configuring Tunnel Interfaces on MX Series Routers

---

Because the MX Series routers do not support Tunnel Services PICs, you create tunnel interfaces on MX Series routers by including the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic number {
    tunnel-services {
      bandwidth (1g | 10g | 20g | 40g);
    }
  }
}
```

**fpc slot-number** is the slot number of the DPC, MPC, or MIC. On the MX80 router, the range is 0 through 1. On other MX series routers, if two SCBs are installed, the range is 0 through 11. If three SCBs are installed, the range is 0 through 5 and 7 through 11.

The **pic number** On MX80 routers, if the FPC is 0, the PIC number can only be 0. If the FPC is 1, the PIC range is 0 through 3. For all other MX series routers, the range is 0 through 3.

**bandwidth (1g | 10g | 20g | 40g)** is the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.



**NOTE:** When you use MPCs and MICs, tunnel interfaces are soft interfaces and allow as much traffic as the forwarding-path allows, so it is advantageous to setup tunnel services without artificially limiting traffic by use of the **bandwidth** option. However, you *must* specify **bandwidth** when configuring tunnel services for MX Series routers with DPCs or FPCs. The GRE key option is not supported on the tunnel interfaces for DPCs on MX960 routers.

Bandwidth rates of 20 gigabits per second and 40 gigabits per second require use of an MX Series router with the 100-Gigabit Ethernet Modular Port Concentrator (MPC) and the 100-Gigabit CFP MIC.

---

- **1g** indicates that the maximum bandwidth available for tunnel traffic is 1 gigabit per second. For MPCs and MICs, this bandwidth is not reserved for tunnel traffic and can be shared by the network interfaces. For DPCs, this bandwidth is reserved and cannot be shared by the network interfaces.
- **10g** indicates that the maximum bandwidth available for tunnel traffic is 10 gigabits per second. For MPCs and MICs, this bandwidth is not reserved for tunnel traffic and can be shared by the network interfaces. For DPCs, this bandwidth is reserved and cannot be shared by the network interfaces.
- **20g** indicates that the maximum bandwidth available for tunnel traffic is 20 gigabits per second. For MPCs and MICs, this bandwidth is not reserved for tunnel traffic and can be shared by the network interfaces. For DPCs, this bandwidth is reserved and cannot be shared by the network interfaces.
- **40g** indicates that the maximum bandwidth available for tunnel traffic is 40 gigabits per second. For MPCs and MICs, this bandwidth is not reserved for tunnel traffic and can be shared by the network interfaces. For DPCs, this bandwidth is reserved and cannot be shared by the network interfaces.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 10-Gigabit Ethernet 4-port DPC.

To verify that the tunnel interfaces have been created, issue the **show interfaces terse** operational mode command. For more information, see the [CLI Explorer](#). The bandwidth that you specify determines the port number of the tunnel interfaces that are created. When you specify a bandwidth of **1g**, the port number is always 10. When you specify any other bandwidth, the port number is always 0.



**NOTE:** Ingress queueing and tunnel services cannot be configured on the same MPC as it causes PFE forwarding to stop. Each feature can, however, be configured and used separately.

#### Related Documentation

- *Example: Configuring Tunnel Interfaces on a Gigabit Ethernet 40-Port DPC*
- *Example: Configuring Tunnel Interfaces on a 10-Gigabit Ethernet 4-Port DPC*
- *Example: Configuring Tunnel Interfaces on the MPC3E*
- *bandwidth (Tunnel Services)*
- *tunnel-services*
- *[edit chassis] Hierarchy Level*

## Configuring Tunnel Interfaces on T4000 Routers

To create tunnel interfaces on a T4000 Core Router, include the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic number {
    tunnel-services {
      bandwidth bandwidth-value;
    }
  }
}
```

**fpc slot-number** denotes the slot number of the FPC. On the T4000 router, the range is 0 through 7.



NOTE:

- This applies only to the T4000 Type 5 FPC. If any other type of FPC is configured in this slot, this configuration is ignored and no tunnel physical interface is created.
- When you use Type 5 FPCs, the tunnel interfaces are soft interfaces and allow as much traffic as the forwarding-path allows. So, it is advantageous to setup tunnel services without artificially limiting traffic by setting the **bandwidth** statement.

---

**pic number** on the T4000 router is 0 or 1.

**bandwidth bandwidth-value** is the amount of bandwidth to reserve for the tunnel traffic on each Packet Forwarding Engine. The bandwidth value accepted includes every multiple of 10g up to 100g.

If you specify a bandwidth that is not compatible, tunnel services are not activated. For example, you cannot specify a bandwidth of 1 Gbps for a Packet Forwarding Engine on a 100-Gigabit Ethernet PIC with CFP.

To verify that the tunnel interfaces have been created, issue the **show interfaces terse** operational mode command. For more information, see the *Junos Interfaces Command Reference*.

**Related  
Documentation**

- *bandwidth (Tunnel Services)*
- *tunnel-services*
- *[edit chassis] Hierarchy Level*

## CHAPTER 2

# Encapsulating One Protocol over Another Using GRE Interfaces

- [GRE Keepalive Time Overview on page 9](#)
- [Configuring GRE Keepalive Time on page 9](#)
- [Enabling Fragmentation on GRE Tunnels on page 12](#)

## GRE Keepalive Time Overview

Generic routing encapsulation (GRE) tunnel interfaces do not have a built-in mechanism for detecting when a tunnel is down. You can enable keepalive messages to serve as the detection mechanism.

Keepalives can be configured on the physical or on the logical interface. If configured on the physical interface, keepalives are sent on all logical interfaces that are part of the physical interface. If configured on a individual logical interface, keepalives are only sent to that logical interface. In addition to configuring a keepalive, you must configure the hold time.

### **Related Documentation**

- [Configuring GRE Keepalive Time on page 9](#)
- [keepalive-time on page 99](#)
- [hold-time on page 97](#)

## Configuring GRE Keepalive Time

- [Configuring Keepalive Time and Hold time for a GRE Tunnel Interface on page 10](#)
- [Display GRE Keepalive Time Configuration on page 10](#)
- [Display Keepalive Time Information on a GRE Tunnel Interface on page 11](#)

## Configuring Keepalive Time and Hold time for a GRE Tunnel Interface

You can configure the keepalives on a generic routing encapsulation (GRE) tunnel interface by including both the **keepalive-time** statement and the **hold-time** statement at the **[edit protocols oam gre-tunnel interface *interface-name*]** hierarchy level.



**NOTE:** For proper operation of keepalives on a GRE interface, you must also include the **family inet** statement at the **[edit interfaces *interface-name* unit *unit*]** hierarchy level. If you do not include this statement, the interface is marked as down.

To configure a GRE tunnel interface:

1. Configure the GRE tunnel interface at **[edit interfaces *interface-name* unit *unit-number*]** hierarchy level, where the interface name is gr-x/y/z, and the family is set as **inet**.

```
user@host# set interfaces interface-name unit unit-number family family-name
```

2. Configure the rest of the GRE tunnel interface options as explained in *Configuring a GRE Tunnel Interface Between a PE and CE Router* or *Configuring a GRE Tunnel Interface Between PE Routers* based on requirement.

To configure keepalive time for a GRE tunnel interface:

1. Configure the Operation, Administration, and Maintenance (OAM) protocol at the **[edit protocols]** hierarchy level for the GRE tunnel interface.

```
[edit]
user@host# edit protocols oam
```

2. Configure the GRE tunnel interface option for OAM protocol.

```
[edit protocols oam]
user@host# edit gre-tunnel interface interface-name
```

3. Configure the keepalive time from 1 through 50 seconds for the GRE tunnel interface.

```
[edit protocols oam gre-tunnel interface interface-name]
user@host# set keepalive-time seconds
```

4. Configure the hold time from 5 through 250 seconds. Note that the hold time must be at least twice the keepalive time.

```
[edit protocols oam gre-tunnel interface interface-name]
user@host# set hold-time seconds
```

## Display GRE Keepalive Time Configuration

**Purpose** Display the configured keepalive time value as 10 and hold time value as 30 on a GRE tunnel interface (for example, gr-1/1/10.1).

**Action** To display the configured values on the GRE tunnel interface, run the **show oam gre-tunnel** command at the **[edit protocols]** hierarchy level:





---

**NOTE:**

When the hold time expires:

- The GRE tunnel will stay up even though the interface cannot send or receive traffic.
  - The Link status will be Up and the Gre keepalives adjacency state will be Down.
- 

**Meaning** The current status information of a GRE tunnel interface with keepalive time and hold time parameters is displayed as expected when the hold time expires.

**Related Documentation**

- [GRE Keepalive Time Overview on page 9](#)
- [keepalive-time on page 99](#)
- [hold-time on page 97](#)

---

## Enabling Fragmentation on GRE Tunnels

---

To enable fragmentation of IPv4 packets in generic routing encapsulation (GRE) tunnels, include the **clear-dont-fragment-bit** statement and a maximum transmission unit (MTU) setting for the tunnel as part of an existing GRE configuration at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
gr-fpc/pic/port {
  unit logical-unit-number {
    clear-dont-fragment-bit;
    ...
  }
  family inet {
    mtu 1000;
    ...
  }
}
```

This statement clears the Don't Fragment (DF) bit in the packet header, regardless of the packet size. If the packet size exceeds the tunnel MTU value, the packet is fragmented before encapsulation. The maximum MTU size configurable on the AS or Multiservices PIC is 9192 bytes.



---

**NOTE:** The **clear-dont-fragment-bit** statement is supported only on MX Series routers and all M Series routers except the M320 router.

---

Fragmentation is enabled only on IPv4 packets being encapsulated in IPv4-based GRE tunnels.



.....

**NOTE:** This configuration is supported only on GRE tunnels on AS or Multiservices interfaces. If you commit `gre-fragmentation` as the encapsulation type on a standard Tunnel PIC interface, the following console log message appears when the PIC comes online:

`gr-fpc/pic/port: does not support this encapsulation`

The Packet Forwarding Engine updates the IP identification field in the outer IP header of GRE-encapsulated packets, so that reassembly of the packets is possible after fragmentation. The previous CLI constraint check that required you to configure either the `clear-dont-fragment-bit` statement or a tunnel key with the `allow-fragmentation` statement is no longer enforced.

When you configure the `clear-dont-fragment-bit` statement on an interface with the MPLS protocol family enabled, you must specify an MTU value. This MTU value must not be greater than maximum supported value (9192).

.....

**Related Documentation**

- [Configuring Unicast Tunnels on page 17](#)



## CHAPTER 3

# Encapsulating One IP Packet over Another Using IP-IP Interfaces

- [Configuring IPv6-over-IPv4 Tunnels on page 15](#)
- [Example: Configuring an IPv6-over-IPv4 Tunnel on page 15](#)

## Configuring IPv6-over-IPv4 Tunnels

---

If you have a Tunnel PIC installed in your M Series or T Series router, you can configure IPv6-over-IPv4 tunnels. To define a tunnel, you configure a unicast tunnel across an existing IPv4 network infrastructure. IPv6/IPv4 packets are encapsulated in IPv4 headers and sent across the IPv4 infrastructure through the configured tunnel. You manually configure configured tunnels on each end point.

On SRX Series devices, Generic Routing Encapsulation (GRE) and IP-IP tunnels use internal interfaces, `gr-0/0/0` and `ip-0/0/0`, respectively. The Junos OS creates these interfaces at system bootup; they are not associated with a physical interface.

IPv6-over-IPv4 tunnels are defined in RFC 2893, *Transition Mechanisms for IPv6 Hosts and Routers*. For information about configuring a unicast tunnel, see [“Configuring Unicast Tunnels” on page 17](#). For an IPv6-over-IPv4 tunnel configuration example, see [“Example: Configuring an IPv6-over-IPv4 Tunnel” on page 15](#).

### Related Documentation

- [Tunnel Services Overview on page 3](#)
- [Example: Configuring an IPv6-over-IPv4 Tunnel on page 15](#)

## Example: Configuring an IPv6-over-IPv4 Tunnel

---

Configure a tunnel on both sides of the connection.

Configuration on	[edit]
Router 1	interfaces {
	gr-1/0/0 {
	unit 0 {
	tunnel {
	source 10.19.2.1;
	destination 10.19.3.1;
	}
	}

	<pre>        family inet6 {             address 2001:DB8:1:1/126;         }     } }</pre>
<b>Configuration on Router 2</b>	<pre>[edit] interfaces {     gr-1/0/0 {         unit 0 {             tunnel {                 source 10.19.3.1;                 destination 10.19.2.1;             }             family inet6 {                 address 2001:DB8:2:1/126;             }         }     } }</pre>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tunnel Services Overview on page 3</a></li><li>• <a href="#">Configuring IPv6-over-IPv4 Tunnels on page 15</a></li></ul>

## CHAPTER 4

# Filtering Unicast Packets Through Multicast Tunnel Interfaces

- [Configuring Unicast Tunnels on page 17](#)
- [Examples: Configuring Unicast Tunnels on page 22](#)
- [Restricting Tunnels to Multicast Traffic on page 23](#)

## Configuring Unicast Tunnels

---

To configure a unicast tunnel, you configure a **gr-** interface (to use GRE encapsulation) or an **ip-** interface (to use IP-IP encapsulation) and include the **tunnel** and **family** statements:

```
gr-fpc/pic/port or ip-fpc/pic/port {  
  unit logical-unit-number {  
    copy-tos-to-outer-ip-header;  
    reassemble-packets;  
    tunnel {  
      allow-fragmentation;  
      destination destination-address;  
      do-not-fragment;  
      key number;  
      routing-instance {  
        destination routing-instance-name;  
      }  
      source address;  
      ttl number;  
    }  
    family family {  
      address address {  
        destination address;  
      }  
    }  
  }  
}
```

You can configure these statements at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

You can configure multiple logical units for each GRE or IP-IP interface, and you can configure only one tunnel per unit.



**NOTE:** On M Series and T Series routers, you can configure the interface on a service PIC or a tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC.

Each tunnel interface must be a point-to-point interface. Point to point is the default interface connection type, so you do not need to include the **point-to-point** statement in the logical interface configuration.

You must specify the tunnel's destination and source addresses. The remaining statements are optional.



**NOTE:** For transit packets exiting the tunnel, forwarding path features, such as reverse path forwarding (RPF), forwarding table filtering, source class usage, destination class usage, and stateless firewall filtering, are not supported on the interfaces you configure as tunnel sources, but are supported on tunnel-pic interfaces.

However, class-of-service (CoS) information obtained from the GRE or IP-IP header is carried over the tunnel and is used by the re-entering packets. For more information, see the *Class of Service Feature Guide for Routing Devices*.

To prevent an invalid configuration, the Junos OS disallows setting the address specified by the source or destination statement at the [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* tunnel] hierarchy level to be the same as the interface's own subnet address, specified by the address statement at the [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* family *family-name*] hierarchy level.

To set the time-to-live (TTL) field that is included in the encapsulating header, include the **ttl** statement. If you explicitly configure a TTL value for the tunnel, you must configure it to be one larger than the number of hops in the tunnel. For example, if the tunnel has seven hops, you must configure a TTL value of 8.

You must configure at least one family on the logical interface. To enable MPLS over GRE tunnel interfaces, you must include the **family mpls** statement in the GRE interface configuration. In addition, you must include the appropriate statements at the [edit **protocols**] hierarchy level to enable Resource Reservation Protocol (RSVP), MPLS, and label-switched paths (LSPs) over GRE tunnels. Unicast tunnels are bidirectional.

A configured tunnel cannot go through Network Address Translation (NAT) at any point along the way to the destination. For more information, see [“Examples: Configuring Unicast Tunnels” on page 22](#) and the *MPLS Applications Feature Guide for Routing Devices*.

For a GRE tunnel, the default is to set the ToS bits in the outer IP header to all zeros. To have the Routing Engine copy the ToS bits from the inner IP header to the outer, include

the **copy-tos-bits-to-outer-ip-header** statement. (This inner-to-outer ToS bits copying is already the default behavior for IP-IP tunnels.)

For GRE tunnel interfaces on Adaptive Services or Multiservices interfaces, you can configure additional tunnel attributes, as described in the following sections:

- [Configuring a Key Number on GRE Tunnels on page 19](#)
- [Enabling Fragmentation on GRE Tunnels on page 20](#)
- [Specifying an MTU Setting for the Tunnel on page 20](#)
- [Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 21](#)
- [Configuring Packet Reassembly on page 21](#)

## Configuring a Key Number on GRE Tunnels

For Adaptive Services and Multiservices interfaces on M Series and T Series routers, you can assign a key value to identify an individual traffic flow within a GRE tunnel, as defined in RFC 2890, *Key and Sequence Number Extensions to GRE*. However, only one key is allowed for each tunnel source and destination pair.

Each IP version 4 (IPv4) packet entering the tunnel is encapsulated with the GRE tunnel key value. Each IPv4 packet exiting the tunnel is verified by the GRE tunnel key value and de-encapsulated. The Adaptive Services or Multiservices PIC drops packets that do not match the configured key value.

To assign a key value to a GRE tunnel interface, include the **key** statement:

```
key number;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* tunnel]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

The key number can be 0 through 4,294,967,295. You must configure the same GRE tunnel key value on tunnel endpoints.

The following example illustrates the use of the key statement in a GRE tunnel configuration:

```
interfaces {
  gr-1/2/0 {
    unit 0 {
      tunnel {
        source 10.58.255.193;
        destination 10.58.255.195;
        key 1234;
      }
    }
    ...
    family inet {
      mtu 1500;
      address 10.200.0.1/30;
    }
  }
}
```

```

    ...
  }
}
}

```

## Enabling Fragmentation on GRE Tunnels

For GRE tunnel interfaces on Adaptive Services and Multiservices interfaces only, you can enable fragmentation of IPv4 packets in GRE tunnels.

By default, IPv4 traffic transmitted over GRE tunnels is not fragmented. To enable fragmentation of IPv4 packets in GRE tunnels, include the **clear-dont-fragment-bit** statement:

```
clear-dont-fragment-bit;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

When you include the **clear-dont-fragment-bit** statement in the configuration, the don't-fragment (DF) bit is cleared on all packets, even packets that do not exceed the tunnel maximum transmission unit (MTU). If the packet's size exceeds the tunnel's MTU value, the packet is fragmented before encapsulation. If the packet's size does not exceed the tunnel's MTU value, the packet is not fragmented.



**NOTE:** The Packet Forwarding Engine updates the IP identification field in the outer IP header of GRE-encapsulated packets, so that reassembly of the packets is possible after fragmentation. The previous CLI constraint check that required you to configure either the **clear-dont-fragment-bit** statement or a tunnel key with the **allow-fragmentation** statement is no longer enforced.

You can also clear the DF bit in packets transmitted over IP Security (IPsec) tunnels. For more information, see *Configuring IPsec Rules*.

## Specifying an MTU Setting for the Tunnel

To enable key numbers and fragmentation on GRE tunnels (as described in [“Configuring a Key Number on GRE Tunnels” on page 19](#) and [“Enabling Fragmentation on GRE Tunnels” on page 20](#)), you must also specify an MTU setting for the tunnel.

To specify an MTU setting for the tunnel, include the **mtu** statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* family inet]

- [edit logical-system *logical-system-name* interfaces *gr-fpc/pic/port* unit *logical-unit-number* family inet]

For more information about MTU settings, see the *Junos OS Network Interfaces Library for Routing Devices*.

## Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header

Unlike IP-IP tunnels, GRE tunnels do not copy the ToS bits to the outer IP header by default. To have the Routing Engine copy the inner ToS bits to the outer IP header (which is required for some tunneled routing protocols) on packets sent by the Routing Engine, include the **copy-tos-to-outer-ip-header** statement at the logical unit hierarchy level of a GRE interface. This example copies the inner ToS bits to the outer IP header on a GRE tunnel:

```
[edit interfaces]
gr-0/0/0 {
  unit 0 {
    copy-tos-to-outer-ip-header;
    family inet;
  }
}
```

## Configuring Packet Reassembly

On GRE tunnel interfaces only, you can enable reassembly of fragmented tunnel packets. To activate this capability, include the **reassemble-packets** statement:

```
reassemble-packets;
```

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For each tunnel you configure on the interface, you can enable or disable fragmentation of GRE packets by including the **allow-fragmentation** or **do-not-fragment** statement:

```
allow-fragmentation;
do-not-fragment;
```

You can configure these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* tunnel]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* tunnel]

If you configure **allow-fragmentation** on a tunnel, it clears the DF bit in the outer IP header, enabling post fragmentation of GRE-encapsulated packets if the packet size exceeds the maximum transmission unit (MTU) value for the egress interface. By default, packets that exceed the MTU size are dropped and post fragmentation of GRE packets is disabled.



**NOTE:** Whenever you configure **allow-fragmentation** on a tunnel, you must also include either the **tunnel key** or the **clear-dont-fragment-bit** statement. This configuration enables the router to send affected packets to the PIC so that the correct IP header can be placed in the fragments. Otherwise, on the reassembly side some packets might be lost when fragments arrive in the PIC out of sequence at high speeds.

Starting with Junos OS Release 14.2, you can configure the generic routing encapsulation (GRE) tunnel interfaces on MX Series routers with Trio-based FPCs (MPCs) to support IP packet reassembly. The IP packet is fragmented over a GRE tunnel when the packet size exceeds the maximum transmission unit (MTU) defined for the connection. When a GRE packet passes through a link that contains an MTU less than that of the packet size, the packet is fragmented. The other endpoint of the tunnel needs to reassemble it before GRE processing starts. You can configure the GRE interfaces to reassemble the fragmented packets at the endpoint of the tunnel before they can be further processed on the network by including the **reassemble-packets** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. You can view the reassembly statistics by using the **show services inline ip-reassembly stastics <fpc fpc-slot | pfe pfe-slot>** command. MX Series routers can be used as the data center edge or data center interconnect devices in network environments as needed to enable interconnection of data centers across an IP backbone using GRE encapsulation. Inline IP reassembly is supported on MX80, MX240, MX480, MX960, MX2010, MX2020, and MX104 routers. The line modules compatible with the corresponding MX Series routers that support the reassembly of GRE packets are MPC1, MPC2, MPC3, MPC4, and MPC-16X10GE. Until Junos OS Release 14.1, reassembly of IP fragments received at GRE tunnels is supported only on MX Series routers with MS-DPCs.

- Related Documentation**
- [Tunnel Services Overview on page 3](#)
  - [Examples: Configuring Unicast Tunnels on page 22](#)

## Examples: Configuring Unicast Tunnels

Configure two unnumbered IP-IP tunnels:

```
[edit interfaces]
ip-0/3/0 {
  unit 0 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.253;
    }
    family inet;
  }
  unit 1 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.254;
    }
  }
}
```

```

        family inet;
    }
}

```

Configure numbered tunnel interfaces by including an address at the **[edit interfaces ip-0/3/0 unit (0 | 1) family inet]** hierarchy level:

```

[edit interfaces]
ip-0/3/0 {
  unit 0 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.253;
    }
    family inet {
      address 10.5.5.1/30;
    }
  }
  unit 1 {
    tunnel {
      source 192.168.4.18;
      destination 192.168.4.254;
    }
    family inet {
      address 10.6.6.100/30;
    }
  }
}

```

Configure an MPLS over GRE tunnel by including the **family mpls** statement at the **[edit interfaces gr-1/2/0 unit 0]** hierarchy level:

```

[edit interfaces]
gr-1/2/0 {
  unit 0 {
    tunnel {
      source 192.168.1.1;
      destination 192.168.1.2;
    }
    family inet {
      address 10.1.1.1/30;
    }
    family mpls;
  }
}

```

- Related Documentation**
- [Tunnel Services Overview on page 3](#)
  - [Configuring Unicast Tunnels on page 17](#)

## Restricting Tunnels to Multicast Traffic

For interfaces that carry IPv4 or IP version 6 (IPv6) traffic, you can configure a tunnel interface to allow multicast traffic only. To configure a multicast-only tunnel, include the **multicast-only** statement:

**multicast-only;**

You can configure this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a 224/8 or greater prefix, the packet is dropped and a counter is incremented.

You can configure this property on GRE, IP-IP, PIM, and multicast tunnel (**mt**) interfaces only.



**NOTE:** If your router has a Tunnel Services PIC, the Junos OS automatically configures one multicast tunnel interface (**mt**) for each virtual private network (VPN) you configure. You do not need to configure multicast tunnel interfaces.

---

**Related  
Documentation**

- [Tunnel Services Overview on page 3](#)
- [Configuring Unicast Tunnels on page 17](#)

## CHAPTER 5

# Connecting Logical Systems Using Logical Tunnel Interfaces

- [Configuring Logical Tunnel Interfaces on page 25](#)
- [Example: Configuring Logical Tunnels on page 26](#)
- [Redundant Logical Tunnels Overview on page 28](#)
- [Configuring Redundant Logical Tunnels on page 31](#)
- [Example: Configuring Redundant Logical Tunnels on page 32](#)

## Configuring Logical Tunnel Interfaces

---

Logical tunnel (**lt-**) interfaces provide quite different services depending on the host router:

- On M Series, MX Series, and T Series routers, logical tunnel interfaces allow you to connect logical systems, virtual routers, or VPN instances. M Series and T Series routers must be equipped with a Tunnel Services PIC or an Adaptive Services Module (only available on M7i routers). MX Series routers must be equipped with a Trio MPC/MIC module. For more information about connecting these applications, see the *Junos OS VPNs Library for Routing Devices*.
- On SRX Series Services Gateways, the logical tunnel interface is used to interconnect logical systems. See the *Junos OS Logical Systems Configuration Guide for Security Devices*.

For M Series, MX Series, and T Series routers, see the following section:

- [Connecting Logical Systems on page 25](#)

## Connecting Logical Systems

To connect two logical systems, you configure a logical tunnel interface on both logical systems. Then you configure a peer relationship between the logical tunnel interfaces, thus creating a point-to-point connection.

To configure a point-to-point connection between two logical systems, configure the logical tunnel interface by including the **lt-fpc/pic/port** statement:

```
lt-fpc/pic/port {
```

```
unit logical-unit-number {  
    encapsulation encapsulation;  
    peer-unit unit-number; # peering logical system unit number  
    dlcid dlcid-number;  
    family (inet | inet6 | iso | mpls);  
}
```

You can include this statement at the following hierarchy levels:

- **[edit interfaces]**
- **[edit logical-systems *logical-system-name* interfaces]**

When configuring logical tunnel interfaces, note the following:

- You can configure each logical tunnel interface with one of the following encapsulation types: Ethernet, Ethernet circuit cross-connect (CCC), Ethernet VPLS, Frame Relay, Frame Relay CCC, VLAN, VLAN CCC, or VLAN VPLS.
- You can configure the IP, IPv6, International Organization for Standardization (ISO), or MPLS protocol family.
- The peering logical interfaces must belong to the same logical tunnel interface derived from the Tunnel Services PIC or Adaptive Services Module.
- You can configure only one peer unit for each logical interface. For example, unit 0 cannot peer with both unit 1 and unit 2.
- To enable the logical tunnel interface, you must configure at least one physical interface statement.
- Logical tunnels are not supported with Adaptive Services, Multiservices, or Link Services PICs (but they are supported on the Adaptive Services Module on M7i routers, as noted above).
- On M Series routers other than the M40e router, logical tunnel interfaces require an Enhanced Flexible PIC Concentrator (FPC).
- On MX Series routers, logical tunnel interfaces require Trio MPC/MIC modules. They do not require a Tunnel Services PIC in the same system.

For more information about configuring logical systems, see the *Junos OS Routing Protocols Library for Routing Devices*.

**Related  
Documentation**

- [Tunnel Services Overview on page 3](#)
- [Example: Configuring Logical Tunnels on page 26](#)

---

## Example: Configuring Logical Tunnels

Configure three logical tunnels:

```
[edit interfaces]  
lt-4/2/0 {  
    description "Logical tunnel interface connects three logical systems";
```

```

}
[edit logical-systems]
lr1 {
  interfaces lt-4/2/0 {
    unit 12 {
      peer-unit 21; #Peering with lr2
      encapsulation frame-relay;
      dlci 612;
      family inet;
    }
    unit 13 {
      peer-unit 31; #Peering with lr3
      encapsulation frame-relay-ccc;
      dlci 613;
    }
  }
}
lr2 {
  interfaces lt-4/2/0 {
    unit 21 {
      peer-unit 12; #Peering with lr1
      encapsulation frame-relay-ccc;
      dlci 612;
    }
    unit 23 {
      peer-unit 32; #Peering with lr3
      encapsulation frame-relay;
      dlci 623;
    }
  }
}
lr3 {
  interfaces lt-4/2/0 {
    unit 31 {
      peer-unit 13; #Peering with lr1
      encapsulation frame-relay;
      dlci 613;
      family inet;
    }
    unit 32 {
      peer-unit 23; #Peering with lr2
      encapsulation frame-relay-ccc;
      dlci 623;
    }
  }
}

```

- Related Documentation**
- [Tunnel Services Overview on page 3](#)
  - [Configuring Logical Tunnel Interfaces on page 25](#)

## Redundant Logical Tunnels Overview

---

You can connect two devices, such as an access-facing device and a core-facing device, through logical tunnels. To provide redundancy for the tunnels, you can create and configure multiple physical logical tunnels and add them to a virtual redundant logical tunnel.



**NOTE:** Redundant logical tunnels are supported only on MX Series routers with MPCs. MX Series Virtual Chassis does not support redundant logical tunnels.

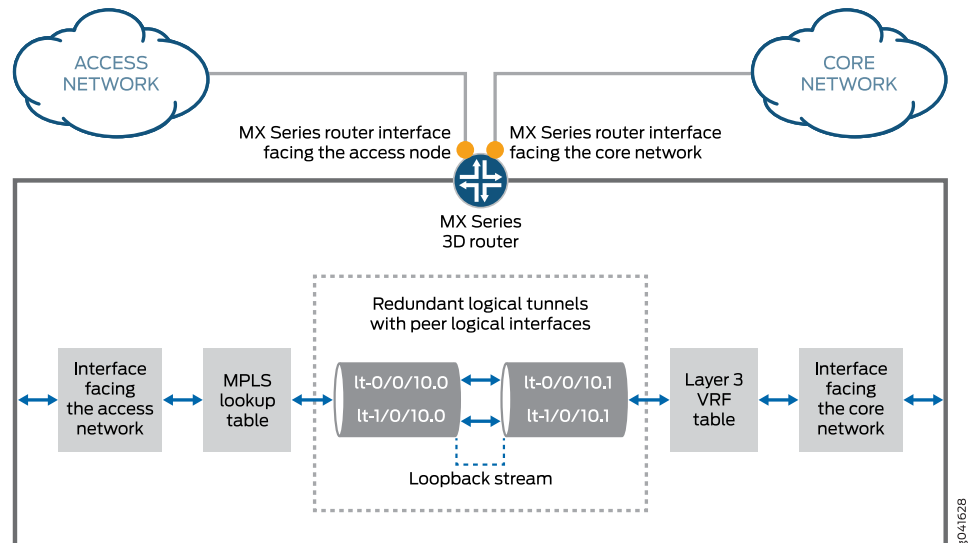


**NOTE:** Redundant logical tunnels are not supported currently in MX Series Virtual Chassis.

For example, in an MPLS access network, you can configure multiple pseudowires between an access node and an MX Series router with MPCs and add them to a redundant logical tunnel. You can then add multiple logical tunnels to the redundant logical tunnel.

[Figure 1 on page 29](#) shows a redundant logical tunnel between the access node and the MX Series router.

**Figure 1: Redundant Logical Tunnels**



The redundant logical tunnel has peer logical interfaces at each end, `lt0.0` and `lt0.1`. You can configure router features on these interfaces for the redundant logical tunnel and its members.

Each member logical tunnel has peer logical interfaces. In [Figure 1 on page 29](#), `lt-0/0/10.0` and `lt-0/0/10.1` are peers.

The MX Series router performs IP lookup in the Layer 3 VPN routing and forwarding (VRF) table on the router where the pseudowires that are grouped in logical tunnels terminate.

## Redundant Logical Tunnel Configuration

In Junos Releases 13.3R1, 13.3R2, and 14.1R1 you can create up to 16 redundant logical tunnels, depending on the number of Packet Forwarding Engines and the number of loopback interfaces on each Packet Forwarding Engine on your device. For Junos Release 13.3R3, 14.1R2, and 14.2 the valid range for device-count is from 1 to 255.

You can add up to 32 logical tunnels as members of a redundant logical tunnel.

When you add more than two members to the redundant logical tunnel, they are in active mode. The traffic is load-balanced over all the tunnel members.

When you add only two members to the redundant logical tunnel, you can configure the members in one of these ways:

- Both members in active mode
- One member in active mode and the other in backup mode

## Redundant Logical Tunnel Failure Detection and Failover

A logical tunnel fails and is removed from the redundant logical tunnel group, and the backup logical tunnel becomes active due to one of these events:

- A hardware failure on the MPC module occurs.
- An MPC failure occurs due to a microkernel crash.
- The MPC module is administratively shut down and removed from the redundant logical tunnel.
- A power failure on the MPC module occurs.



.....

**NOTE:** You can decrease the time it takes for failure detection and failover to occur. Configure the `enhanced-ip` statement at the `[edit chassis network-services]` hierarchy level to enable Packet Forwarding Engine liveliness detection.

.....

### Related Documentation

- [Example: Configuring Redundant Logical Tunnels on page 32](#)
- [Pseudowire Subscriber Logical Interfaces Overview](#)
- [Configuring Logical Tunnel Interfaces on page 25](#)
- [Configuring Redundant Logical Tunnels on page 31](#)
- [Configuring a Pseudowire Subscriber Logical Interface Device](#)

## Configuring Redundant Logical Tunnels

Use redundant logical tunnels to provide redundancy for logical tunnels between two devices, such as an access-facing device and a core-facing device.

When configuring redundant logical tunnel interfaces, note the following:

- In Junos OS Release 13.3 or later, you can configure redundant logical tunnels only on MX Series routers with MPCs.

In Junos Releases 13.3R1, 13.3R2, and 14.1R1 you can create up to 16 redundant logical tunnels, depending on the number of Packet Forwarding Engines and the number of loopback interfaces on each Packet Forwarding Engine on your device. For Junos Release 13.3R3, 14.1R2, and 14.2 the valid range for device-count is from 1 to 255. The command is shown below

**set chassis redundancy-group interface-type redundant-logical-tunnel device-count *[number]*;**

You can add up to 32 logical tunnels as members.

- When a logical tunnel with an existing configuration joins a redundant logical tunnel, you must configure the redundant logical tunnel with the settings from the existing configuration.
- You can add member logical tunnels to a parent logical tunnel for redundancy.
- When you add more than two logical tunnels to the redundant logical tunnel, the members are in active mode by default.
- When you add only two members, you can configure the members in one of these ways:
  - Both members in active mode
  - One member in active mode and the other in backup mode

To configure a redundant logical tunnel between two devices:

1. Create the logical tunnel and redundant logical tunnel interfaces.

```
[edit chassis]
user@host# set redundancy-group interface-type redundant-logical-tunnel
device-count count
user@host# set fpc slot-number pic number tunnel-services bandwidth 1g
```

2. Bind the member logical tunnels to the redundant logical tunnel.

```
[edit interfaces]
user@host# set interface-name redundancy-group member-interface interface-name
```

3. Configure the redundant logical tunnel interfaces.
4. Attach the redundant logical tunnel interface to a Layer 2 circuit.
5. Add the peer redundant logical tunnel interface to a Layer 3 VRF instance.
6. Configure MPLS and LDP in the pseudowires and the Layer 3 VPN.

```
[edit protocols]
user@host# set mpls no-cspf
user@host# set mpls interface all
user@host# set ldp interface all
```

7. Configure BGP in the Layer 3 VPN.
8. Configure OSPF on the core-facing interfaces and the router local loopback interface.
9. Set the policy options for BGP.
10. Set the router ID and the autonomous system (AS) number.

**Related  
Documentation**

- [Example: Configuring Redundant Logical Tunnels on page 32](#)
- [Redundant Logical Tunnels Overview on page 28](#)

---

## Example: Configuring Redundant Logical Tunnels

---

This example shows how to configure redundant logical tunnels in an MPLS access network.

- [Requirements on page 32](#)
- [Overview on page 32](#)
- [Configuration on page 33](#)
- [Verification on page 39](#)

### Requirements

In Junos OS Release 13.3 or later, you can configure redundant logical tunnels only on MX Series routers with MPCs.

### Overview

When a logical tunnel with an existing configuration joins a redundant logical tunnel, you must configure the redundant logical tunnel with the settings from the existing configuration.

You can add member logical tunnels to a parent logical tunnel for redundancy.

On MX Series routers with MPCs, you can configure redundant logical tunnels as follows:

- In Junos Releases 13.3R1, 13.3R2, and 14.1R1 you can create up to 16 redundant logical tunnels, depending on the number of Packet Forwarding Engines and the number of loopback interfaces on each Packet Forwarding Engine on your device. For Junos Releases 13.3R3, 14.1R2, and 14.2 the valid range for device-count is from 1 to 255. The command is shown below

```
set chassis redundancy-group interface-type redundant-logical-tunnel device-count  
[number];
```

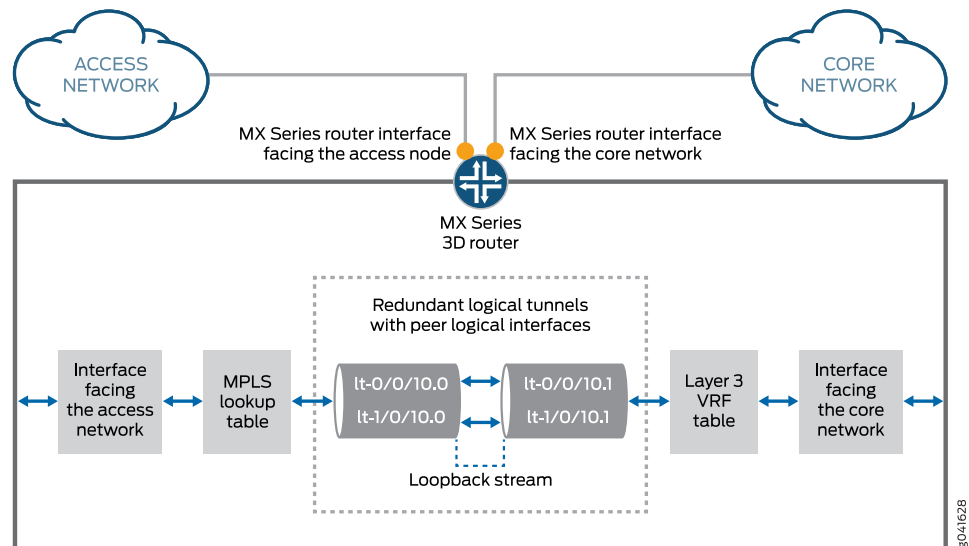
- You can add up to 32 logical tunnels as members.

- When you add more than two logical tunnels to a redundant logical tunnel, the members are in active mode by default.
- When you add only two members, you can configure the members in one of these ways:
  - Both members in active mode
  - One member in active mode and the other in backup mode

## Topology

Figure 1 on page 29 shows a redundant logical tunnel between the access node and the MX Series router in an MPLS access network.

Figure 2: Redundant Logical Tunnels



The redundant logical tunnel has peer logical interfaces at each end, `rlt0.0` and `rlt0.1`. You can configure router features on these interfaces for the redundant logical tunnel and its members.

Each member logical tunnel has peer logical interfaces on the access-facing and core-facing devices. In Figure 1 on page 29, `lt-0/0/10.0` and `lt-0/0/10.1` are peers.

The MX Series router performs IP lookup in the Layer 3 VPN routing and forwarding (VRF) table on the router where the pseudowires that are grouped in logical tunnels terminate.

## Configuration

### CLI Quick Configuration

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

```
set chassis redundancy-group interface-type redundant-logical-tunnel device-count 4
set chassis fpc 1 pic 0 tunnel-services bandwidth 1g
```

```
set chassis fpc 1 pic 2 tunnel-services bandwidth 1g
set interfaces rlt0 redundancy-group member-interface lt-1/0/10
set interfaces rlt0 redundancy-group member-interface lt-2/0/10
set interfaces rlt0 unit 0 description "Towards Layer 2 Circuit"
set interfaces rlt0 unit 0 encapsulation vlan-ccc
set interfaces rlt0 unit 0 vlan-id 600
set interfaces rlt0 unit 0 peer-unit 1
set interfaces rlt0 unit 0 family ccc
set interfaces rlt0 unit 1 description "Towards Layer 3 VRF"
set interfaces rlt0 unit 1 encapsulation vlan
set interfaces rlt0 unit 1 vlan-id 600
set interfaces rlt0 unit 1 peer-unit 0
set interfaces rlt0 unit 1 family inet address 10.10.10.2/24
set protocols l2circuit neighbor 2.2.2.2 interface rlt0.0 virtual-circuit-id 100
set protocols l2circuit neighbor 2.2.2.2 interface rlt0.0 no-control-word
set routing-instances pe-vrf instance-type vrf
set routing-instances pe-vrf interface rlt0.1
set routing-instances pe-vrf route-distinguisher 65056:1
set routing-instances pe-vrf vrf-import VPN-A-Import
set routing-instances pe-vrf vrf-export VPN-A-Export
set routing-instances pe-vrf vrf-table-label
set routing-instances pe-vrf protocols ospf export VPN-A-Import
set routing-instances pe-vrf protocols ospf area 0.0.0.0 interface rlt0.1
set protocols mpls no-cspf
set protocols mpls interface all
set protocols ldp interface all
set protocols bgp export local-routes
set protocols bgp group internal type internal
set protocols bgp group internal local-address 3.3.3.3
set protocols bgp group internal family inet any
set protocols bgp group internal family inet-vpn unicast
set protocols bgp group internal neighbor 4.4.4.4
set protocols ospf area 0.0.0.0 interface ge-5/3/8.0
set protocols ospf area 0.0.0.0 interface ge-5/2/5.0
set protocols ospf area 0.0.0.0 interface lo0.3 passive
set policy-options policy-statement VPN-A-Export term a then community add VPN-A
set policy-options policy-statement VPN-A-Export term a then accept
set policy-options policy-statement VPN-A-Export term b then reject
set policy-options policy-statement VPN-A-Import term a from protocol bgp
set policy-options policy-statement VPN-A-Import term a from community VPN-A
set policy-options policy-statement VPN-A-Import term a then accept
set policy-options policy-statement VPN-A-Import term b then reject
set policy-options policy-statement local-routes then accept
set policy-options community VPN-A members target:100:100
set routing-options router-id 3.3.3.3
set routing-options autonomous-system 65056
```

#### Step-by-Step Procedure

In this example, all the logical tunnels are in active mode.

1. Create the logical tunnel and redundant logical tunnel interfaces.

```
[edit chassis]
user@host# set redundancy-group interface-type redundant-logical-tunnel
device-count 4
user@host# set fpc 1 pic 0 tunnel-services bandwidth 1g
user@host# set fpc 1 pic 2 tunnel-services bandwidth 1g
```

2. Bind the member logical tunnels to the redundant logical tunnel.
 

```
[edit interfaces]
user@host# set rlt0 redundancy-group member-interface lt-1/0/10
user@host# set rlt0 redundancy-group member-interface lt-2/0/10
```
3. Configure the redundant logical tunnel interfaces.
 

```
[edit interfaces]
user@host# set rlt0 unit 0 description "Towards Layer 2 Circuit"
user@host# set rlt0 unit 0 encapsulation vlan-ccc
user@host# set rlt0 unit 0 vlan-id 600
user@host# set rlt0 unit 0 peer-unit 1
user@host# set rlt0 unit 0 family ccc

user@host# set rlt0 unit 1 description "Towards Layer 3 VRF"
user@host# set rlt0 unit 1 encapsulation vlan
user@host# set rlt0 unit 1 vlan-id 600
user@host# set rlt0 unit 1 peer-unit 0
user@host# set rlt0 unit 1 family inet address 10.10.10.2/24
```
4. Attach rlt0.0 to a Layer 2 circuit.
 

```
[edit protocols]
user@host# set l2circuit neighbor 2.2.2.2 interface rlt0.0 virtual-circuit-id 100
user@host# set l2circuit neighbor 2.2.2.2 interface rlt0.0 no-control-word
```
5. Add rlt0.1 to a Layer 3 VRF instance.
 

```
[edit routing-instances]
user@host# set pe-vrf instance-type vrf
user@host# set pe-vrf interface rlt0.1
user@host# set pe-vrf route-distinguisher 65056:1
user@host# set pe-vrf vrf-import VPN-A-Import
user@host# set pe-vrf vrf-export VPN-A-Export
user@host# set pe-vrf vrf-table-label
user@host# set pe-vrf protocols ospf export VPN-A-Import
user@host# set pe-vrf protocols ospf area 0.0.0.0 interface rlt0.1
```
6. Configure MPLS and LDP in the pseudowires and the Layer 3 VPN.
 

```
[edit protocols]
user@host# set mpls no-cspf
user@host# set mpls interface all
user@host# set ldp interface all
```
7. Configure BGP in the Layer 3 VPN.
 

```
[edit protocols]
user@host# set bgp export local-routes
user@host# set bgp group internal type internal
user@host# set bgp group internal local-address 3.3.3.3
user@host# set bgp group internal family inet any
user@host# set bgp group internal family inet-vpn unicast
user@host# set bgp group internal neighbor 4.4.4.4
```
8. Configure OSPF on the core-facing interfaces and the router local loopback interface.
 

```
[edit protocols]
user@host# set ospf area 0.0.0.0 interface ge-5/3/8.0
```

```
user@host# set ospf area 0.0.0.0 interface ge-5/2/5.0
user@host# set ospf area 0.0.0.0 interface lo0.3 passive
```

9. Set the policy options for BGP.

```
[edit policy-options]
user@host# set policy-statement VPN-A-Export term a then community add VPN-A
user@host# set policy-statement VPN-A-Export term a then accept
user@host# set policy-statement VPN-A-Export term b then reject
user@host# set policy-statement VPN-A-Import term a from protocol bgp
user@host# set policy-statement VPN-A-Import term a from community VPN-A
user@host# set policy-statement VPN-A-Import term a then accept
user@host# set policy-statement VPN-A-Import term b then reject
user@host# set policy-statement local-routes then accept
user@host# set community VPN-A members target:100:100
```

10. Set the router ID and the autonomous system (AS) number.

```
[edit routing-options]
user@host# set router-id 3.3.3.3
user@host# set autonomous-system 65056
```

---

## Results

From configuration mode, confirm your configuration by entering the following commands:

- **show chassis**
- **show interfaces**
- **show policy-options**
- **show protocols**
- **show routing-instances**
- **show routing-options**

If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show chassis
redundancy-group {
  interface-type {
    redundant-logical-tunnel {
      device-count 4;
    }
  }
}
fpc 1 {
  pic 0 {
    tunnel-services {
      bandwidth 1g;
    }
  }
}
fpc 1 {
  pic 2 {
```

```
tunnel-services {
  bandwidth 1g;
}
}

user@host# show interfaces rlt0
redundancy-group {
  member-interface lt-1/0/10;
  member-interface lt-2/0/10;
}
unit 0 {
  description "Towards Layer 2 Circuit";
  encapsulation vlan-ccc;
  vlan-id 600;
  peer-unit 1;
  family ccc;
}
unit 1 {
  description "Towards Layer 3 VRF";
  encapsulation vlan;
  vlan-id 600;
  peer-unit 0;
  family inet {
    address 10.10.10.2/24;
  }
}

user@host# show protocols l2circuit
neighbor 2.2.2.2 {
  interface rlt0.0 {
    virtual-circuit-id 100;
    no-control-word;
  }
}

user@host# show protocols
mpls {
  no-cspf;
  interface all;
}
bgp {
  export local-routes;
  group internal {
    type internal;
    local-address 3.3.3.3;
    family inet {
      any;
    }
    family inet-vpn {
      unicast;
    }
    neighbor 4.4.4.4;
  }
}
ospf {
  area 0.0.0.0 {
```

```
interface ge-5/3/8.0;
interface ge-5/2/5.0;
interface lo0.3 {
    passive;
}
}
}
ldp {
    interface all;
}
l2circuit {
    neighbor 2.2.2.2 {
        interface rlt0.0 {
            virtual-circuit-id 100;
            no-control-word;
        }
    }
}
}

user@host# routing-instances
pe-vrf {
    instance-type vrf;
    interface rlt0.1;
    route-distinguisher 65056:1;
    vrf-import VPN-A-Import;
    vrf-export VPN-A-Export;
    vrf-table-label;
    protocols {
        ospf {
            export VPN-A-Import;
            area 0.0.0.0 {
                interface rlt0.1;
            }
        }
    }
}

user@host# policy-options
policy-statement VPN-A-Export {
    term a {
        then {
            community add VPN-A;
            accept;
        }
    }
    term b {
        then reject;
    }
}
policy-statement VPN-A-Import {
    term a {
        from {
            protocol bgp;
            community VPN-A;
        }
        then accept;
    }
}
```

```

    term b {
      then reject;
    }
  }
  policy-statement local-routes {
    then accept;
  }
  community VPN-A members target:100:100;

user@host# routing-options
router-id 3.3.3.3;
autonomous-system 65056;

```

## Verification

Confirm that the configuration is working properly.

- [Verifying the Redundant Logical Tunnel Configuration on page 39](#)
- [Verifying the Layer 2 Circuit on page 39](#)
- [Verifying OSPF Neighbors on page 40](#)
- [Verifying the BGP Group on page 40](#)
- [Verifying the BGP Routes in the Routing Table on page 40](#)

### Verifying the Redundant Logical Tunnel Configuration

**Purpose** Verify that the redundant logical tunnel with the child logical tunnel interfaces are created with the correct encapsulations.

**Action** user@host# run show interfaces terse | match rlt0

lt-1/0/10.0	up	up	container-->	rlt0.0	
lt-1/0/10.1	up	up	container-->	rlt0.1	
lt-2/0/10.0	up	up	container-->	rlt0.0	
lt-2/0/10.1	up	up	container-->	rlt0.1	
rlt0	up	up			
rlt0.0	up	up	ccc		
rlt0.1	up	up	inet		10.10.10.2/24

### Verifying the Layer 2 Circuit

**Purpose** Verify that the Layer 2 circuit is up.

**Action** user@host# run show l2circuit connections  
Layer-2 Circuit Connections:

Legend for connection status (St)

EI -- encapsulation invalid	NP -- interface h/w not present
MM -- mtu mismatch	Dn -- down
EM -- encapsulation mismatch	VC-Dn -- Virtual circuit Down
CM -- control-word mismatch	Up -- operational
VM -- vlan id mismatch	CF -- Call admission control failure
OL -- no outgoing label	IB -- TDM incompatible bitrate
NC -- intf encaps not CCC/TCC	TM -- TDM misconfiguration
BK -- Backup Connection	ST -- Standby Connection
CB -- rcvd cell-bundle size bad	SP -- Static Pseudowire
LD -- local site signaled down	RS -- remote site standby
RD -- remote site signaled down	HS -- Hot-standby Connection
XX -- unknown	

Legend for interface status

Up -- operational

Dn -- down

Neighbor: 2.2.2.2

Interface	Type	St	Time last up	# Up trans
rlt0.0(vc 100)	rmt	Up	Aug 8 00:28:04 2013	1
Remote PE: 2.2.2.2, Negotiated control-word: No				
Incoming label: 299776, Outgoing label: 299776				
Negotiated PW status TLV: No				
Local interface: rlt0.0, Status: Up, Encapsulation: VLAN				

### Verifying OSPF Neighbors

**Purpose** Verify that routers are adjacent and able to exchange OSPF data.

**Action** user@host# run show ospf neighbor

Address	Interface	State	ID	Pri	Dead
30.30.30.2	ge-5/2/5.0	Full	4.4.4.4	128	38
20.20.20.1	ge-5/3/8.0	Full	2.2.2.2	128	38

### Verifying the BGP Group

**Purpose** Verify that the BGP group is created.

**Action** user@host# run show bgp group internal

Group Type: Internal	AS: 65056	Local AS: 65056
Name: internal	Index: 0	Flags: <Export Eval>
Export: [ local-routes ]		
Holdtime: 0		
Total peers: 1	Established: 1	
4.4.4.4+179		
inet.0: 1/6/3/0		
inet.2: 0/0/0/0		
bgp.l3vpn.0: 2/2/2/0		
pe-vrf.inet.0: 2/2/2/0		

### Verifying the BGP Routes in the Routing Table

**Purpose** Verify that the BGP routes are in the pe-vrf.inet.0 routing table.

**Action**    user@host# run show route protocol bgp table pe-vrf.inet.0  
pe-vrf.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)  
+ = Active Route, - = Last Active, \* = Both

50.50.50.0/24        \*[BGP/170] 01:18:14, localpref 100, from 4.4.4.4  
                      AS path: I, validation-state: unverified  
                      > to 30.30.30.2 via ge-5/2/5.0, Push 16

50.50.51.0/24        \*[BGP/170] 01:18:14, MED 2, localpref 100, from 4.4.4.4  
                      AS path: I, validation-state: unverified  
                      > to 30.30.30.2 via ge-5/2/5.0, Push 16

- Related Documentation**
- [Configuring Redundant Logical Tunnels on page 31](#)
  - [Redundant Logical Tunnels Overview on page 28](#)



## CHAPTER 6

# Configuring Layer 2 Ethernet Services over GRE Tunnel Interfaces

- [Layer 2 Services Over GRE Tunnel Interfaces on MX Series with MS-MPCs and MS-MICs on page 43](#)
- [Format of GRE Frames and Processing of GRE Interfaces for Layer 2 Ethernet Packets on page 44](#)
- [Guidelines for Configuring Layer 2 Ethernet Traffic Over GRE Tunnels on page 45](#)
- [Sample Scenarios of Configuring Layer 2 Ethernet Traffic Over GRE Tunnels on page 46](#)
- [Configuring Layer 2 Services over GRE Logical Interfaces in Bridge Domains on page 47](#)
- [Example: Configuring Layer 2 Services Over GRE Logical Interfaces in Bridge Domains on page 48](#)

### Layer 2 Services Over GRE Tunnel Interfaces on MX Series with MS-MPCs and MS-MICs

Starting in Junos OS Release 15.1, you can configure Layer 2 Ethernet services over GRE interfaces (***gr-fpc/pic/port*** to use GRE encapsulation).

The outputs of the **show bridge mac-table** and **show vpls mac-table** commands have been enhanced to display the MAC addresses learned on a GRE logical interface and the status of MAC address learning properties in the MAC address and MAC flags fields. Also, the **L2 Routing Instance** and **L3 Routing Instance** fields are added to the output of the **show interfaces gr** command to display the names of the routing instances associated with the GRE interfaces are displayed.

To enable Layer 2 Ethernet packets to be terminated on GRE tunnels, you must configure the bridge domain protocol family on the **gr-** interfaces and associate the **gr-** interfaces with the bridge domain. You must configure the GRE interfaces as core-facing interfaces, and they must be access or trunk interfaces. To configure the bridge domain family on **gr-** interfaces, include the **family bridge** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. To associate the **gr-** interface with a bridge domain, include the **interface gr-fpc/pic/port** statement at the **[edit routing-instances routing-instance-name bridge-domains bridge-domain-name]** hierarchy level.

You can associate GRE interfaces in a bridge domain with the corresponding VLAN ID or list of VLAN IDs in a bridge domain by including the **vlan-id (all | none | number)** statement or the **vlan-id-list [vlan-id-numbers]** statement at the **[edit bridge-domains**

**bridge-domain-name**] hierarchy level. The VLAN IDs configured for the bridge domain must match with the VLAN IDs that you configure for GRE interfaces by using the **vlan-id (all | none | number)** statement or the **vlan-id-list [ *vlan-id-numbers* ]** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level. You can also configure GRE interfaces within a bridge domain associated with a virtual switch instance. Layer 2 Ethernet packets over GRE tunnels are also supported with the GRE key option. The gre-key match condition allows a user to match against the GRE key field, which is an optional field in GRE encapsulated packets. The key can be matched as a single key value, a range of key values, or both.

**Related** •  
**Documentation**

## Format of GRE Frames and Processing of GRE Interfaces for Layer 2 Ethernet Packets

The GRE frame contains the outer MAC header, outer IP header, GRE header, original layer 2 frame, and frame checksum (FCS).

In the outer MAC header, the following fields are present:

- The outer destination MAC address is set as the next-hop MAC address
- The outer source MAC address is set as the source address of the MX Series router that functions as the gateway
- The outer VLAN tag information

In the outer IP header, the following fields are contained:

- The outer source address is set as the source address of the MX Series router gateway
- The outer destination address is set as the remote GRE tunnel address
- The outer protocol type is set as 47 (encapsulation type is GRE)
- The VLAN ID configuration within the bridge domain updates the VLAN ID of the original Layer 2 header

The gr-interface supports GRE encapsulation over IP, which is supported over Layer 3 over GRE. Support for bridging over GRE enables you to configure bridge domain families on gr- interfaces and also enable integrated routing and bridging (IRB) on gr- interfaces. The device control daemon (dcd) that controls the physical and logical interface processes enables the processing of bridge domain families under the GRE interfaces. The kernel supports IRB to send and receive packets on IRB interfaces.

The Packet Forwarding Engine supports the Layer 2 encapsulation and decapsulation over GRE interfaces. The chassis daemon is responsible for creating the GRE physical interface when an FPC comes online and triggering the deletion of the GRE interfaces when the FPC goes offline. The kernel receives the GRE logical interface that is added over the underlying physical interface and propagates the GRE logical interface to other clients, including the Packet Forwarding Engine to create the Layer 2 over GRE data path in the hardware. In addition, it adds the GRE logical interface into a bridge domain. The Packet Forwarding Engine receives interprocess communication message (IPC) from

the kernel and adds the interface into the forwarding plane. The existing MTU size for the GRE interface is increased by 22 bytes for the L2 header addition (6 DMAC + 6 SMAC + 4 CVLAN + 4 SVLAN + 2 EtherType)

**Related  
Documentation**

## Guidelines for Configuring Layer 2 Ethernet Traffic Over GRE Tunnels

Observe the following guidelines while configuring Layer 2 packets to be transmitted over GRE tunnel interfaces on MX Series routers with MS-MPCs and MS-MICs:

- For integrated routing and bridging (IRB) to work, at least one Layer 2 interface must be up and active, and it must be associated with the bridge domain as an IRB interface along with a GRE Layer 2 logical interface. This configuration is required to leverage the existing broadcast infrastructure of Layer 2 with IRB.
- Graceful Routing Engine switchover (GRES) is supported and unified ISSU is not currently supported.
- MAC addresses learned from the GRE networks are learned on the bridge domain interfaces associated with the gr-fpc/pic/port.unit logical interface. The MAC addresses are learned on GRE logical interfaces and the Layer 2 token used for forwarding is the token associated with the GRE interface. Destination MAC lookup yields an L2 token, which causes the next-hop lookup. This next-hop is used to forward the packet.
- The GRE tunnel encapsulation and decapsulation next-hops are enhanced to support this functionality. The GRE tunnel encapsulation next-hop is used to encapsulate the outer IP and GRE headers with the incoming L2 packet. The GRE tunnel decapsulation next-hop is used to decapsulate the outer IP and GRE headers, parse the inner Layer 2 packet, and set the protocol as bridge for further bridge domain properties processing in the Packet Forwarding Engine.
- The following packet flows are supported:
  - As part of Layer 2 packet flows, L2 unicast from L2 to GRE, L2 unicast from GRE to L2, Layer 2 broadcast, unknown unicast, and multicast (L2 BUM) from L2 to GRE, and L2 BUM from GRE to L2 are supported.
  - As part of Layer 3 packet flows, L3 Unicast from L2 to GRE, L3 Unicast from GRE to L2, L3 Multicast from L2 to GRE, L3 Multicast from GRE to L2, and L3 Multicast from Internet to GRE and L2 are supported.
- Support for L2 control protocols is not available.
- At the GRE decapsulation side, packets destined to the tunnel IP are processed and decapsulated by the forwarding plane, and inner L2 packets are processed. MAC learned packets are generated for control plane processing for newly learned MAC entries. However, these entries are throttled for MAC learning.
- 802.1x authentication can be used to validate the individual endpoints and protect them from unauthorized access.

- With the capability to configure bridge domain families on GRE tunnel interfaces, the maximum number of GRE interfaces supported depends on the maximum number of tunnel devices allocated, where each tunnel device can host up to 4000 logical interfaces. The maximum number of logical tunnel interfaces supported is not changed with the support for Layer 2 GRE tunnels. For example, in a 4x10 MIC on MX960 routers, 8000 tunnel logical interfaces can be created.
- The tunnels are pinned to a specific Packet Forwarding Engine instance.
- Statistical information for GRE Layer 2 tunnels is displayed in the output of the **show interfaces gr-fpc/pic/port** command.
- Only trunk and access mode configuration is supported for the bridge family of GRE interfaces; subinterface style configuration is not supported.
- You can enable a connection to a traditional Layer 2 network. Connection to a VPLS network is not currently supported. IRB in bridge domains with GRE interfaces is supported.
- Virtual switch instances are supported.
- Configuration of the GRE Key and using it to perform the hash load-balancing at the GRE tunnel-initiated and transit routers is supported.

**Related  
Documentation** •

---

## Sample Scenarios of Configuring Layer 2 Ethernet Traffic Over GRE Tunnels

You can configure Layer 2 Ethernet services over GRE interfaces (**gr-fpc/pic/port** to use GRE encapsulation). This topic contains the following sections that illustrate sample network deployments that support Layer 2 packets over GRE tunnel interfaces:

### GRE Tunnels with an MX Series Router as the Gateway in Layer 3

You can configure an MX Series router as the gateway that contains GRE tunnels configured to connect to legacy switches on the one end and to a Layer 3 network on the other end. The Layer 3 network in turn can be linked with multiple servers on a LAN where the GRE tunnel is terminated from the WAN.

### GRE Tunnels With an MX Series Router as the Gateway and Aggregator

You can configure an MX Series router as the gateway with GRE tunnels configured and also with aggregation specified. The gateway can be connected to legacy switches on one end of the network and the aggregator can be connected to a top-of-rack (ToR) switch, as a QFX Series device, which handles GRE tunneled packets with load balancing. The ToR switch can be connected, in turn, over a Layer 3 GRE tunnel to several servers in data centers.

### GRE Tunnels with MX Series Gateways for Enterprise and Data Center Servers

You can configure an MX Series router as the gateway with GRE tunnels configured. Over the Internet, GRE tunnels connect multiple gateways, which are MX routers, to servers

in enterprises where the GRE tunnel is terminated from the WAN on one end, and to servers in data centers on the other end.

The following configuration scenarios are supported for Layer 2 Ethernet over GRE tunnels:

- In a Layer 2 Ethernet over GRE with VPLS environment, an MX Series router supports Layer 2 over GRE tunnels (without the MPLS layer) and terminate these tunnels into a VPLS or an routed VLAN interface (RVI) into a L3VPN. The tunnels serve to cross the cable modem termination system (CMTS) and cable modem CM infrastructure transparently, up to the MX Series router that serves as the gateway. Every GRE tunnel terminates over a VLAN interface, a VPLS instance, or an IRB interface.
- In a Layer 2 Ethernet over GRE without VPLS environment, Layer 2 VPN encapsulations are needed for conditions that do not involve these protocols. Certain data center users terminate the other end of GRE tunnels directly on the servers on the LAN, while an MX Series router functions as the gateway router between the WAN and LAN. This type of termination of tunnels enables users to build overlay networks within the data center without having to configure end-user VLANs, IP addresses, and other network parameters on the underlying switches. Such a setup simplifies data center network design and provisioning.

Related •  
Documentation

## Configuring Layer 2 Services over GRE Logical Interfaces in Bridge Domains

You can configure Layer 2 Ethernet services over GRE interfaces (*gr-fpc/pic/port* to use GRE encapsulation).

To configure a GRE tunnel interface, associate it in a bridge domain within a virtual-switch instance, and specify the amount of bandwidth reserved for tunnel services traffic:

1. Configure GRE tunnel interface and specify the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.

[edit]

```
user@host# set chassis fpc slot-number pic slot-number tunnel-services bandwidth
(1g | 10g | 20g | 40g)
```

2. Configure the interfaces and their VLAN IDs.

[edit]

```
user@host# set interfaces gr-interface-name unit logical-unit-number family
family-name address address
user@host# set interfaces gr-interface-name unit logical-unit-number family
family-name interface-mode trunk
user@host# set interfaces gr-interface-name unit logical-unit-number family
family-name vlan-id-list vlan-id-list
user@host# set interfaces gr-interface-name unit logical-unit-number tunnel source
source-address
user@host# set interfaces gr-interface-name unit logical-unit-number tunnel destination
destination-address
```

3. Create a virtual switch instance with a bridge domain and associate the GRE logical interfaces.

```
[edit routing-instances]
user@host# set routing-instance-name instance-type virtual-switch
user@host# set routing-instance-name interface interface-name unit
logical-unit-number
user@host# set routing-instance-name bridge-domains bridge-domain-name vlan-id
vlan-id
```

The VLAN IDs configured for the bridge domain must match with the VLAN IDs that you configure for GRE interfaces by using the **vlan-id (all | none | number)** statement or the **vlan-id-list [ vlan-id-numbers ]** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level.

Related •  
Documentation

## Example: Configuring Layer 2 Services Over GRE Logical Interfaces in Bridge Domains

This example illustrates how you can configure GRE logical interfaces in a bridge domain. You can also configure a virtual switch instance associated with a bridge domain and include GRE interfaces in the bridge domain. This type of configuration enables you to specify Layer 2 Ethernet packets to be terminated on GRE tunnels. In a Layer 2 Ethernet over GRE with VPLS environment, an MX Series router supports Layer 2 over GRE tunnels (without the MPLS layer) and terminate these tunnels into a VPLS or a routed VLAN interface (RVI) into a L3VPN. The tunnels serve to cross the cable modem termination system (CMTS) and cable modem CM infrastructure transparently, up to the MX Series router that serves as the gateway. Every GRE tunnel terminates over a VLAN interface, a VPLS instance, or an IRB interface.

- [Requirements on page 48](#)
- [Overview on page 48](#)
- [Configuration on page 49](#)
- [Verification on page 51](#)

### Requirements

This example uses the following hardware and software components:

- An MX Series router
- Junos OS Release 15.1R1 or later running on an MX Series router with MPCs.

### Overview

GRE encapsulates packets into IP packets and redirects them to an intermediate host, where they are de-encapsulated and routed to their final destination. Because the route to the intermediate host appears to the inner datagrams as one hop, Juniper Networks EX Series Ethernet switches can operate as if they have a virtual point-to-point connection with each other. GRE tunnels allow routing protocols like RIP and OSPF to forward data

packets from one switch to another switch across the Internet. In addition, GRE tunnels can encapsulate multicast data streams for transmission over the Internet.

Ethernet frames have all the essentials for networking, such as globally unique source and destination addresses, error control, and so on. Ethernet frames can carry any kind of packet. Networking at Layer 2 is protocol independent (independent of the Layer 3 protocol). If more of the end-to-end transfer of information from a source to a destination can be done in the form of Ethernet frames, more of the benefits of Ethernet can be realized on the network. Networking at Layer 2 can be a powerful adjunct to IP networking, but it is not usually a substitute for IP networking.

Consider a sample network topology in which a GRE tunnel interface is configured with the bandwidth set as 10 gigabits per second for tunnel traffic on each Packet Forwarding Engine. The GRE interface, gr-0/1/10.0, is specified with the source address of 2.2.2.2 and the destination address of 2.2.2.1. Two gigabit Ethernet interfaces, ge-0/1/2.0 and ge-0/1/6.0, are also configured. A virtual switch instance, VS1, is defined and a bridge domain, bd0, is associated with VS1. The bridge domain contains the VLAN ID of 10. The GRE interface is configured as a trunk interface and associated with the bridge domain, bd0. With such a setup, Layer 2 Ethernet services can be terminated over GRE tunnel interfaces in virtual switch instances that contain bridge domains.

## Configuration

To configure a GRE tunnel interface, associate it in a bridge domain within a virtual-switch instance, and specify the amount of bandwidth reserved for tunnel services traffic.

### CLI Quick Configuration

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

```
set chassis fpc 0 pic 1 tunnel-services bandwidth 1g
set chassis network-services enhanced-ip
set interfaces ge-0/1/2 unit 0 family inet address 2.2.2.2/30
set interfaces ge-0/1/6 unit 0 family bridge interface-mode trunk
set interfaces ge-0/1/6 unit 0 family bridge vlan-id-list 1-100
set interfaces gr-0/1/10 unit 0 tunnel source 2.2.2.2
set interfaces gr-0/1/10 unit 0 tunnel destination 2.2.2.1
set interfaces gr-0/1/10 unit 0 family bridge interface-mode trunk
set interfaces gr-0/1/10 unit 0 family bridge vlan-id-list 1-100
set routing-instances VS1 instance-type virtual-switch
set routing-instances VS1 bridge-domains bd0 vlan-id 10
set routing-instances VS1 interface ge-0/1/2.0
set routing-instances VS1 interface ge-0/1/6.0
set routing-instances VS1 interface gr-0/1/10.0
```

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

To configure GRE logical tunnel interfaces for Layer 2 services in bridge domains:

1. Configure GRE tunnel interface and specify the amount of bandwidth to reserve for tunnel traffic on each Packet Forwarding Engine.

```
[edit]
user@host# set chassis fpc 0 pic 1 tunnel-services bandwidth 1g
user@host# set chassis network-services enhanced-ip
```

2. Configure the interfaces and their VLAN IDs.

```
[edit]
user@host# set interfaces ge-0/1/2 unit 0 family inet address 2.2.2.2/30
user@host# set interfaces ge-0/1/6 unit 0 family bridge interface-mode trunk
user@host# set interfaces ge-0/1/6 unit 0 family bridge vlan-id-list 1-100
user@host# set interfaces gr-0/1/10 unit 0 tunnel source 2.2.2.2
user@host# set interfaces gr-0/1/10 unit 0 tunnel destination 2.2.2.1
user@host# set interfaces gr-0/1/10 unit 0 family bridge interface-mode trunk
user@host# set interfaces gr-0/1/10 unit 0 family bridge vlan-id-list 1-100
```

3. Configure the bridge domain in a virtual switch instance and associate the GRE interface with it.

```
[edit]
user@host# set routing-instances VS1 instance-type virtual-switch
user@host# set routing-instances VS1 bridge-domains bd0 vlan-id 10
user@host# set routing-instances VS1 interface ge-0/1/2.0
user@host# set routing-instances VS1 interface ge-0/1/6.0
user@host# set routing-instances VS1 interface gr-0/1/10.0
```

**Results** Display the results of the configuration:

```
user@host> show configuration

chassis {
  fpc 0 {
    pic 1 {
      tunnel-services {
        bandwidth 1g;
      }
    }
  }
  network-services enhanced-ip;
}
interfaces {
  ge-0/1/2 {
    unit 0 {
      family inet {
        address 2.2.2.2/30;
      }
    }
  }
  ge-0/1/6 {
    unit 0 {
      family bridge {
```

```

        interface-mode trunk;
        vlan-id-list 1-100;
    }
}
gr-0/1/10 {
    unit 0 {
        tunnel {
            source 2.2.2.2;
            destination 2.2.2.1;
        }
        family bridge {
            interface-mode trunk;
            vlan-id-list 1-100;
        }
    }
}
}
VS-1 {
    instance-type virtual-switch;
    interface ge-0/1/2.0;
    interface ge-0/1/6.0;
    interface gr-0/1/10.0;
    bridge-domains {
        bd0 {
            vlan-id 10;
        }
    }
}
}

```

## Verification

To confirm that the configuration is working properly, perform these tasks:

- [Verifying the MAC Addresses Learned on GRE Interfaces on page 51](#)
- [Verifying the MAC Address Learning Status on page 52](#)

### Verifying the MAC Addresses Learned on GRE Interfaces

**Purpose** Display the MAC addresses learned on a GRE logical interface.

**Action** From operational mode, use the **show bridge mac-table** command

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

Routing instance : default-switch
Bridging domain : vlan-1, VLAN : 1
MAC          MAC          Logical
address      flags         interface
00:01:01:00:01:f7  D,SE      gr-1/2/10.0
00:03:00:32:01:f7  D,SE      gr-1/2/10.0
00:00:21:11:11:10  DL        ge-1/0/0.0
00:00:21:11:11:11  DL        ge-1/1/0.0

```

Routing instance : default-switch

```

Bridging domain : vlan-2, VLAN : 2
MAC              MAC      Logical
address          flags    interface
00:02:01:33:01:f7 D,SE  gr-1/2/10.1
00:00:21:11:21:10 DL      ge-1/0/0.1
00:00:21:11:21:11 DL      ge-1/1/0.1

```

**Meaning** The output displays MAC addresses learned on GRE logical tunnels.

### Verifying the MAC Address Learning Status

**Purpose** Display the status of MAC address learning properties in the MAC address and MAC flags fields.

**Action** From operational mode, enter the **show vpls mac-table** command.

MAC flags (S -static MAC, D -dynamic MAC, L -locally learned  
SE -Statistics enabled, NM -Non configured MAC, R -Remote PE MAC)

```

Routing instance : vpls_4site:1000
Bridging domain : __vpls_4site:1000__, MAC      Logical
address          flags    interface
00:01:01:00:01:f4 D,SE  ge-4/2/0.1000
00:02:01:33:01:f4 D,SE  lsi.1052004
00:03:00:32:01:f4 D,SE  lsi.1048840
00:04:00:14:01:f4 D,SE  lsi.1052005
00:02:01:33:02:f7 D,SE  gr-1/2/10.10

```

```

user@host> show interfaces gr-2/2/10
Physical interface: gr-2/2/10, Enabled, Physical link is Up
  Interface index: 214, SNMP ifIndex: 690
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
    3.0.0.1:3.0.0.254:47:df:64:0000000000000000 Encapsulation: GRE-NULL
    L2 Routing Instance: vs1, L3 Routing Instance: default
    Copy-tos-to-outer-ip-header: Off
    Gre keepalives configured: Off, Gre keepalives adjacency state: down
    Input packets : 2
    Output packets: 0
    Protocol bridge, MTU: 1476
    Flags: Sendbroadcast-pkt-to-re
    Addresses, Flags: Is-Preferred Is-Primary
    Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255

```

```

user@host> show interfaces gr-2/2/10.0
  Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
    Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
    3.0.0.1:3.0.0.254:47:df:64:0000000000000000 Encapsulation: GRE-NULL
    L2 Routing Instance: vs1, L3 Routing Instance: default
    Copy-tos-to-outer-ip-header: Off
    Gre keepalives configured: Off, Gre keepalives adjacency state: down
    Input packets : 2

```

```
Output packets: 0
Protocol bridge, MTU: 1476
Flags: Sendbcst-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255
```

**Meaning** The output displays the status of MAC address learning properties in the MAC address and MAC flags fields. The output displays the names of the routing instances associated with the GRE interfaces are displayed.



## CHAPTER 7

# Understanding Default PIM Tunnel Configurations

- [Configuring PIM Tunnels on page 55](#)

## Configuring PIM Tunnels

---

PIM tunnels are enabled automatically on routers that have a tunnel PIC and on which you enable PIM sparse mode. You do not need to configure the tunnel interface.

PIM tunnels are unidirectional.

In PIM sparse mode, the first-hop router encapsulates packets destined for the rendezvous point (RP) router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the RP. The RP then de-encapsulates the packets and transmits them through its multicast tree. To perform the encapsulation and de-encapsulation, the first-hop and RP routers must be equipped with Tunnel PICs.

The Junos OS creates two interfaces to handle PIM tunnels:

- **pe**—Encapsulates packets destined for the RP. This interface is present on the first-hop router.
- **pd**—De-encapsulates packets at the RP. This interface is present on the RP.



**NOTE:** The **pe** and **pd** interfaces do not support class-of-service (CoS) configurations.

### Related Documentation

- [Tunnel Services Overview on page 3](#)



## CHAPTER 8

# Facilitating VRF Table Lookup Using Virtual Loopback Tunnel Interfaces

- [Configuring Virtual Loopback Tunnels for VRF Table Lookup on page 57](#)
- [Configuring Tunnel Interfaces for Routing Table Lookup on page 59](#)
- [Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 59](#)
- [Example: Virtual Routing and Forwarding \(VRF\) and Service Configuration on page 60](#)

## Configuring Virtual Loopback Tunnels for VRF Table Lookup

To enable egress filtering, you can either configure filtering based on the IP header, or you can configure a virtual loopback tunnel on routers equipped with a Tunnel PIC. [Table 4 on page 57](#) describes each method.

**Table 4: Methods for Configuring Egress Filtering**

Method	Interface Type	Configuration Guidelines	Comments
Filter traffic based on the IP header	Nonchannelized Point-to-Point Protocol / High Level Data Link Control (PPP/HDLC) core-facing SONET/SDH interfaces	Include the <b>vrf-table-label</b> statement at the <b>[edit routing-instances instance-name]</b> hierarchy level.  For more information, see the <i>Junos OS VPNs Library for Routing Devices</i> .	There is no restriction on customer-edge (CE) router-to-provider edge (PE) router interfaces.
Configure a virtual loopback tunnel on routers equipped with a Tunnel PIC	All interfaces	See the guidelines in this section.	Router must be equipped with a Tunnel PIC.  There is no restriction on the type of core-facing interface used or CE router-to-PE router interface used.  You cannot configure a virtual loopback tunnel and the <b>vrf-table-label</b> statement at the same time.

You can configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels. You might want to enable this functionality so you can do either of the following:

- Forward traffic on a PE router to CE device interface, in a shared medium, where the CE device is a Layer 2 switch without IP capabilities (for example, a metro Ethernet switch).

The first lookup is done based on the VPN label to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to forward packets to the correct end hosts on the shared medium.

- Perform egress filtering at the egress PE router.

The first lookup on the VPN label is done to determine which VRF table to refer to, and the second lookup is done on the IP header to determine how to filter and forward packets. You can enable this functionality by configuring output filters on the VRF interfaces.

To configure a virtual loopback tunnel to facilitate VRF table lookup based on MPLS labels, you specify a virtual loopback tunnel interface name and associate it with a routing instance that belongs to a particular routing table. The packet loops back through the virtual loopback tunnel for route lookup. To specify a virtual loopback tunnel interface name, you configure the virtual loopback tunnel interface at the **[edit interfaces]** hierarchy level and include the **family inet** and **family mpls** statements:

```
vt-fpc/pic/port {
  unit 0 {
    family inet;
    family mpls;
  }
  unit 1 {
    family inet;
  }
}
```

To associate the virtual loopback tunnel with a routing instance, include the virtual loopback tunnel interface name at the **[edit routing-instances]** hierarchy level:

```
interface vt-fpc/pic/port;
```



**NOTE:** On virtual loopback tunnel interfaces, none of the logical interface statements except the **family** statement is supported. Note that you can configure only **inet** and **mpls** families, and you cannot configure IPv4 or IPv6 addresses on virtual loopback tunnel interfaces. Also, virtual loopback tunnel interfaces do not support class-of-service (CoS) configurations.

#### Related Documentation

- [Tunnel Services Overview on page 3](#)
- [Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup on page 59](#)

## Configuring Tunnel Interfaces for Routing Table Lookup

To configure tunnel interfaces to facilitate routing table lookups for VPNs, you specify a tunnel's endpoint IP addresses and associate them with a routing instance that belongs to a particular routing table. This enables the Junos OS to search in the appropriate routing table for the route prefix, because the same prefix can appear in multiple routing tables. To configure the destination VPN, include the **routing-instance** statement:

```
routing-instance {
  destination routing-instance-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *gr-fpc/pic/port* unit *logical-unit-number* tunnel]
- [edit logical-systems *logical-system-name* interfaces *gr-fpc/pic/port* unit *logical-unit-number* tunnel]

This configuration indicates that the tunnel's destination address is in routing instance **routing-instance-name**. By default, the tunnel route prefixes are assumed to be in the default Internet routing table **inet.0**.



**NOTE:** If you configure a virtual loopback tunnel interface and the **vrf-table-label** statement on the same routing instance, the **vrf-table-label** statement takes precedence over the virtual loopback tunnel interface. For more information, see “Configuring Virtual Loopback Tunnels for VRF Table Lookup” on page 57.

For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*.

### Related Documentation

- [Tunnel Services Overview on page 3](#)
- [destination \(Routing Instance\) on page 89](#)

## Example: Configuring a Virtual Loopback Tunnel for VRF Table Lookup

Configure a virtual loopback tunnel for VRF table lookup:

```
[edit routing-instances]
routing-instance-1 {
  instance-type vrf;
  interface vt-1/0/0.0;
  interface so-0/2/2.0;
  route-distinguisher 2:3;
  vrf-import VPN-A-import;
  vrf-export VPN-A-export;
  routing-options {
    static {
      route 10.0.0.0/8 next-hop so-0/2/2.0;
    }
  }
}
```

```
    }  
  }  
  routing-instance-2 {  
    instance-type vrf;  
    interface vt-1/0/0.1;  
    interface so-0/3/2.0;  
    route-distinguisher 4:5;  
    vrf-import VPN-B-import;  
    vrf-export VPN-B-export;  
    routing-options {  
      static {  
        route 10.0.0.0/8 next-hop so-0/3/2.0;  
      }  
    }  
  }  
}  
[edit interfaces]  
vt-1/0/0 {  
  unit 0 {  
    family inet;  
    family mpls;  
  }  
  unit 1 {  
    family inet;  
  }  
}
```

- Related Documentation**
- [Tunnel Services Overview on page 3](#)
  - [Configuring Virtual Loopback Tunnels for VRF Table Lookup on page 57](#)

---

## Example: Virtual Routing and Forwarding (VRF) and Service Configuration

The following example combines virtual routing and forwarding (VRF) and services configuration:

```
[edit policy-options]  
policy-statement test-policy {  
  term t1 {  
    then reject;  
  }  
}  
[edit routing-instances]  
test {  
  interface ge-0/2/0.0;  
  interface sp-1/3/0.20;  
  instance-type vrf;  
  route-distinguisher 10.58.255.1:37;  
  vrf-import test-policy;  
  vrf-export test-policy;  
  routing-options {  
    static {  
      route 0.0.0.0/0 next-table inet.0;  
    }  
  }  
}
```

```
[edit interfaces]
ge-0/2/0 {
  unit 0 {
    family inet {
      service {
        input service-set nat-me;
        output service-set nat-me;
      }
    }
  }
}
sp-1/3/0 {
  unit 0 {
    family inet;
  }
  unit 20 {
    family inet;
    service-domain inside;
  }
  unit 21 {
    family inet;
    service-domain outside;
  }
}
[edit services]
stateful-firewall {
  rule allow-any-input {
    match-direction input;
    term t1 {
      then accept;
    }
  }
}
nat {
  pool hide-pool {
    address 10.58.16.100;
    port automatic;
  }
  rule hide-all-input {
    match-direction input;
    term t1 {
      then {
        translated {
          source-pool hide-pool;
          translation-type source napt-44;
        }
      }
    }
  }
}
service-set nat-me {
  stateful-firewall-rules allow-any-input;
  nat-rules hide-all-input;
  interface-service {
    service-interface sp-1/3/0.20;
  }
}
```

}

## CHAPTER 9

# Enabling a VPN to Travel Through a Non-MPLS Network Using Dynamic Tunnels

- [Configuring Dynamic Tunnels on page 63](#)

## Configuring Dynamic Tunnels

---

A VPN that travels through a non-MPLS network requires a GRE tunnel. This tunnel can be either a static tunnel or a dynamic tunnel. A static tunnel is configured manually between two PE routers. A dynamic tunnel is configured using BGP route resolution.

When a router receives a VPN route that resolves over a BGP next hop that does not have an MPLS path, a GRE tunnel can be created dynamically, allowing the VPN traffic to be forwarded to that route. Only GRE IPv4 tunnels are supported.

To configure a dynamic tunnel between two PE routers, include the **dynamic-tunnels** statement:

```
dynamic-tunnels tunnel-name {  
    destination-networks prefix;  
    source-address address;  
}
```

You can configure this statement at the following hierarchy levels:

- [edit routing-options]
- [edit routing-instances *routing-instance-name* routing-options]
- [edit logical-systems *logical-system-name* routing-options]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options]

For more information about configuring routing options or BGP, see the *Junos OS Routing Protocols Library for Routing Devices*. For more information about VPNs, see the *Junos OS VPNs Library for Routing Devices*.

### Related Documentation

- [Tunnel Services Overview on page 3](#)

- [dynamic-tunnels on page 92](#)

## PART 2

# Encryption Services

- [Overview on page 67](#)
- [Sending Encrypted Traffic Through Tunnels on page 69](#)
- [Configuring Redundancy in Case of Service Failure on page 77](#)



## CHAPTER 10

# Overview

- [Encryption Overview on page 67](#)
- [Configuring an ES Tunnel Interface for a Layer 3 VPN on page 67](#)

## Encryption Overview

---

The IP Security (IPsec) architecture provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates.

IPsec defines a security association (SA) and key management framework that can be used with any network layer protocol. The SA specifies what protection policy to apply to traffic between two IP-layer entities. For more information, see the *Junos OS Administration Library for Routing Devices*. The standards are defined in the following RFCs:

- RFC 2401, *Security Architecture for the Internet Protocol*
- RFC 2406, *IP Encapsulating Security Payload (ESP)*

### Related Documentation

- [Configuring Encryption Interfaces on page 69](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 71](#)
- [Configuring ES PIC Redundancy on page 77](#)
- [Configuring IPsec Tunnel Redundancy on page 78](#)

## Configuring an ES Tunnel Interface for a Layer 3 VPN

---

To configure an ES tunnel interface for a Layer 3 VPN, you need to configure an ES tunnel interface on the provider edge (PE) router and on the customer edge (CE) router. You also need to configure IPsec on the PE and CE routers. For more information about configuring an ES tunnel for a Layer 3 VPN, see the *Junos OS VPNs Library for Routing Devices*.

**Related  
Documentation**

- [Encryption Overview on page 67](#)
- [Configuring Encryption Interfaces on page 69](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 71](#)
- [Configuring ES PIC Redundancy on page 77](#)
- [Configuring IPsec Tunnel Redundancy on page 78](#)

# Sending Encrypted Traffic Through Tunnels

- [Configuring Encryption Interfaces on page 69](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 71](#)

## Configuring Encryption Interfaces

---

When you configure the encryption interface, you associate the configured SA with a logical interface. This configuration defines the tunnel, including the logical unit, tunnel addresses, maximum transmission unit (MTU), optional interface addresses, and the name of the IPsec SA to apply to traffic. To configure an encryption interface, include the following statements at the `[edit interfaces es-fpc/pic/port unit logical-unit-number]` hierarchy level:

```
family inet {
  ipsec-sa ipsec-sa; # name of security association to apply to packet
  address address; # local interface address inside local VPN
  destination address; # destination address inside remote VPN
}
tunnel {
  source source-address;
  destination destination-address;
}
```

The addresses configured as the tunnel source and destination are the addresses in the outer IP header of the tunnel.



**NOTE:** You must configure the tunnel source address locally on the router, and the tunnel destination address must be a valid address for the security gateway terminating the tunnel.

The ES Physical Interface Card (PIC) is supported on M Series and T Series routers.

The SA must be a valid tunnel-mode SA. The interface address and destination address listed are optional. The destination address allows the user to configure a static route to

encrypt traffic. If a static route uses that destination address as the next hop, traffic is forwarded through the portion of the tunnel in which encryption occurs.

## Specifying the Security Association Name for Encryption Interfaces

The security association is the set of properties that defines the protocols for encrypting Internet traffic. To configure encryption interfaces, you specify the SA name associated with the interface by including the **ipsec-sa** statement at the **[edit interfaces es-fpc/pic/port unit logical-unit-number family inet]** hierarchy level:

```
ipsec-sa sa-name;
```

For information about configuring the security association, see [“Configuring Filters for Traffic Transiting the ES PIC” on page 71](#).

## Configuring the MTU for Encryption Interfaces

The protocol MTU value for encryption interfaces must always be less than the default interface MTU value of 3900 bytes; the configuration fails to commit if you select a greater value. To set the MTU value, include the **mtu** statement at the **[edit interfaces interface-name unit logical-unit-number family inet]** hierarchy level:

```
mtu bytes;
```

For more information, see the *Junos OS Network Interfaces Library for Routing Devices*.

## Example: Configuring an Encryption Interface

Configure an IPsec tunnel as a logical interface on the ES PIC. The logical interface specifies the tunnel through which the encrypted traffic travels. The **ipsec-sa** statement associates the security profile with the interface.

```
[edit interfaces]
es-0/0/0 {
  unit 0 {
    tunnel {
      source 10.5.5.5; # tunnel source address
      destination 10.6.6.6; # tunnel destination address
    }
    family inet {
      ipsec-sa manual-sa1; # name of security association to apply to packet
      mtu 3800;
      address 10.1.1.8/32 { # local interface address inside local VPN
        destination 10.2.2.254; # destination address inside remote VPN
      }
    }
  }
}
```

### Related Documentation

- [Encryption Overview on page 67](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 71](#)
- [Configuring ES PIC Redundancy on page 77](#)
- [Configuring IPsec Tunnel Redundancy on page 78](#)

## Configuring Filters for Traffic Transiting the ES PIC

This section contains the following topics:

- [Traffic Overview on page 71](#)
- [Configuring the Security Association on page 72](#)
- [Configuring an Outbound Traffic Filter on page 73](#)
- [Applying the Outbound Traffic Filter on page 74](#)
- [Configuring an Inbound Traffic Filter on page 74](#)
- [Applying the Inbound Traffic Filter to the Encryption Interface on page 75](#)

### Traffic Overview

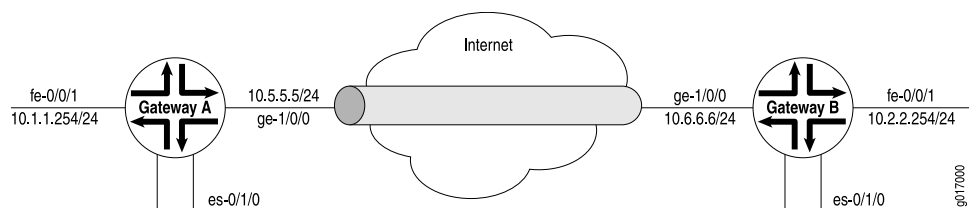
Traffic configuration defines the traffic that must flow through the tunnel. You configure outbound and inbound firewall filters, which identify and direct traffic to be encrypted and confirm that decrypted traffic parameters match those defined for the given tunnel. The outbound filter is applied to the LAN or WAN interface for the incoming traffic you want to encrypt. The inbound filter is applied to the ES PIC to check the policy for traffic coming in from the remote host. Because of the complexity of configuring a router to forward packets, no automatic checking is done to ensure that the configuration is correct.



**NOTE:** The valid firewall filters statements for IPsec are **destination-port**, **source-port**, **protocol**, **destination-address**, and **source-address**.

In [Figure 3 on page 71](#), Gateway A protects the network 10.1.1.0/24, and Gateway B protects the network 10.2.2.0/24. The gateways are connected by an IPsec tunnel. For more information about firewalls, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

**Figure 3: Example: IPsec Tunnel Connecting Security Gateways**



The SA and ES interface for security Gateway A are configured as follows:

```
[edit security ipsec]
security-association manual-sa1 {
  manual {
    direction bidirectional {
      protocol esp;
      spi 2312;
      authentication {
        algorithm hmac-md5-96;
```

```
        key ascii-text 1234123412341234;
    }
    encryption {
        algorithm 3des-cbc;
        key ascii-text 123456789009876543211234;
    }
}
}
[edit interfaces es-0/1/0]
unit 0 {
    tunnel {
        source 10.5.5.5;
        destination 10.6.6.6;
    }
    family inet {
        ipsec-sa manual-sa1;
        address 10.1.1.8/32 {
            destination 10.2.2.254;
        }
    }
}
```

## Configuring the Security Association

To configure the SA, include the **security-association** statement at the **[edit security]** hierarchy level:

```
security-association name {
    mode (tunnel | transport);
    manual {
        direction (inbound | outbound | bi-directional) {
            auxiliary-spi auxiliary-spi-value;
            spi spi-value;
            protocol (ah | esp | bundle);
            authentication {
                algorithm (hmac-md5-96 | hmac-sha1-96);
                key (ascii-text key | hexadecimal key);
            }
            encryption {
                algorithm (des-cbc | 3des-cbc);
                key (ascii-text key | hexadecimal key);
            }
        }
        dynamic {
            replay-window-size (32 | 64);
            ipsec-policy policy-name;
        }
    }
}
```

For more information about configuring an SA, see the *Junos OS Administration Library for Routing Devices*. For information about applying the SA to an interface, see [“Specifying the Security Association Name for Encryption Interfaces” on page 70](#).

## Configuring an Outbound Traffic Filter

To configure the outbound traffic filter, include the **filter** statement at the **[edit firewall]** hierarchy level:

```
filter filter-name {
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
```

For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Example: Configuring an Outbound Traffic Filter

Firewall filters for outbound traffic direct the traffic through the desired IPsec tunnel and ensure that the tunneled traffic goes out the appropriate interface (see [Figure 3 on page 71](#)). Here, an outbound firewall filter is created on security Gateway A; it identifies the traffic to be encrypted and adds it to the input side of the interface that carries the internal virtual private network (VPN) traffic:

```
[edit firewall]
filter ipsec-encrypt-policy-filter {
  term term1 {
    from {
      source-address { # local network
        10.1.1.0/24;
      }
      destination-address { # remote network
        10.2.2.0/24;
      }
    }
  }
  then ipsec-sa manual-sa1; # apply SA name to packet
  term default {
    then accept;
  }
}
```



**NOTE:** The source address, port, and protocol on the outbound traffic filter must match the destination address, port, and protocol on the inbound traffic filter. The destination address, port, and protocol on the outbound traffic filter must match the source address, port, and protocol on the inbound traffic filter.

## Applying the Outbound Traffic Filter

After you have configured the outbound firewall filter, you apply it by including the **filter** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level:

```
filter {  
    input filter-name;  
}
```

---

### Example: Applying the Outbound Traffic Filter

Apply the outbound traffic filter. The outbound filter is applied on the Fast Ethernet interface at the **[edit interfaces *fe-0/0/1* unit 0 family inet]** hierarchy level. Any packet matching the IPsec action term (**term 1**) on the input filter (**ipsec-encrypt-policy-filter**), configured on the Fast Ethernet interface, is directed to the ES PIC interface at the **[edit interfaces *es-0/1/0* unit 0 family inet]** hierarchy level. So, if a packet arrives from the source address **10.1.1.0/24** and goes to the destination address **10.2.2.0/24**, the Packet Forwarding Engine directs the packet to the ES PIC interface, which is configured with the **manual-sa1** SA. The ES PIC receives the packet, applies the **manual-sa1** SA, and sends the packet through the tunnel.

The router must have a route to the tunnel end point; add a static route if necessary.

```
[edit interfaces]  
fe-0/0/1 {  
    unit 0 {  
        family inet {  
            filter {  
                input ipsec-encrypt-policy-filter;  
            }  
            address 10.1.1.254/24;  
        }  
    }  
}
```

## Configuring an Inbound Traffic Filter

To configure an inbound traffic filter, include the **filter** statement at the **[edit firewall]** hierarchy level:

```
filter filter-name {  
    term term-name {  
        from {  
            match-conditions;  
        }  
        then {  
            action;  
            action-modifiers;  
        }  
    }  
}
```

For more information, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Example: Configuring an Inbound Traffic Filter

Configure an inbound firewall filter. This filter performs the final IPsec policy check and is created on security gateway A. The policy check ensures that only packets that match the traffic configured for this tunnel are accepted.

```
[edit firewall]
filter ipsec-decrypt-policy-filter {
  term term1 { # perform policy check
    from {
      source-address { # remote network
        10.2.2.0/24;
      }
      destination-address { # local network
        10.1.1.0/24;
      }
    }
  }
  then accept;
```

### Applying the Inbound Traffic Filter to the Encryption Interface

After you create the inbound firewall filter, you can apply it to the ES PIC. To apply the filter to the ES PIC, include the **filter** statement at the **[edit interfaces es-fpc/pic/port unit logical-unit-number family inet filter]** hierarchy level:

```
filter {
  input filter;
}
```

The input filter is the name of the filter applied to received traffic. For a configuration example, see “[Example: Configuring an Inbound Traffic Filter](#)” on page 75. For more information about firewall filters, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*.

### Example: Applying the Inbound Traffic Filter to the Encryption Interface

Apply the inbound firewall filter (**ipsec-decrypt-policy-filter**) to the decrypted packet to perform the final policy check. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and decrypts the incoming packet.

The Packet Forwarding Engine directs IPsec packets to the ES PIC. It uses the packet's security parameter index (SPI), protocol, and destination address to look up the SA configured on one of the ES interfaces. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and is used to decrypt the incoming packet. When the packets are processed (decrypted, authenticated, or both), the input firewall filter (**ipsec-decrypt-policy-filter**) is applied on the decrypted packet to perform the final policy check. **term1** defines the decrypted (and verified) traffic and performs the required policy check. For information about **term1**, see “[Example: Configuring an Inbound Traffic Filter](#)” on page 75.



**NOTE:** The inbound traffic filter is applied after the ES PIC has processed the packet, so the decrypted traffic is defined as any traffic that the remote gateway is encrypting and sending to this router. IKE uses this filter to determine the policy required for a tunnel. This policy is used during the negotiation with the remote gateway to find the matching SA configuration.

```
[edit interfaces]
es-1/2/0 {
  unit 0 {
    tunnel {
      source 10.5.5.5; # tunnel source address
      destination 10.6.6.6; # tunnel destination address
    }
    family inet {
      filter {
        input ipsec-decrypt-policy-filter;
      }
      ipsec-sa manual-sa1; # SA name applied to packet
      address 10.1.1.8/32 { # local interface address inside local VPN
        destination 10.2.2.254; # destination address inside remote VPN
      }
    }
  }
}
```

**Related  
Documentation**

- [Encryption Overview on page 67](#)
- [Configuring Encryption Interfaces on page 69](#)
- [Configuring ES PIC Redundancy on page 77](#)
- [Configuring IPsec Tunnel Redundancy on page 78](#)

# Configuring Redundancy in Case of Service Failure

- [Configuring ES PIC Redundancy on page 77](#)
- [Configuring IPsec Tunnel Redundancy on page 78](#)

## Configuring ES PIC Redundancy

You can configure ES PIC redundancy on M Series and T Series routers that have multiple ES PICs. With ES PIC redundancy, one ES PIC is active and another ES PIC is on standby. When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and SAs, and acts as the new next hop for IPsec traffic. Reestablishment of tunnels on the backup ES PIC does not require new Internet Key Exchange (IKE) negotiations. If the primary ES PIC comes online, it remains in standby and does not preempt the backup. To determine which PIC is currently active, use the **show ipsec redundancy** command.



**NOTE:** ES PIC redundancy is supported on M Series and T Series routers.

To configure an ES PIC as the backup, include the **backup-interface** statement at the **[edit interfaces fpc/pic/port es-options]** hierarchy level:

```
backup-interface es-fpc/pic/port;
```

## Example: Configuring ES PIC Redundancy

After you create the inbound firewall filter, apply it to the master ES PIC. Here, the inbound firewall filter (**ipsec-decrypt-policy-filter**) is applied on the decrypted packet to perform the final policy check. The IPsec **manual-sa1** SA is referenced at the **[edit interfaces es-1/2/0 unit 0 family inet]** hierarchy level and decrypts the incoming packet. This example does not show SA and filter configuration. For information about SA and filter configuration, see the *Junos OS Administration Library for Routing Devices*, the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices*, and “Example: Configuring an Inbound Traffic Filter” on page 75.

```
[edit interfaces]
es-1/2/0 {
  es-options {
```

```

    backup-interface es-1/0/0;
  }
  unit 0 {
    tunnel {
      source 10.5.5.5;
      destination 10.6.6.6;
    }
    family inet {
      ipsec-sa manual-sa1;
      filter {
        input ipsec-decrypt-policy-filter;
      }
      address 10.1.1.8/32 {
        destination 10.2.2.254;
      }
    }
  }
}

```

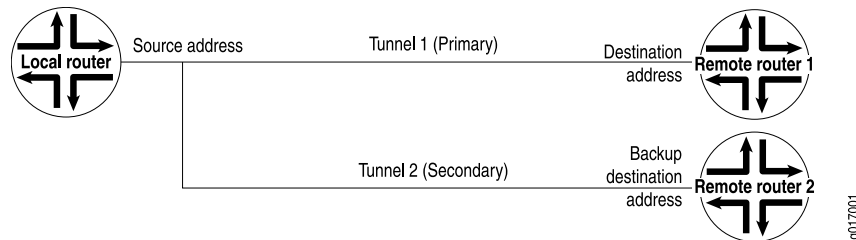
#### Related Documentation

- [Encryption Overview on page 67](#)
- [Configuring Encryption Interfaces on page 69](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 71](#)
- [Configuring IPsec Tunnel Redundancy on page 78](#)

## Configuring IPsec Tunnel Redundancy

You can configure IPsec tunnel redundancy by specifying a backup destination address. The local router sends keepalives to determine the remote site's reachability. When the peer is no longer reachable, a new tunnel is established. For up to 60 seconds during failover, traffic is dropped without notification being sent. [Figure 4 on page 78](#) shows IPsec primary and backup tunnels.

Figure 4: IPsec Tunnel Redundancy



To configure IPsec tunnel redundancy, include the **backup-destination** statement at the **[edit interfaces unit *logical-unit-number* tunnel]** hierarchy level:

```

backup-destination address;
destination address;
source address;

```



NOTE: Tunnel redundancy is supported on M Series and T Series routers.

The primary and backup destinations must be on different routers.

The tunnels must be distinct from each other and policies must match.

For more information about tunnels, see *Tunnel Properties*.

**Related  
Documentation**

- [Encryption Overview on page 67](#)
- [Configuring Encryption Interfaces on page 69](#)
- [Configuring Filters for Traffic Transiting the ES PIC on page 71](#)
- [Configuring ES PIC Redundancy on page 77](#)



## PART 3

# Configuration Statements and Operational Commands

- Configuration Statements on page 83
- Operational Commands on page 113



## CHAPTER 13

# Configuration Statements

- [address \(Interfaces\) on page 84](#)
- [allow-fragmentation on page 85](#)
- [apply-groups-except on page 85](#)
- [backup-destination on page 86](#)
- [backup-interface on page 86](#)
- [copy-tos-to-outer-ip-header on page 87](#)
- [core-facing on page 87](#)
- [destination \(Interfaces\) on page 88](#)
- [destination \(Routing Instance\) on page 89](#)
- [destination \(Tunnel Remote End\) on page 89](#)
- [destination-networks on page 90](#)
- [do-not-fragment on page 91](#)
- [dynamic-tunnels on page 92](#)
- [es-options on page 93](#)
- [family on page 94](#)
- [family bridge on page 95](#)
- [family bridge \(GRE Interfaces\) on page 96](#)
- [filter on page 97](#)
- [hold-time \(OAM\) on page 97](#)
- [interfaces on page 98](#)
- [ipsec-sa on page 98](#)
- [keepalive-time on page 99](#)
- [key on page 100](#)
- [multicast-only on page 100](#)
- [peer-unit on page 101](#)
- [peer-certificate-type on page 101](#)
- [reassemble-packets on page 102](#)
- [redundancy-group \(Interfaces\) on page 102](#)

- [redundancy-group \(Logical Tunnels\) on page 103](#)
- [routing-instance on page 104](#)
- [routing-instances on page 104](#)
- [routing-options on page 105](#)
- [source on page 105](#)
- [source on page 106](#)
- [source-address on page 106](#)
- [ttl on page 107](#)
- [tunnel on page 108](#)
- [tunnel on page 109](#)
- [unit \(Interfaces\) on page 110](#)
- [unit \(Interfaces\) on page 111](#)

---

## address (Interfaces)

---

<b>Syntax</b>	<code>address <i>address</i> {     <i>destination address</i>; }</code>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">family</a> <i>family</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the interface address.
<b>Options</b>	<i>address</i> —Address of the interface.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li></ul>

## allow-fragmentation

<b>Syntax</b>	allow-fragmentation;
<b>Hierarchy Level</b>	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.
<b>Description</b>	Enable fragmentation of generic routing encapsulation (GRE) encapsulated packets regardless of maximum transmission unit (MTU) value.
<b>Default</b>	By default, the GRE-encapsulated packets are dropped if the packet size exceeds the MTU setting of the egress interface.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">reassemble-packets on page 102</a></li> <li>• <a href="#">Configuring Packet Reassembly on page 21</a></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> </ul>

## apply-groups-except

<b>Syntax</b>	apply-groups-except <i>values</i> ;
<b>Hierarchy Level</b>	[edit interfaces]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Disable inheritance of a configuration group.
<b>Options</b>	<i>value</i> —.
<b>Required Privilege Level</b>	configure—To enter configuration mode, but other required privilege levels depend on where the statement is located in the configuration hierarchy.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>groups</i></li> <li>• <i>Disabling Inheritance of a Junos OS Configuration Group</i></li> <li>• <a href="#">Tunnel Services Overview on page 3</a></li> <li>• <a href="#">Configuring Tunnel Interfaces on MX Series Routers on page 6</a></li> </ul>

## backup-destination

---

<b>Syntax</b>	<code>backup-destination <i>destination-address</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">tunnel</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">tunnel</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.
<b>Description</b>	For tunnel interfaces, specify the remote address of the backup tunnel.
<b>Options</b>	<i>destination-address</i> —Address of the remote side of the connection.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">destination (Interfaces) on page 88</a></li><li>• <a href="#">destination (Tunnel Remote End) on page 89</a></li><li>• <a href="#">Configuring IPsec Tunnel Redundancy on page 78</a></li></ul>

## backup-interface

---

<b>Syntax</b>	<code>backup-interface <i>interface-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">es-options</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure a backup ES Physical Interface Card (PIC). When the primary ES PIC has a servicing failure, the backup becomes active, inherits all the tunnels and security associations (SAs), and acts as the new next hop for IPsec traffic.
<b>Options</b>	<i>interface-name</i> —Name of ES interface to serve as the backup.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring ES PIC Redundancy on page 77</a></li></ul>

## copy-tos-to-outer-ip-header

---

<b>Syntax</b>	copy-tos-to-outer-ip-header;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 8.2.
<b>Description</b>	For GRE tunnel interfaces only, enable the inner IP header's ToS bits to be copied to the outer IP packet header.
<b>Default</b>	If you omit this statement, the ToS bits in the outer IP header are set to 0.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring a GRE Tunnel to Copy ToS Bits to the Outer IP Header on page 21</a></li> </ul>

## core-facing

---

<b>Syntax</b>	core-facing;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> ],
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Specifies that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, the interface is inserted into the core-facing default mesh group where traffic from pseudowires that belong to the default mesh group is not forwarded on the core-facing link.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Junos OS Broadband Subscriber Management and Services Library</i></li> <li>• <a href="#">Tunnel Services Overview on page 3</a></li> <li>• <a href="#">Configuring Tunnel Interfaces on MX Series Routers on page 6</a></li> </ul>

## destination (Interfaces)

---

<b>Syntax</b>	<code>destination address;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit logical-unit-number</a> tunnel], [edit interfaces <i>interface-name</i> <a href="#">unit logical-unit-number</a> family inet address <i>address</i> ], [edit interfaces <i>interface-name</i> <a href="#">unit logical-unit-number</a> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit logical-unit-number</a> family inet address <i>address</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>For CoS on ATM interfaces, specify the remote address of the connection.</p> <p>For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.</p> <p>For tunnel and encryption interfaces, specify the remote address of the tunnel.</p>
<b>Options</b>	<b><i>address</i></b> —Address of the remote side of the connection.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Linear RED Profiles on ATM Interfaces</a></li><li>• <a href="#">Multilink and Link Services Logical Interface Configuration Overview</a></li><li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li><li>• <a href="#">Configuring Traffic Sampling</a></li><li>• <a href="#">Configuring Flow Monitoring</a></li><li>• <a href="#">Configuring Unicast Tunnels on page 17</a></li></ul>

## destination (Routing Instance)

<b>Syntax</b>	<code>destination <i>routing-instance-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> <i>routing-instance</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the destination routing instance that points to the routing table containing the tunnel destination address.
<b>Default</b>	The default Internet routing table <b>inet.0</b> .
<b>Options</b>	<i>routing-instance-name</i> —Name of the destination routing instance.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Tunnel Interfaces for Routing Table Lookup on page 59</a></li> </ul>

## destination (Tunnel Remote End)

<b>Syntax</b>	<code>destination <i>destination-address</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	For tunnel interfaces, specify the remote address of the tunnel.
<b>Options</b>	<i>destination-address</i> —Address of the remote side of the connection.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Unicast Tunnels on page 17</a></li> <li>• <a href="#">Configuring Traffic Sampling</a></li> <li>• <a href="#">Configuring Flow Monitoring</a></li> </ul>

## destination-networks

---

<b>Syntax</b>	<code>destination-networks <i>prefix</i>;</code>
<b>Hierarchy Level</b>	<code>[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>],</code> <code>[edit logical-systems <i>logical-system-name</i> <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit routing-instances <i>routing-instance-name</i> <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>],</code> <code>[edit <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i>],</code> <code>[edit <a href="#">routing-options</a> dynamic-tunnels <i>tunnel-name</i> rsvp-te <i>entry</i>]</code>
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Specify the IPv4 prefix range for the destination network. Only tunnels within the specified IPv4 prefix range can be created.
<b>Options</b>	<i>prefix</i> —Destination prefix of the network.
<b>Required Privilege Level</b>	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring GRE Tunnels for Layer 3 VPNs</a></li><li>• <a href="#">Configuring Dynamic Tunnels on page 63</a></li><li>• <a href="#">Configuring RSVP Automatic Mesh</a></li></ul>

## do-not-fragment

---

<b>Syntax</b>	do-not-fragment;
<b>Hierarchy Level</b>	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.
<b>Description</b>	Disable fragmentation of GRE encapsulated packets. Set the do-not-fragment (DF) bit on the packets entering the GRE tunnel so that they do not get fragmented anywhere in the path.
<b>Default</b>	By default, fragmentation is disabled.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">reassemble-packets on page 102</a></li> <li>• <a href="#">Configuring Packet Reassembly on page 21</a></li> <li>• <i>Junos OS Services Interfaces Library for Routing Devices</i></li> </ul>

## dynamic-tunnels

---

Syntax	<pre>dynamic-tunnels <i>tunnel-name</i> {   <i>destination-networks</i> <i>prefix</i>;   gre;   rsvp-te <i>entry-name</i> {     <i>destination-networks</i> <i>network-prefix</i>;     label-switched-path-template (Multicast) {       default-template;       <i>template-name</i>;     }   }   source-address <i>address</i>; }</pre>
Hierarchy Level	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> <i>routing-options</i> ], [edit logical-systems <i>logical-system-name</i> <i>routing-options</i> ], [edit routing-instances <i>routing-instance-name</i> <i>routing-options</i> ], [edit <i>routing-options</i> ]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Description	Configure a dynamic tunnel between two PE routers.
Options	<i>tunnel-name</i> —Name of the dynamic tunnel.  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	<ul style="list-style-type: none"><li>Example: Configuring a Two-Tiered Virtualized Data Center for Large Enterprise Networks</li></ul>

---

## es-options

---

<b>Syntax</b>	<pre>es-options {     backup-interface <i>interface-name</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>On ES interfaces, configure ES interface-specific interface properties.</p> <p>The <b>backup-interface</b> statement is explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring ES PIC Redundancy on page 77</a></li></ul>

## family

---

<b>Syntax</b>	<code>family <i>family</i> {     <b>ipsec-sa</b> <i>sa-name</i>; }</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure protocol family information for the logical interface.
<b>Options</b>	<p><b>family</b>—Protocol family:</p> <ul style="list-style-type: none"><li>• <b>ccc</b>—Circuit cross-connect protocol suite</li><li>• <b>inet</b>—IP version 4 suite</li><li>• <b>inet6</b>—IP version 6 suite</li><li>• <b>iso</b>—Open Systems Interconnection (OSI) International Organization for Standardization (ISO) protocol suite</li><li>• <b>mlfr-end-to-end</b>—Multilink Frame Relay FRF.15</li><li>• <b>mlfr-uni-nni</b>—Multilink Frame Relay FRF.16</li><li>• <b>multilink-ppp</b>—Multilink Point-to-Point Protocol</li><li>• <b>mpls</b>—MPLS</li><li>• <b>tcc</b>—Translational cross-connect protocol suite</li><li>• <b>tnp</b>—Trivial Network Protocol</li><li>• <b>vpls</b>—Virtual private LAN service</li></ul> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li></ul>

---

## family bridge

---

<b>Syntax</b>	<pre>family bridge {     apply-groups <i>value</i>;     apply-groups-except <i>value</i>;     core-facing;     interface-mode <i>access</i>   <i>trunk</i>;     vlan-id <i>vlan-id</i> ;     vlan-id-list <i>vlan-id-range</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces]
<b>Release Information</b>	Statement introduced in Junos OS Release 15.1.
<b>Description</b>	Family bridge is used when you want a port that has more than one logical unit, each with the same or different encapsulations. Bridge domains are associated with gre interface with the corresponding BD VLAN.
<b>Options</b>	The remaining statements are explained separately.
<b>Required Privilege Level</b>	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tunnel Services Overview on page 3</a></li><li>• <a href="#">Configuring Tunnel Interfaces on MX Series Routers on page 6</a></li></ul>

## family bridge (GRE Interfaces)

---

<b>Syntax</b>	<pre>family bridge {     interface-mode (access   trunk);     core-facing;     vlan-id <i>number</i>;     vlan-id-list <i>vlan-list</i>; }</pre>
<b>Hierarchy Level</b>	[edit interfaces gr-fpc/pic/port.unit]
<b>Release Information</b>	Statement at the [edit interfaces gr-fpc/pic/port.unit] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
<b>Description</b>	Configure the bridge domain family on GRE interfaces. To enable Layer 2 Ethernet packets to be terminated on GRE tunnels, you must configure the bridge domain protocol family on the <b>gr-</b> interfaces and associate the <b>gr-</b> interfaces with the bridge domain. You must configure the GRE interfaces as core-facing interfaces, and they must be access or trunk interfaces. To configure the bridge domain family on <b>gr-</b> interfaces, include the <b>family bridge</b> statement at the [edit interfaces gr-fpc/pic/port unit <i>logical-unit-number</i> ] hierarchy level.
<b>Options</b>	<p><b>interface-mode</b>—Specify the type of VLAN tagging on packets that the interface accepts.</p> <p><b>access</b>—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the <b>vlan-id</b> statement.</p> <p><b>trunk</b>—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the <b>vlan-id</b> or <b>vlan-id-list</b> statement.</p> <p><b>core-facing</b>—Specify that the VLAN is physically connected to a core-facing ISP router and ensures that the network does not improperly treat the interface as a client interface. When specified, the interface is inserted into the core-facing default mesh group where traffic from pseudowires that belong to the default mesh group is not forwarded on the core-facing link.</p> <p><b>vlan-id <i>number</i></b>—Individual VLAN IDs separated by a space.</p> <p><b>vlan-id-list <i>vlan-list</i></b>—Starting VLAN ID and ending VLAN ID in an inclusive range. Separate the starting VLAN ID and ending VLAN ID with a hyphen.</p>
<b>Required Privilege Level</b>	<p><b>interface</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>VLANs Within a Bridge Domain or VPLS Instance</i></li></ul>

## filter

---

<b>Syntax</b>	filter { input <i>filter-name</i> ; output <i>filter-name</i> ; }
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Define the filters to be applied on an interface.
<b>Options</b>	<p><b>input <i>filter-name</i></b>—Identifier for the input filter.</p> <p><b>output <i>filter-name</i></b>—Identifier for the output filter.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Filters for Traffic Transiting the ES PIC on page 71</a></li> </ul>

## hold-time (OAM)

---

<b>Syntax</b>	hold-time <i>seconds</i> ;
<b>Hierarchy Level</b>	<p>[edit protocols oam],</p> <p>[edit protocols oam gre-tunnel interface <i>interface-name</i>]</p>
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	Length of time the originating end of a GRE tunnel waits for keepalive packets from the other end of the tunnel before marking the tunnel as operationally down.
<b>Options</b>	<p><b><i>seconds</i></b>—Hold-time value.</p> <p><b>Default:</b> 5 seconds</p> <p><b>Range:</b> 5 through 250 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">GRE Keepalive Time Overview on page 9</a></li> <li>• <a href="#">Configuring GRE Keepalive Time on page 9</a></li> <li>• <a href="#">keepalive-time on page 99</a></li> </ul>

## interfaces

---


<b>Syntax</b>	<code>interfaces { ... }</code>
<b>Hierarchy Level</b>	[edit]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure interfaces on the router.
<b>Default</b>	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li></ul>

## ipsec-sa

---

<b>Syntax</b>	<code>ipsec-sa <i>sa-name</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>es-fpc/pic/port</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Specify the IP Security (IPsec) SA name associated with the interface.
<b>Options</b>	<i>sa-name</i> —IPsec SA name.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li><li>• <i>Junos OS Administration Library for Routing Devices</i></li></ul>

## keepalive-time

<b>Syntax</b>	keepalive-time <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols oam], [edit protocols oam gre-tunnel interface <i>interface-name</i> ], [edit protocols oam gre-tunnel interface <i>interface-name.unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 10.2.
<b>Description</b>	Time difference between consecutive keepalive packets in a GRE tunnel.
<div>  <p><b>NOTE:</b> Support for GRE keepalive packets on MPC line cards became available as of Junos OS Release 11.4.</p> </div>	
<b>Options</b>	<p><b>seconds</b>—Keepalive time value.</p> <p><b>Default:</b> 1 second</p> <p><b>Range:</b> 1 through 50 seconds</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">GRE Keepalive Time Overview on page 9</a></li> <li>• <a href="#">Configuring GRE Keepalive Time on page 9</a></li> <li>• <a href="#">hold-time on page 97</a></li> </ul>

## key

---

<b>Syntax</b>	<code>key number;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">tunnel</a> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> <a href="#">tunnel</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Identify an individual traffic flow within a tunnel, as defined in RFC 2890, <i>Key and Sequence Number Extensions to GRE</i> . On M Series and T Series routers, you can configure the GRE interface on an Adaptive Services, Multiservices, or Tunnel PIC. On MX Series routers, configure the interface on a Multiservices DPC.
<b>Options</b>	<b>number</b> —Value of the key. <b>Range:</b> 0 through 4,294,967,295
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring a Key Number on GRE Tunnels on page 19</a></li></ul>

## multicast-only

---

<b>Syntax</b>	<code>multicast-only;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> family inet]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure the unit and family so that the interface can transmit and receive multicast traffic only. You can configure this property on the IP family only.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Restricting Tunnels to Multicast Traffic on page 23</a></li><li>• <a href="#">tunnel on page 109</a></li></ul>

## peer-unit

---

<b>Syntax</b>	<code>peer-unit <i>unit-number</i>;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure a peer relationship between two logical systems.
<b>Options</b>	<i>unit-number</i> —Peering logical system unit number.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Logical Tunnel Interfaces on page 25</a></li> </ul>

## peer-certificate-type

---

<b>Syntax</b>	<code>peer-certificate-type (pkcs7   x509-signature);</code>
<b>Hierarchy Level</b>	[edit services ipsec-vpn ike policy <i>policy-name</i> ]
<b>Release Information</b>	Statement introduced in Release 15.1 for MX Series routers.
<b>Description</b>	(MX Series routers only) Specify a preferred type of certificate (PKCS7 or X509). By default, X509 encoding format is used. With the flexibility to configure the encoding format in which certificate requests are sent to the peer, you can determine the type of certificate to be used depending on the type supported by the peer. For example, if the peer does not support PKCS7, certificate authentication cannot occur unless you configure the same type on MX Series routers as the initiator or sender.
<b>Options</b>	<ul style="list-style-type: none"> <li>• <b>pkcs7</b>—Public-Key Cryptography Standard #7.</li> <li>• <b>x509-signature</b>—X509 is an ITU-T standard for public key infrastructure.</li> </ul>
<b>Required Privilege Level</b>	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IKE Policies</a></li> </ul>

## reassemble-packets

---

<b>Syntax</b>	reassemble-packets;
<b>Hierarchy Level</b>	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 9.2.
<b>Description</b>	Enable reassembly of fragmented tunnel packets on generic routing encapsulation (GRE) tunnel interfaces.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Packet Reassembly on page 21</a></li></ul>

## redundancy-group (Interfaces)

---

<b>Syntax</b>	redundancy-group { member-interface <i>interface-name</i> { (active   backup); } }
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> ]
<b>Release Information</b>	Statement introduced in Junos OS Release 13.3.
<b>Description</b>	Configure member logical tunnels of redundant logical tunnels only on MX Series 3D Universal Edge Routers.
<b>Options</b>	<b>active</b> —Set the interface to the active mode.  <b>backup</b> —Set the interface to the backup mode.  The remaining statement is explained separately.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To view this statement in the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Redundant Logical Tunnels on page 32</a></li><li>• <a href="#">Configuring Redundant Logical Tunnels on page 31</a></li><li>• <a href="#">Redundant Logical Tunnels Overview on page 28</a></li><li>• <a href="#">redundancy-group (Logical Tunnels) on page 103</a></li></ul>

## redundancy-group (Logical Tunnels)

<b>Syntax</b>	<pre> redundancy-group {   interface-type {     redundant-logical-tunnel {       device <i>count</i>;     }   } } </pre>
<b>Hierarchy Level</b>	[edit chassis]
<b>Release Information</b>	<p>Statement introduced in Junos OS Release 13.3.</p> <p>Support for up to 255 redundant logical tunnels added to Junos OS Release 13.3R3 and 14.1R2.</p>
<b>Description</b>	Configure redundant logical tunnels only on MX Series 3D Universal Edge Routers.
<b>Options</b>	<p><b><i>count</i></b>—Specify the number of the redundant logical tunnels. For Junos OS Release 13.3R1, 13.3R2, and 14.1R1 the valid range is from 1 through 16. For Junos OS Release 13.3R3 and 14.1R2 the valid range is from 1 through 255.</p> <p>—</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To view this statement in the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Redundant Logical Tunnels on page 32</a></li> <li>• <a href="#">Configuring Redundant Logical Tunnels on page 31</a></li> <li>• <a href="#">Redundant Logical Tunnels Overview on page 28</a></li> <li>• <a href="#">redundancy-group (Interfaces) on page 102</a></li> </ul>

## routing-instance

---

<b>Syntax</b>	<code>routing-instance {     <b>destination</b> <i>routing-instance-name</i>; }</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.
<b>Description</b>	Specify the destination routing instance that points to the routing table containing the tunnel destination address.
<b>Default</b>	The default Internet routing table <b>inet.0</b> .
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Tunnel Interfaces for Routing Table Lookup on page 59</a></li></ul>

## routing-instances

---

<b>Syntax</b>	<code>routing-instances <i>routing-instance-name</i> { ... }</code>
<b>Hierarchy Level</b>	[edit], [edit logical-systems <i>logical-system-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	Configure an additional routing entity for a router or switch. You can create multiple instances of BGP, IS-IS, OSPF, OSPF version 3 (OSPFv3), and RIP for a router or switch.
<b>Default</b>	Routing instances are disabled for the router or switch.
<b>Options</b>	<b><i>routing-instance-name</i></b> —Name of the routing instance, a maximum of 31 characters. The remaining statements are explained separately.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring EVPN Routing Instances</a></li><li>• <a href="#">Configuring Routing Instances on PE Routers in VPNs</a></li></ul>

## routing-options

---

<b>Syntax</b>	routing-options { ... }
	For information on complete list of <b>routing-options</b> , see <i>[edit routing-options] Hierarchy Level</i> .
<b>Hierarchy Level</b>	[edit], [edit logical-systems <i>logical-system-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> ], [edit routing-instances <i>routing-instance-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	Configure protocol-independent routing properties.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Protocol-Independent Routing Properties Feature Guide for Routing Devices</i></li> <li>• <i>[edit routing-options] Hierarchy Level</i></li> </ul>

## source

---

<b>Syntax</b>	source <i>source-address</i> ;
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>family</b> inet <b>address</b> <i>address</i> ], [edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> tunnel]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	For tunnel and encryption interfaces, specify the source address.
<b>Options</b>	<b><i>source-address</i></b> —Address of the source side of the connection.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li> <li>• <i>Configuring Traffic Sampling</i></li> <li>• <i>Configuring Flow Monitoring</i></li> </ul>

## source

---

<b>Syntax</b>	<code>source source-address;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> <b>tunnel</b> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches. Statement introduced in Junos OS Release 13.2 for the QFX Series. Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	Specify the source address of the tunnel.
<b>Default</b>	If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.
<b>Options</b>	<b>source-address</b> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tunnel Services Overview on page 3</a></li></ul>

## source-address

---

<b>Syntax</b>	<code>source-address address;</code>
<b>Hierarchy Level</b>	[edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> routing-options <b>dynamic-tunnels</b> <i>tunnel-name</i> ], [edit logical-systems <i>logical-system-name</i> routing-options <b>dynamic-tunnels</b> <i>tunnel-name</i> ], [edit routing-instances <i>routing-instance-name</i> routing-options <b>dynamic-tunnels</b> <i>tunnel-name</i> ], [edit routing-options <b>dynamic-tunnels</b> <i>tunnel-name</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.3 for ACX Series routers.
<b>Description</b>	Configure the tunnel source address.
<b>Options</b>	<b>address</b> —Name of the source address.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Dynamic Tunnels on page 63</a></li></ul>

---

## ttl

---

<b>Syntax</b>	<code>ttl value;</code>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <a href="#">unit number</a> <a href="#">tunnel</a> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1 for EX Series switches.
<b>Description</b>	Set the time-to-live value bit in the header of the outer IP packet.
<b>Options</b>	<b>value</b> —Time-to-live value. <b>Range:</b> 0 through 255 <b>Default:</b> 64
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Tunnel Properties</a></li><li>• <a href="#">Tunnel Services Overview on page 3</a></li></ul>

## tunnel

---

<b>Syntax</b>	<pre>tunnel {   backup-destination destination-address;   destination destination-address;   routing-instance {     destination routing-instance-name;   }   source source-address;   ttl number; }</pre>
<b>Hierarchy Level</b>	[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i> ]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4.
<b>Description</b>	<p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).</p> <p>The statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li><li>• <i>Tunnel Properties</i></li><li>• <i>Junos OS VPNs Library for Routing Devices</i></li></ul>

## tunnel

<b>Syntax</b>	<pre> tunnel {     allow-fragmentation;     backup-destination address;     destination destination-address;     do-not-fragment;     key number;     routing-instance {         destination routing-instance-name;     }     source source-address;     ttl number; } </pre>
<b>Hierarchy Level</b>	<p>[edit interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> <b>unit</b> <i>logical-unit-number</i>]</p>
<b>Release Information</b>	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p>
<b>Description</b>	<p>Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or virtual private networks (VPNs).</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Encryption Interfaces on page 69</a></li> <li>• <i>Junos OS VPNs Library for Routing Devices</i></li> </ul>

## unit (Interfaces)

---

**Syntax**    `unit logical-unit-number {  
              family inet {  
                  ipsec-sa sa-name;  
              }  
              tunnel {  
                  backup-destination destination-address;  
                  destination destination-address;  
                  routing-instance {  
                      destination routing-instance-name;  
                  }  
                  source source-address;  
                  ttl number;  
              }  
          }`

**Hierarchy Level**    [edit [interfaces](#) *interface-name*]

**Release Information**    Statement introduced before Junos OS Release 7.4.

**Description**    Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

**Options**    *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,384

The remaining statements are explained separately.

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

**Related Documentation**

- [Configuring Encryption Interfaces on page 69](#)
- *Junos OS Network Interfaces Library for Routing Devices*
- *Junos OS Network Interfaces Library for Routing Devices* for other statements that do not affect services interfaces.

## unit (Interfaces)

<b>Syntax</b>	<pre> unit logical-unit-number {   peer-unit unit-number;   reassemble-packets;   tunnel {     allow-fragmentation;     backup-destination address;     destination destination-address;     do-not-fragment;     key number;     routing-instance {       destination routing-instance-name;     }     source source-address;     ttl number;   } } </pre>
<b>Hierarchy Level</b>	[edit <a href="#">interfaces</a> interface-name], [edit logical-systems logical-system-name <a href="#">interfaces</a> interface-name]
<b>Release Information</b>	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 15.1X53-D10 for QFX10000 switches.
<b>Description</b>	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
<b>Options</b>	<p><b>logical-unit-number</b>—Number of the logical unit.</p> <p><b>Range:</b> 0 through 16,384</p> <p>The remaining statements are explained separately.</p>
<b>Required Privilege Level</b>	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces.</li> </ul>



## CHAPTER 14

# Operational Commands

- clear ike security-associations
- clear ipsec security-associations
- request ipsec switch
- request security certificate enroll (Signed)
- request security certificate enroll (Unsigned)
- request security key-pair
- request system certificate add
- show ike security-associations
- show interfaces (Encryption)
- show interfaces (GRE)
- show interfaces (IP-over-IP)
- show interfaces (Logical Tunnel)
- show interfaces (Multicast Tunnel)
- show interfaces (PIM)
- show interfaces (Virtual Loopback Tunnel)
- show ipsec certificates
- show ipsec redundancy
- show ipsec security-associations
- show system certificate

## clear ike security-associations

---

<b>Syntax</b>	clear ike security-associations <destination-ip-address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(Encryption interface on M Series and T Series routers only) Clear information about the current Internet Key Exchange (IKE) security association. This command is valid for dynamic security associations only. For IKEv2, this command creates new security associations for IKE SA and IPSEC SAs.
<b>Options</b>	<b>none</b> —Clear all IKE security associations.  <b>destination-ip-address</b> —(Optional) Clear the IKE security association at the specified destination address.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"><li>• <a href="#">show ike security-associations on page 123</a></li></ul>
<b>List of Sample Output</b>	<a href="#">clear ike security-associations on page 114</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### clear ike security-associations

```
user@host> clear ike security-associations
```

## clear ipsec security-associations

<b>Syntax</b>	<code>clear ipsec security-associations</code> <code>&lt;sa-name&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(Encryption interface on M Series and T Series routers only) Clear information about the current IP Security (IPsec) security association. This command is valid for dynamic security associations only. For IKEv1, this command creates new security associations for IKE SA and IPSEC SAs.
<b>Options</b>	<b>none</b> —Clear all IPsec security associations.  <b>sa-name</b> —(Optional) Clear the specified security association.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ipsec security-associations on page 170</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">clear ipsec security-associations on page 115</a>
<b>Output Fields</b>	See the <a href="#">show ipsec security-associations</a> for an explanation of output fields.

## Sample Output

### clear ipsec security-associations

The following output from the **show ipsec security-associations detail** command is displayed before and after the **clear ipsec security-associations** command is issued:

```
user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

Direction: inbound, SPI: 242379418, State: Installed
Mode: tunnel, Type: dynamic
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expires in 22979 seconds
Hard lifetime: Expires in 28739 seconds

Direction: outbound, SPI: 368592771, State: Installed
Mode: tunnel, Type: dynamic
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expires in 22979 seconds
Hard lifetime: Expires in 28739 seconds

user@host> clear ipsec security-associations

user@host> show ipsec security-associations detail
Security association: sa-dynamic, Interface family: Up

Direction: inbound, SPI: 1031597683, State: Installed
```

Mode: tunnel, Type: dynamic  
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None  
Soft lifetime: Expires in 23037 seconds  
Hard lifetime: Expires in 28797 seconds

Direction: outbound, SPI: 1618419878, State: Installed  
Mode: tunnel, Type: dynamic  
Protocol: ESP, Authentication: hmac-md5-96, Encryption: None  
Soft lifetime: Expires in 23037 seconds  
Hard lifetime: Expires in 28797 seconds

## request ipsec switch

---


<b>Syntax</b>	<code>request ipsec switch (interface &lt;es-fpc/pic/port&gt;   security-associations &lt;sa-name&gt;)</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series, PTX Series, and T Series routers and EX Series switches only) Manually switch from the primary to the backup encryption services interface, or switch from the primary to the backup IP Security (IPsec) tunnel.
<b>Options</b>	<code>interface &lt;es-fpc/pic/port&gt;</code> —Switch to the backup encryption interface. <code>security-associations &lt;sa-name&gt;</code> —Switch to the backup tunnel.
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">show ipsec redundancy on page 168</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">request ipsec switch security-associations on page 117</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request ipsec switch security-associations

```
user@host> request ipsec switch security-associations sa-private
```

## request security certificate enroll (Signed)

<b>Syntax</b>	request security certificate enroll filename <i>filename</i> subject <i>subject</i> alternative-subject <i>alternative-subject</i> certification-authority <i>certification-authority</i> encoding (binary   pem) key-file <i>key-file</i> domain-name <i>domain-name</i>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Obtain a signed certificate from a certificate authority (CA). The signed certificate validates the CA and the owner of the certificate. The results are saved in a specified file to the <code>/var/etc/ikecert</code> directory.
<div>  <p><b>NOTE:</b> For FIPS mode, the digital security certificates must be compliant with the National Institute of Standards and Technology (NIST) SP 800-131A standard. The <code>request security key-pair</code> command is deprecated and not available with Junos in FIPS mode because it generates RSA and DSA keys with sizes of 512 and 1024 bits that are not compliant with the NIST SP 800-131A standard.</p> </div>	
<b>Options</b>	<p><b>filename <i>filename</i></b>—File that stores the certificate.</p> <p><b>subject <i>subject</i></b>—Distinguished name (<b>dn</b>), which consists of a set of components—for example, an organization (<b>o</b>), an organization unit (<b>ou</b>), a country (<b>c</b>), and a locality (<b>l</b>).</p> <p><b>alternative-subject <i>alternative-subject</i></b>—Tunnel source address.</p> <p><b>certification-authority <i>certification-authority</i></b>—Name of the certificate authority profile in the configuration.</p> <p><b>encoding (binary   pem)</b>—File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default format is binary.</p> <p><b>key-file <i>key-file</i></b>—File containing a local private key.</p> <p><b>domain-name <i>domain-name</i></b>—Fully qualified domain name.</p>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<code>request security certificate enroll filename subject alternative-subject certification-authority key-file domain-name (Signed)</code> on page 119
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

request security certificate enroll filename subject alternative-subject certification-authority key-file  
domain-name (Signed)

```
user@host> request security certificate enroll filename host.crt subject c=uk,o=london  
alternative-subject 10.50.1.4 certification-authority verisign key-file host-1.prv domain-name  
host.example.com  
CA name: example.com CA file: ca_verisign  
local pub/private key pair: host.prv  
subject: c=uk,o=london domain name: host.example.com  
alternative subject: 10.50.1.4  
Encoding: binary  
Certificate enrollment has started. To view the status of your enrollment, check  
the key management process (kmd) log file at /var/log/kmd. <-----
```

## request security certificate enroll (Unsigned)

---


<b>Syntax</b>	<code>request security certificate enroll filename <i>filename</i> ca-file <i>ca-file</i> ca-name <i>ca-name</i> encoding (binary   perm) url <i>url</i></code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Obtain a certificate from a certificate authority (CA). The results are saved in a specified file to the <code>/var/etc/ikecert</code> directory.
<b>Options</b>	<b>filename <i>filename</i></b> —File that stores the public key certificate.  <b>ca-file <i>ca-file</i></b> —Name of the certificate authority profile in the configuration.  <b>ca-name <i>ca-name</i></b> —Name of the certificate authority.  <b>encoding (binary   pem)</b> —File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default value is <b>binary</b> .  <b>url <i>url</i></b> —Certificate authority URL.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security certificate enroll filename ca-file ca-name url (Unsigned) on page 120</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### request security certificate enroll filename ca-file ca-name url (Unsigned)

```
user@host> request security certificate enroll filename ca_verisign ca-file verisign ca-name
example.com urlxyzcompany URL
http://<verisign ca-name xyzcompany url>/cgi-bin/pkiclient.exe CA name: example.com
CA file: verisign Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
the key management process (kmd) log file at /var/log/kmd. <-----
```

## request security key-pair

<b>Syntax</b>	<code>request security key-pair <i>filename</i></code> <code>&lt;size <i>key-size</i>&gt;</code> <code>&lt;type (rsa   dsa)&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches.
<b>Description</b>	(Encryption interface on M Series and T Series routers and EX Series switches only) Generate a public and private key pair for a digital certificate.
<div>  <p><b>NOTE:</b> The <code>request security-certificates</code> command is deprecated and are not available with Junos in FIPS mode because security certificates are not compliant with the NIST SP 800-131A standard.</p> </div>	
<b>Options</b>	<p><b><i>filename</i></b>—Name of a file in which to store the key pair.</p> <p><b><i>size key-size</i></b>—(Optional) Key size, in bits. The key size can be <b>512</b>, <b>1024</b>, or <b>2048</b>. The default value is <b>1024</b>.</p> <p><b><i>type</i></b>—(Optional) Algorithm used to encrypt the key:</p> <ul style="list-style-type: none"> <li>• <b>rsa</b>—RSA algorithm. This is the default.</li> <li>• <b>dsa</b>—Digital signature algorithm with Secure Hash Algorithm (SHA).</li> </ul>
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request security key-pair on page 121</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

## Sample Output

### request security key-pair

```
user@host> request security key-pair security-key-file
```

## request system certificate add

---

<b>Syntax</b>	<code>request system certificate add (<i>filename</i>   terminal)</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
<b>Description</b>	(Encryption interface on M Series and T Series routers, PTX Series, and QFX Series switches only) Add a certificate provided by the Juniper Networks certificate authority (CA).
<b>Options</b>	<b><i>filename</i></b> —Filename (URL, local, or remote). <b><i>terminal</i></b> —Use login terminal.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">request system certificate add terminal on page 122</a>
<b>Output Fields</b>	When you enter this command, you are provided feedback on the status of your request.

### Sample Output

#### [request system certificate add terminal](#)

```
user@host> request system certificate add terminal
```

## show ike security-associations

<b>Syntax</b>	show ike security-associations <brief   detail> <peer-address>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(Encryption interface on M Series and T Series routers only) Display information about Internet Key Exchange (IKE) security associations.
<b>Options</b>	<p><b>none</b>—Display standard information about all IKE security associations.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>peer-address</b>—(Optional) Display IKE security associations for the specified peer address.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ike security-associations on page 114</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ike security-associations on page 126</a> <a href="#">show ike security-associations detail on page 126</a>
<b>Output Fields</b>	<p><a href="#">Table 5 on page 123</a> lists the output fields for the <b>show ike security-associations</b> command. Output fields are listed in the approximate order in which they appear.</p>

**Table 5: show ike security-associations Output Fields**

Field Name	Field Description	Level of Output
<b>IKE peer</b>	Remote end of the IKE negotiation.	<b>detail</b>
<b>Role</b>	Part played in the IKE session. The router triggering the IKE negotiation is the initiator, and the router accepting the first IKE exchange packets is the responder.	<b>detail</b>
<b>Remote Address</b>	Responder's address.	none specified
<b>State</b>	State of the IKE security association: <ul style="list-style-type: none"> <li>• <b>Matured</b>—The IKE security association is established.</li> <li>• <b>Not matured</b>—The IKE security association is in the process of negotiation.</li> </ul>	none specified
<b>Initiator cookie</b>	When the IKE negotiation is triggered, a random number is sent to the remote node.	All levels

Table 5: show ike security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Responder cookie</b>	<p>The remote node generates its own random number and sends it back to the initiator as a verification that the packets were received.</p> <p>Of the numerous security services available, protection against denial of service (DoS) is one of the most difficult to address. A “cookie” or anticlogging token (ACT) is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity. An exchange prior to CPU-intensive public key operations can thwart some DoS attempts (such as simple flooding with invalid IP source addresses).</p>	All levels
<b>Exchange type</b>	<p>Specifies the number of messages in an IKE exchange, and the payload types that are contained in each message. Each exchange type provides a particular set of security services, such as anonymity of the participants, perfect forward secrecy of the keying material, and authentication of the participants. Junos OS supports two types of exchanges:</p> <ul style="list-style-type: none"> <li>• <b>Main</b>—The exchange is done with six messages. <b>Main</b> encrypts the payload, protecting the identity of the neighbor.</li> <li>• <b>Aggressive</b>—The exchange is done with three messages. <b>Aggressive</b> does not encrypt the payload, leaving the identity of the neighbor unprotected.</li> </ul>	All Levels
<b>Authentication method</b>	Type of authentication determines which payloads are exchanged and when they are exchanged. The Junos OS supports only <b>pre-shared keys</b> .	<b>detail</b>
<b>Local</b>	Prefix and port number of the local end.	<b>detail</b>
<b>Remote</b>	Prefix and port number of the remote end.	<b>detail</b>
<b>Lifetime</b>	Number of seconds remaining until the IKE security association expires.	<b>detail</b>
<b>Algorithms</b>	<p>Header for the IKE algorithms output.</p> <ul style="list-style-type: none"> <li>• <b>Authentication</b>—Type of authentication algorithm used: <b>md5</b> or <b>sha1</b>.</li> <li>• <b>Encryption</b>—Type of encryption algorithm used: <b>des-cbc</b>, <b>3des-cbc</b>, or <b>None</b>.</li> <li>• <b>Pseudo random function</b>—Function that generates highly unpredictable random numbers: <b>hmac-md5</b> or <b>hmac-sha1</b>.</li> </ul>	<b>detail</b>
<b>Traffic statistics</b>	<p>Number of bytes and packets received and transmitted on the IKE security association.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the IKE security association.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the IKE security association.</li> </ul>	<b>detail</b>

Table 5: show ike security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Flags</b>	Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> <li>• <b>caller notification sent</b>—Caller program notified about the completion of the IKE negotiation.</li> <li>• <b>waiting for done</b>—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.</li> <li>• <b>waiting for remove</b>—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.</li> <li>• <b>waiting for policy manager</b>—Negotiation is waiting for a response from the policy manager.</li> </ul>	<b>detail</b>
<b>IPsec security associates</b>	Number of IPsec security associations created and deleted with this IKE security association.	<b>detail</b>
<b>Phase 2 negotiations in progress</b>	Number of phase 2 IKE negotiations in progress and status information: <ul style="list-style-type: none"> <li>• <b>Negotiation type</b>—Type of phase 2 negotiation. The Junos OS currently supports <b>quick mode</b>.</li> <li>• <b>Message ID</b>—Unique identifier for a phase 2 negotiation.</li> <li>• <b>Local identity</b>—Identity of the local phase 2 negotiation. The format is <i>id-type-name (proto-name:port-number,[O..id-data-len] = iddata-presentation)</i></li> <li>• <b>Remote identity</b>—Identity of the remote phase 2 negotiation. The format is <i>id-type-name (proto-name:port-number,[O..id-data-len] = iddata-presentation)</i></li> <li>• <b>Flags</b>—Notification to the key management process of the status of the IKE negotiation: <ul style="list-style-type: none"> <li>• <b>caller notification sent</b>—Caller program notified about the completion of the IKE negotiation.</li> <li>• <b>waiting for done</b>—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.</li> <li>• <b>waiting for remove</b>—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.</li> <li>• <b>waiting for policy manager</b>—Negotiation is waiting for a response from the policy manager.</li> </ul> </li> </ul>	<b>detail</b>

## Sample Output

### show ike security-associations

```
user@host> show ike security-associations
Remote Address  State      Initiator cookie  Responder cookie  Exchange type
4.4.4.4         Matured          93870456fa000011 723a20713700003e Main
```

### show ike security-associations detail

```
user@host> show ike security-associations detail
IKE peer 4.4.4.4
  Role: Initiator, State: Matured
  Initiator cookie: cf22bd81a7000001, Responder cookie: fe83795c2800002e
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 4.4.4.5:500, Remote: 4.4.4.4:500
  Lifetime: Expires in 187 seconds
  Algorithms:
    Authentication      : md5
    Encryption          : 3des-cbc
    Pseudo random function: hmac-md5
  Traffic statistics:
    Input bytes  :          1000
    Output bytes :          1280
    Input packets:           5
    Output packets:          9
  Flags: Caller notification sent
  IPsec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

Negotiation type: Quick mode, Role: Initiator, Message ID: 3582889153
  Local: 4.4.4.5:500, Remote: 4.4.4.4:500
  Local identity: ipv4_subnet(tcp:80,[0..7]=10.1.1.0/24)
  Remote identity: ipv4_subnet(tcp:100,[0..7]=10.1.2.0/24)
  Flags: Caller notification sent, Waiting for done
```

## show interfaces (Encryption)

<b>Syntax</b>	<pre>show interfaces es-fpc/pic/port:channel &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(M Series and T Series routers only) Display status information about the specified encryption interface.
<b>Options</b>	<p><b>es-fpc/pic/port:channel</b>—Display standard status information about the specified encryption interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces (Encryption) on page 130</a> <a href="#">show interfaces brief (Encryption) on page 130</a> <a href="#">show interfaces detail (Encryption) on page 130</a> <a href="#">show interfaces extensive (Encryption) on page 131</a>
<b>Output Fields</b>	Table 6 on page 127 lists the output fields for the <b>show interfaces</b> (Encryption) command. Output fields are listed in the approximate order in which they appear.

**Table 6: Encryption show interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>

Table 6: Encryption show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Encapsulation being used on the interface.	All levels
<b>Link-level type</b>	Encapsulation being used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Input rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output rate</b>	Output rate in bps and pps.	None specified
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface.</li> <li>• <b>Anti-replay failures</b>—Total number of antireplay failures seen on all tunnels configured on the ES PIC.</li> <li>• <b>Authentication</b>—Total number of authentication failures seen on all tunnels configured on the ES PIC.</li> </ul>	<b>detail extensive</b>
<b>Egress queues</b>	Total number of egress queues supported on the specified interface.	<b>detail extensive</b>
<b>Queue counters</b>	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> <li>• <b>Queued packets</b>—Number of queued packets.</li> <li>• <b>Transmitted packets</b>—Number of transmitted packets.</li> <li>• <b>Dropped packets</b>—Number of packets dropped by the ASIC's RED mechanism.</li> </ul>	<b>detail extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>

Table 6: Encryption show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface. Possible values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
IP-Header	IP header of the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
<i>protocol-family</i>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	brief
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified
Traffic statistics	Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit 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.	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. 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.	detail extensive
Transit statistics	Statistics for traffic transiting the router. 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.	detail extensive
Protocol	Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , <b>mpls</b> .	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	detail extensive
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> . Address	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none

Table 6: Encryption show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces (Encryption)

```

user@host> show interfaces es-0/3/0
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 71
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45)
  Flags: Hardware-Down Point-To-Point SNMP-Traps
  IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 3800
  Flags: None
  Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
  Destination: 10.10.0.2, Local: 10.10.0.1

```

### show interfaces brief (Encryption)

```

user@host> show interfaces es-0/3/0 brief
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface es-0/3/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps
  IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC
  inet 10.10.0.1 --> 10.10.0.2s

```

### show interfaces detail (Encryption)

```

user@host> show interfaces es-0/3/0 detail
Physical interface: es-0/3/0, Enabled, Physical link is Up
  Interface index: 138, SNMP ifIndex: 71, Generation: 21
  Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   : 0 0 bps
    Output bytes  : 0 0 bps

```

```

Input packets:                0                0 pps
Output packets:               0                0 pps
Anti-replay failures         : 0
Authentication failures      : 0
Egress queues: 4 supported, 4 in use
Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort          0                0                0
1 expedited-fo         0                0                0
2 assured-forw         0                0                0
3 network-cont         0                0                0

```

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45) (Generation 9)

Flags: Hardware-Down Point-To-Point SNMP-Traps

IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC

Traffic statistics:

```

Input bytes :                0
Output bytes :                0
Input packets:               0
Output packets:              0

```

Local statistics:

```

Input bytes :                0
Output bytes :                0
Input packets:               0
Output packets:              0

```

Transit statistics:

```

Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:               0                0 pps
Output packets:              0                0 pps

```

Protocol inet, MTU: 3800, Generation: 22, Route table: 0

Flags: None

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 10.10.0.2, Local: 10.10.0.1, Broadcast: Unspecified,  
Generation: 26

### show interfaces extensive (Encryption)

```
user@host> show interfaces es-0/3/0 extensive
```

Physical interface: es-0/3/0, Enabled, Physical link is Up

Interface index: 138, SNMP ifIndex: 71, Generation: 21

Type: IPSEC, Link-level type: IPSEC-over-IP, MTU: 3900, Speed: 800mbps

Hold-times : Up 0 ms, Down 0 ms

Device flags : Present Running

Interface flags: Point-To-Point SNMP-Traps

Statistics last cleared: Never

Traffic statistics:

```

Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:               0                0 pps
Output packets:              0                0 pps

```

Anti-replay failures : 0

Authentication failures : 0

Egress queues: 4 supported, 4 in use

```

Queue counters:      Queued packets  Transmitted packets  Dropped packets

0 best-effort          0                0                0

```

1 expedited-fo	0	0	0
2 assured-forw	0	0	0
3 network-cont	0	0	0

Logical interface es-0/3/0.0 (Index 70) (SNMP ifIndex 45) (Generation 9)

Flags: Hardware-Down Point-To-Point SNMP-Traps

IP-Header 10.0.10.2:10.0.10.1::df:64:00000000 Encapsulation: IPSEC

Traffic statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Transit statistics:

Input bytes : 0 0 bps

Output bytes : 0 0 bps

Input packets: 0 0 pps

Output packets: 0 0 pps

Protocol inet, MTU: 3800, Generation: 22, Route table: 0


Flags: None

Addresses, Flags: Dest-route-down Is-Preferred Is-Primary

Destination: 10.10.0.2, Local: 10.10.0.1, Broadcast: Unspecified,

Generation: 26

## show interfaces (GRE)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Command introduced in Junos OS Release 13.2 for the QFX Series.</p> <p>Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.</p>
<b>Description</b>	Display status information about the specified generic routing encapsulation (GRE) interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers and EX Series switches, the interface type is <b><i>gr-fpc/pic/port</i></b>.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified output level of interface information.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
	<div>  <p><b>NOTE:</b> You can configure generic routing encapsulation (GRE) interfaces (<i>gre-x/y/z</i>) only for GMPLS control channels. GRE interfaces are not supported or configurable for other applications. For more information about GMPLS, see the <i>MPLS Applications Feature Guide for Routing Devices</i> and the <i>Junos OS, Release 15.1</i>.</p> </div>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (GRE) on page 138</a></p> <p><a href="#">show interfaces brief (GRE) on page 138</a></p> <p><a href="#">show interfaces detail (GRE) on page 138</a></p> <p><a href="#">show interfaces (Layer 2 Services Over GRE Interfaces) on page 139</a></p> <p><a href="#">show interfaces extensive (Layer 2 Services Over GRE Interfaces) on page 139</a></p> <p><a href="#">show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch on page 140</a></p> <p><a href="#">show interfaces extensive (GRE) on page 141</a></p>

**Output Fields** Table 7 on page 134 lists the output fields for the **show interfaces** (GRE) command. Output fields are listed in the approximate order in which they appear.

**Table 7: GRE show interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Type of interface.	All levels
<b>Link-level type</b>	Encapsulation used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Device Flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface Flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Input rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	<p>The number of and the rate at which input and output bytes and packets are received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li><b>Input bytes</b>—Number of bytes received on the interface.</li> <li><b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li><b>Input packets</b>—Number of packets received on the interface.</li> <li><b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Logical Interface</b>		

Table 7: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support.	<b>detail extensive</b>
<b>Flags</b>	<p>Information about the logical interface. Possible values listed in the “Logical Interface Flags” section under <i>Common Output Fields Description</i>. describe general information about the logical interface.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>Reassemble-Pkts</b>—If the <b>Flags</b> field includes this string, the GRE tunnel is configured to reassemble tunnel packets that were fragmented after tunnel encapsulation.</li> </ul>	All levels
<b>IP-Header</b>	<p>IP header of the logical interface. If the <b>tunnel key</b> statement is configured, this information is included in the <b>IP Header</b> entry.</p> <p>GRE-specific information about the logical interface is indicated by the presence or absence of the following value in this field:</p> <ul style="list-style-type: none"> <li>• <b>df</b>—If the <b>IP-Header</b> field includes this string immediately following the 16 bits of identification information (that is, if <b>:df</b> displays after the twelfth byte), the GRE tunnel is configured to allow fragmentation of GRE packets after encapsulation.</li> </ul>	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>L2 Routing Instance</b>	Name of the Layer 2 routing instance associated with the GRE interface.	All levels
<b>L3 Routing Instance</b>	Name of the Layer 3 routing instance associated with the GRE interface.	All levels
<b>Copy-tos-to-outer-ip-header</b>	<p>Status of type of service (ToS) bits in the GRE packet header:</p> <ul style="list-style-type: none"> <li>• <b>On</b>—ToS bits were copied from the payload packet header into the header of the IP packet sent through the GRE tunnel.</li> <li>• <b>Off</b>—ToS bits were not copied from the payload packet header and are set to 0 in the GRE packet header.</li> </ul> <p><b>NOTE:</b> EX Series switches do not support copying ToS bits to the encapsulated packet, so the value of this field is always <b>Off</b> in switch output.</p>	<b>detail extensive</b>
<b>Gre keepalives configured</b>	<p>Indicates whether a GRE keepalive time and hold time are configured for the GRE tunnel.</p> <p><b>NOTE:</b> EX Series switches do not support configuration of GRE tunnel keepalive times and hold times, so the value of this field is always <b>Off</b> in switch output.</p>	<b>detail extensive</b>

Table 7: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Gre keepalives adjacency state</b>	Status of the other end of the GRE tunnel: <b>Up</b> or <b>Down</b> . If keepalive messages are not received by either end of the GRE tunnel within the hold-time period, the GRE keepalive adjacency state is down even when the GRE tunnel is up.	<b>detail extensive</b>
<b>Input packets</b>	Number of packets received on the logical interface.	None specified
<b>Output packets</b>	Number of packets transmitted on the logical interface.	None specified
<b>Traffic statistics</b>	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit 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.</p> <ul style="list-style-type: none"> <li>• <b>Input rate</b>—Rate of bits and packets received on the interface.</li> <li>• <b>Output rate</b>—Rate of bits and packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	Statistics for traffic received from and transmitted to the Routing Engine. 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.	<b>detail extensive</b>
<b>Transit statistics</b>	Statistics for traffic transiting the router. 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.	<b>detail extensive none</b>
<b>Protocol</b>	Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , or <b>mpls</b> .	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>

Table 7: GRE show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces (GRE)

```

user@host> show interfaces gr-1/2/0
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 1476
  Flags: None
  Addresses, Flags: Is-Primary
    Local: 1.10.1.1

```

### show interfaces brief (GRE)

```

user@host> show interfaces gr-1/2/0 brief
Physical interface: gr-1/2/0, Enabled, Physical link is Up
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps

Logical interface gr-1/2/0.0
  Flags: Hardware-Down Point-To-Point SNMP-Traps 0x4000
  IP-Header 10.10.0.2:10.10.0.1:47:df:64:0000000000000000
  Encapsulation: GRE-NULL
  inet 10.100.0.1/30
  mpls

```

### show interfaces detail (GRE)

```

user@host> show interfaces gr-1/2/0 detail
Physical interface: gr-0/0/0, Enabled, Physical link is Up
  Interface index: 132, SNMP ifIndex: 26, Generation: 13
  Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes : 0 0 bps
    Output bytes : 0 0 bps
    Input packets: 0 0 pps
    Output packets: 0 0 pps

Logical interface gr-0/0/0.0 (Index 68) (SNMP ifIndex 47) (Generation 8)
  Flags: Point-To-Point SNMP-Traps 16384
  IP-Header 1.1.1.2:1.1.1.1:47:df:64:0000000000000000 Encapsulation: GRE-NULL
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0

```

```

Output packets:                0
Local statistics:
Input bytes :                   0
Output bytes :                   0
Input packets:                  0
Output packets:                 0
Transit statistics:
Input bytes :                   0          0 bps
Output bytes :                   0          0 bps
Input packets:                  0          0 pps
Output packets:                 0          0 pps
Protocol inet, MTU: 1476, Generation: 12, Route table: 0
Flags: None
Addresses, Flags: Is-Primary
Destination: Unspecified, Local: 1.10.1.1, Broadcast: Unspecified,
Generation: 15

```

### show interfaces (Layer 2 Services Over GRE Interfaces)

```

user@host> show interfaces gr-2/2/10
show interfaces gr-2/2/10
Physical interface: gr-2/2/10, Enabled, Physical link is Up
Interface index: 214, SNMP ifIndex: 690
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Input rate : 0 bps (0 pps)
Output rate : 0 bps (0 pps)

Logical interface gr-2/2/10.0 (Index 342) (SNMP ifIndex 10834)
Flags: Up Point-To-Point SNMP-Traps 0x4000 IP-Header
3.0.0.1:3.0.0.254:47:df:64:0000000000000000 Encapsulation: GRE-NUL
L2 Routing Instance: vs1, L3 Routing Instance: default
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Input packets : 2
Output packets: 0
Protocol bridge, MTU: 1476
Flags: Sendbroadcast-pkt-to-re
Addresses, Flags: Is-Preferred Is-Primary
Destination: 6/8, Local: 6.0.0.1, Broadcast: 6.255.255.255

```

### show interfaces extensive (Layer 2 Services Over GRE Interfaces)

```

user@host> show interfaces gr-2/2/10.0 extensive

Flags: SNMP-Traps Encapsulation: ENET2
L2 Routing Instance: vs1, L3 Routing Instance: default
Traffic statistics:
Input bytes :          58851250
Output bytes :          0
Input packets:        1279375
Output packets:        0
Local statistics:
Input bytes :          0
Output bytes :          0
Input packets:         0
Output packets:         0
Transit statistics:
Input bytes :          58851250          75136 bps

```

```

Output bytes : 0 0 bps
Input packets: 1279375 204 pps
Output packets: 0 0 pps
Protocol bridge, MTU: 1476, Generation: 175, Route table: 7
Flags: Access-Mode

```

### show interfaces detail (GRE) on an EX4200 Virtual Chassis Member Switch

```

user@host> show interfaces gr-2/0/15 detail
Physical interface: gr-2/0/15, Enabled, Physical link is Up
Interface index: 195, SNMP ifIndex: 846, Generation: 198
Type: GRE, Link-level type: GRE, MTU: Unlimited, Speed: 1000mbps
Hold-times : Up 0 ms, Down 0 ms
Current address: 00:1f:12:38:0f:d2, Hardware address: 00:1f:12:38:0f:d2
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Statistics last cleared: 2011-09-14 17:43:15 UTC (00:00:18 ago)
Traffic statistics:
Input bytes : 5600636 0 bps
Output bytes : 5600636 0 bps
Input packets: 20007 0 pps
Output packets: 20007 0 pps
IPv6 transit statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Logical interface gr-2/0/15.0 (Index 75) (SNMP ifIndex 847) (HW Token 4093)
(Generation 140)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 180.20.30.2:180.20.20.3:47:df:64:0000000000000000
Encapsulation: GRE-NULL
Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
Input bytes : 5600886
Output bytes : 2881784
Input packets: 20010
Output packets: 10018
Local statistics:
Input bytes : 398
Output bytes : 264
Input packets: 5
Output packets: 3
Transit statistics:
Input bytes : 5600488 0 bps
Output bytes : 2881520 0 bps
Input packets: 20005 0 pps
Output packets: 10015 0 pps
Protocol inet, Generation: 159, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 90.90.90/24, Local: 90.90.90.10, Broadcast: 90.90.90.255,
Generation: 144

Logical interface gr-2/0/15.1 (Index 80) (SNMP ifIndex 848) (HW Token 4088)
(Generation 150)
Flags: Point-To-Point SNMP-Traps 0x0
IP-Header 160.20.40.2:160.20.30.1:47:df:64:0000000000000000
Encapsulation: GRE-NULL

```

```

Copy-tos-to-outer-ip-header: Off
Gre keepalives configured: Off, Gre keepalives adjacency state: down
Traffic statistics:
  Input bytes :          260
  Output bytes :        2880148
  Input packets:          4
  Output packets:       10002
Local statistics:
  Input bytes :          112
  Output bytes :           0
  Input packets:          2
  Output packets:         0
Transit statistics:
  Input bytes :          148          0 bps
  Output bytes :        2880148      0 bps
  Input packets:           2          0 pps
  Output packets:       10002        0 pps
Protocol inet, Generation: 171, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
  Destination: 70.70.70/24, Local: 70.70.70.10, Broadcast: 70.70.70.255,
  Generation: 160

```

#### show interfaces extensive (GRE)

The output for the **show interfaces extensive** command is identical to that for the **show interfaces detail** command. For sample output, see [show interfaces detail \(GRE\) on page 138](#) and [show interfaces detail \(GRE\) on an EX4200 Virtual Chassis Member Switch on page 140](#).

## show interfaces (IP-over-IP)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified IP-over-IP interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b><i>ip-fpc/pic/port</i></b>.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (IP-over-IP) on page 144</a></p> <p><a href="#">show interfaces brief (IP-over-IP) on page 145</a></p> <p><a href="#">show interfaces detail (IP-over-IP) on page 145</a></p> <p><a href="#">show interfaces extensive (IP-over-IP) on page 145</a></p>
<b>Output Fields</b>	<p><a href="#">Table 8 on page 142</a> lists the output fields for the <b>show interfaces (IP-over-IP)</b> command. Output fields are listed in the approximate order in which they appear.</p>

Table 8: IP-over-IP show interfaces Output Fields

Field	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

Table 8: IP-over-IP show interfaces Output Fields (*continued*)

Field	Field Description	Level of Output
Type	Type of interface.	All levels
Link-level type	Encapsulation used on the physical interface.	All levels
MTU	MTU size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	detail extensive
Device flags	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
Input rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	detail extensive
<b>Logical Interface</b>		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support.	detail extensive
Flags	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
IP Header	IP header of the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified

Table 8: IP-over-IP show interfaces Output Fields (*continued*)

Field	Field Description	Level of Output
<b>Output packets</b>	Number of packets transmitted on the logical interface.	None specified
<b>Traffic statistics</b>	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit 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.</p> <ul style="list-style-type: none"> <li>• <b>Input rate</b>—Rate of bits and packets received on the interface.</li> <li>• <b>Output rate</b>—Rate of bits and packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	Statistics for traffic received from and transmitted to the Routing Engine. 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.	<b>detail extensive</b>
<b>Transit statistics</b>	Statistics for traffic transiting the router. 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.	<b>detail extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , or <b>mpls</b> .	<b>detail extensive none</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. If the protocol is <b>inet</b> , the IP address of the interface is also displayed.	<b>brief</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Routing table in which the logical interface address is located. For example, <b>0</b> refers to the routing table <b>inet.0</b> .	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>

## Sample Output

### show interfaces (IP-over-IP)

```

user@host> show interfaces ip-0/0/0
Physical interface: ip-0/0/0, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: SNMP-Traps
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

  Logical interface ip-0/0/0.0 (Index 69) (SNMP ifIndex 49)
    Flags: Point-To-Point SNMP-Traps 16384
    IP-Header 2.2.2.1:2.2.2.2:4:df:64:00000000 Encapsulation: IPv4-NUL
    Input packets : 0

```

```

Output packets: 0
Protocol inet, MTU: 1480
Flags: None

```

### show interfaces brief (IP-over-IP)

```

user@host> show interfaces ip-0/0/0 brief
Physical interface: ip-0/0/0, Enabled, Physical link is Up
Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
Device flags : Present Running
Interface flags: SNMP-Traps

Logical interface ip-0/0/0.0
Flags: Point-To-Point SNMP-Traps 16384
IP-Header 2.2.2.1:2.2.2.2:4:df:64:00000000 Encapsulation: IPv4-NULl
inet

```

### show interfaces detail (IP-over-IP)

```

user@host> show interfaces ip-0/0/0 detail
Physical interface: ip-0/0/0, Enabled, Physical link is Up
Interface index: 133, SNMP ifIndex: 27, Generation: 14
Type: IPIP, Link-level type: IP-over-IP, MTU: Unlimited, Speed: 800mbps
Hold-times : Up 0 ms, Down 0 ms
Device flags : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

Logical interface ip-0/0/0.0 (Index 69) (SNMP ifIndex 49) (Generation 9)
Flags: Point-To-Point SNMP-Traps 16384
IP-Header 2.2.2.1:2.2.2.2:4:df:64:00000000 Encapsulation: IPv4-NULl
Traffic statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Local statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Transit statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
Protocol inet, MTU: 1480, Generation: 13, Route table: 0
Flags: None

```

### show interfaces extensive (IP-over-IP)

The output for the show interfaces extensive command is identical to that for the show interfaces detail command. For sample output, see [show interfaces detail \(IP-over-IP\) on page 145](#).

## show interfaces (Logical Tunnel)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified logical tunnel interface.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b><i>lt-fpc/pic/port</i></b>.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces extensive (Logical Tunnel) on page 150</a>
<b>Output Fields</b>	<a href="#">Table 9 on page 146</a> lists the output fields for the <b>show interfaces</b> (logical tunnel) command. Output fields are listed in the approximate order in which they appear.

**Table 9: Logical Tunnel show interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

Table 9: Logical Tunnel show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Type</b>	Type of interface. <b>Software-Pseudo</b> indicates a standard software interface with no associated hardware device.	All levels
<b>Link-level type</b>	Encapsulation used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Clocking</b>	Reference clock source: <b>Internal</b> or <b>External</b> when configured. Otherwise, <b>Unspecified</b> .	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Link type</b>	Type of link.	All levels
<b>Link flags</b>	Information about the link. Possible values are described in the "Link Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Physical info</b>	Information about the physical interface.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Current address</b>	Configured MAC address.	<b>detail extensive none</b>
<b>Hardware address</b>	Hardware MAC address.	<b>detail extensive none</b>
<b>Alternate link address</b>	Backup link address.	<b>detail extensive none</b>
<b>Last flapped</b>	Date, time, and how long ago the interface went from down to up. The format is <b>Last flapped: year-month-day hour:minute:second timezone (hour:minute:second ago)</b> . For example, <b>Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)</b> .	<b>detail extensive none</b>
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li><b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface.</li> <li><b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface.</li> </ul>	<b>detail extensive</b>

Table 9: Logical Tunnel show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Input errors</b>	<p>Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Errors</b>—Sum of the incoming frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>Framing errors</b>—Number of packets received with an invalid frame checksum (FCS).</li> <li>• <b>Runts</b>—Number of frames received that are smaller than the runt threshold.</li> <li>• <b>Giants</b>—Number of frames received that are larger than the giant threshold.</li> <li>• <b>Policed discards</b>—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Output errors</b>	<p>Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious:</p> <ul style="list-style-type: none"> <li>• <b>Carrier transitions</b>—Number of times the interface has gone from <b>down</b> to <b>up</b>. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning.</li> <li>• <b>Errors</b>—Sum of the outgoing frame aborts and FCS errors.</li> <li>• <b>Drops</b>—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism.</li> <li>• <b>MTU errors</b>—Number of packets larger than the MTU threshold.</li> <li>• <b>Resource errors</b>—Sum of transmit drops.</li> </ul>	<b>extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels

Table 9: Logical Tunnel show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit 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.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Rate of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Rate of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Rate of packets received on the interface.</li> <li>• <b>Output packets</b>—Rate of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Local statistics</b>	Statistics for traffic received from and transmitted to the Routing Engine. 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.	<b>detail extensive</b>
<b>Transit statistics</b>	Statistics for traffic transiting the router. 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.	<b>detail extensive</b>
<b>Protocol</b>	Protocol family configured on the logical interface, such as <b>iso</b> , <b>inet6</b> , <b>mpls</b> .	<b>detail extensive none</b>
<b>MTU</b>	MTU size on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route table</b>	Route table in which this address exists. For example, <b>Route table:0</b> refers to <b>inet.0</b> .	<b>detail extensive</b>
<b>Flags</b>	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Addresses, Flags</b>	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>Destination</b>	IP address of the remote side of the connection.	<b>detail extensive none</b>
<b>Local</b>	IP address of the logical interface.	<b>detail extensive none</b>
<b>Broadcast</b>	Broadcast address of the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>

## Sample Output

### show interfaces extensive (Logical Tunnel)

```

user@host> show interfaces lt-1/0/0 extensive
Physical interface: lt-1/0/0, Enabled, Physical link is Up
  Interface index: 143, SNMP ifIndex: 70, Generation: 26
  Type: Logical-tunnel, Link-level type: Logical-tunnel, MTU: 0,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Unspecified
  Link flags     : None
  Physical info  : 13
  Hold-times    : Up 0 ms, Down 0 ms
  Current address: 00:90:69:a6:48:7e, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped  : 2004-03-03 15:53:52 PST (22:08:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:                0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0

Logical interface lt-1/0/0.0 (Index 66) (SNMP ifIndex 467) (Generation 3024)
  Flags: Point-To-Point SNMP-Traps 16384 DLCI 100 Encapsulation: FR-NLPID
  Traffic statistics:
    Input bytes :                0
    Output bytes :                0
    Input packets:                0
    Output packets:                0
  Local statistics:
    Input bytes :                0
    Output bytes :                0
    Input packets:                0
    Output packets:                0
  Transit statistics:
    Input bytes :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:                0                0 pps
  Protocol inet, MTU: 4470, Generation: 7034, Route table: 0
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.1.1/24, Local: 10.1.1.1, Broadcast: Unspecified,
    Generation: 2054

```

## show interfaces (Multicast Tunnel)

<b>Syntax</b>	<pre>show interfaces <i>interface-type</i> &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index <i>snmp-index</i>&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified multicast tunnel interface and its logical encapsulation and de-encapsulation interfaces.
<b>Options</b>	<p><b><i>interface-type</i></b>—On M Series and T Series routers, the interface type is <b><i>mt-fpc/pic/port</i></b>.</p> <p><b><i>brief   detail   extensive   terse</i></b>—(Optional) Display the specified level of output.</p> <p><b><i>descriptions</i></b>—(Optional) Display interface description strings.</p> <p><b><i>media</i></b>—(Optional) Display media-specific information about network interfaces.</p> <p><b><i>snmp-index snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b><i>statistics</i></b>—(Optional) Display static interface statistics.</p>
<b>Additional Information</b>	The multicast tunnel interface has two logical interfaces: encapsulation and de-encapsulation. These interfaces are automatically created by the Junos OS for every multicast-enabled VPN routing and forwarding (VRF) instance. The encapsulation interface carries multicast traffic traveling from the edge interface to the core interface. The de-encapsulation interface carries traffic coming from the core interface to the edge interface.
<b>Required Privilege Level</b>	view

**List of Sample Output** [show interfaces \(Multicast Tunnel\) on page 153](#)  
[show interfaces brief \(Multicast Tunnel\) on page 153](#)  
[show interfaces detail \(Multicast Tunnel\) on page 153](#)  
[show interfaces extensive \(Multicast Tunnel\) on page 153](#)  
[show interfaces \(Multicast Tunnel Encapsulation\) on page 155](#)  
[show interfaces \(Multicast Tunnel De-Encapsulation\) on page 155](#)

**Output Fields** Table 10 on page 152 lists the output fields for the **show interfaces** (Multicast Tunnel) command. Output fields are listed in the approximate order in which they appear.

**Table 10: Multicast Tunnel show interfaces Output Fields**

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Type of interface.	All levels
<b>Link-level type</b>	Encapsulation used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>

Table 10: Multicast Tunnel show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Traffic statistics</b>	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	All levels

## Sample Output

### show interfaces (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Interface index: 145, SNMP ifIndex: 41
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Device flags : Present Running
Interface flags: SNMP-Traps
Input rate   : 0 bps (0 pps)
Output rate  : 0 bps (0 pps)
```

### show interfaces brief (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 brief
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Device flags : Present Running
Interface flags: SNMP-Traps
```

### show interfaces detail (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 detail
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Interface index: 145, SNMP ifIndex: 41, Generation: 28
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Hold-times : Up 0 ms, Down 0 ms
Device flags : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
  Input bytes :          170664562          560000 bps
  Output bytes :          112345376          368176 bps
  Input packets:           2439107           1000 pps
  Output packets:          2439120           1000 pps
```

### show interfaces extensive (Multicast Tunnel)

```
user@host> show interfaces mt-1/2/0 extensive
Physical interface: mt-1/2/0, Enabled, Physical link is Up
Interface index: 141, SNMP ifIndex: 529, Generation: 144
Type: Multicast-GRE, Link-level type: GRE, MTU: Unlimited, Speed: 800mbps
Hold-times : Up 0 ms, Down 0 ms
Device flags : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
```

```

Traffic statistics:
Input bytes :          170664562          560000 bps
Output bytes :         112345376          368176 bps
Input packets:          2439107           1000 pps
Output packets:         2439120           1000 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0

```

Logical interface mt-1/2/0.32768 (Index 83) (SNMP ifIndex 556) (Generation 148)

Flags: Point-To-Point SNMP-Traps 0x4000 IP-Header  
 232.1.1.1:10.0.0.6:47:df:64:0000000800000000 Encapsulation: GRE=NULL

```

Traffic statistics:
Input bytes :          170418430
Output bytes :         112070294
Input packets:          2434549
Output packets:         2435593
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Local statistics:
Input bytes :              0
Output bytes :             80442
Input packets:              0
Output packets:            1031
Transit statistics:
Input bytes :          170418430          560000 bps
Output bytes :         111989852          368176 bps
Input packets:          2434549           1000 pps
Output packets:         2434562           1000 pps
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Protocol inet, MTU: 1572, Generation: 182, Route table: 4
Flags: None
Protocol inet6, MTU: 1572, Generation: 183, Route table: 4
Flags: None

```

Logical interface mt-1/2/0.1081344 (Index 84) (SNMP ifIndex 560) (Generation 149)

```

Flags: Point-To-Point SNMP-Traps 0x6000 Encapsulation: GRE=NULL
Traffic statistics:
Input bytes :          246132
Output bytes :         355524
Input packets:           4558
Output packets:          4558
IPv6 transit statistics:
Input bytes :              0
Output bytes :              0
Input packets:              0
Output packets:              0
Local statistics:
Input bytes :          246132
Output bytes :              0

```

```

Input packets:          4558
Output packets:         0
Transit statistics:
Input bytes :           0          0 bps
Output bytes :         355524      0 bps
Input packets:         0          0 pps
Output packets:        4558        0 pps
IPv6 transit statistics:
Input bytes :           0
Output bytes :           0
Input packets:         0
Output packets:        0
Protocol inet, MTU: Unlimited, Generation: 184, Route table: 4
Flags: None
Protocol inet6, MTU: Unlimited, Generation: 185, Route table: 4
Flags: None

```

#### show interfaces (Multicast Tunnel Encapsulation)

```

user@host> show interfaces mt-3/1/0.32768
Logical interface mt-3/1/0.32768 (Index 67) (SNMP ifIndex 0)
  Flags: Point-To-Point SNMP-Traps 0x4000
  IP-Header 239.1.1.1:10.255.70.15:47:df:64:0000000800000000
  Encapsulation: GRE-NULL
Input packets : 0
Output packets: 2
  Protocol inet, MTU: Unlimited
  Flags: None

```

#### show interfaces (Multicast Tunnel De-Encapsulation)

```

user@host> show interfaces mt-3/1/0.49152
Logical interface mt-3/1/0.49152 (Index 74) (SNMP ifIndex 0)
  Flags: Point-To-Point SNMP-Traps 0x6000 Encapsulation: GRE-NULL
Input packets : 0
Output packets: 2
  Protocol inet, MTU: Unlimited
  Flags: None

```

## show interfaces (PIM)

<b>Syntax</b>	show interfaces <i>interface-type</i> <brief   detail   extensive   terse> <descriptions> <media> <snmp-index <i>snmp-index</i> > <statistics>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified Protocol Independent Multicast (PIM) de-encapsulation or PIM encapsulation interface, respectively.
<b>Options</b>	<p><i>interface-type</i>—On M Series and T Series routers, the PIM de-encapsulation interface type is <b>pd-fpc/pic/port</b> and the PIM encapsulation interface type is <b>pe-fpc/pic/port</b>.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index <i>snmp-index</i></b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<p><a href="#">show interfaces (PIM De-Encapsulation) on page 157</a></p> <p><a href="#">show interfaces brief (PIM De-Encapsulation) on page 158</a></p> <p><a href="#">show interfaces detail (PIM De-Encapsulation) on page 158</a></p> <p><a href="#">show interfaces extensive (PIM Encapsulation) on page 158</a></p> <p><a href="#">show interfaces (PIM Encapsulation) on page 158</a></p> <p><a href="#">show interfaces brief (PIM Encapsulation) on page 158</a></p> <p><a href="#">show interfaces detail (PIM Encapsulation) on page 159</a></p> <p><a href="#">show interfaces extensive (PIM Encapsulation) on page 159</a></p>
<b>Output Fields</b>	Table 11 on page 156 lists the output fields for the <b>show interfaces</b> (PIM de-encapsulation or encapsulation) command. Output fields are listed in the approximate order in which they appear.

Table 11: PIM show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 11: PIM show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Enabled</b>	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive</b> none
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive</b> none
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Type of interface.	All levels
<b>Link-level type</b>	Encapsulation used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Device flags</b>	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Interface flags</b>	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>

## Sample Output

### show interfaces (PIM De-Encapsulation)

```

user@host> show interfaces pd-0/0/0
Physical interface: pd-0/0/0, Enabled, Physical link is Up
  Interface index: 130, SNMP ifIndex: 25
  Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
  Device flags   : Present Running

```

```
Interface flags: SNMP-Traps
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

#### show interfaces brief (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0 brief
Physical interface: pd-0/0/0, Enabled, Physical link is Up
Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
Device flags   : Present Running
Interface flags: SNMP-Traps
```

#### show interfaces detail (PIM De-Encapsulation)

```
user@host> show interfaces pd-0/0/0 detail
Physical interface: pd-0/0/0, Enabled, Physical link is Up
Interface index: 130, SNMP ifIndex: 25, Generation: 11
Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
Hold-times      : Up 0 ms, Down 0 ms
Device flags    : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
Input bytes      :                      0                      0 bps
Output bytes     :                      0                      0 bps
Input packets    :                      0                      0 pps
Output packets   :                      0                      0 pps
```

#### show interfaces extensive (PIM Encapsulation)

```
user@host> show interfaces pd-0/0/0 extensive
Physical interface: pd-0/0/0, Enabled, Physical link is Up
Interface index: 130, SNMP ifIndex: 25, Generation: 11
Type: PIMD, Link-level type: PIM-Decapsulator, MTU: Unlimited, Speed: 800mbps
Hold-times      : Up 0 ms, Down 0 ms
Device flags    : Present Running
Interface flags: SNMP-Traps
Statistics last cleared: Never
Traffic statistics:
Input bytes      :                      0                      0 bps
Output bytes     :                      0                      0 bps
Input packets    :                      0                      0 pps
Output packets   :                      0                      0 pps
```

#### show interfaces (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0
Physical interface: pe-0/0/0, Enabled, Physical link is Up
Interface index: 131, SNMP ifIndex: 26
Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
Device flags    : Present Running
Interface flags: SNMP-Traps
Input rate      : 0 bps (0 pps)
Output rate     : 0 bps (0 pps)
```

#### show interfaces brief (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 brief
Physical interface: pe-0/0/0, Enabled, Physical link is Up
Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
Device flags    : Present Running
Interface flags: SNMP-Traps
```

### show interfaces detail (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 detail
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26, Generation: 12
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
```

### show interfaces extensive (PIM Encapsulation)

```
user@host> show interfaces pe-0/0/0 extensive
Physical interface: pe-0/0/0, Enabled, Physical link is Up
  Interface index: 131, SNMP ifIndex: 26, Generation: 12
  Type: PIME, Link-level type: PIM-Encapsulator, MTU: Unlimited, Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Interface flags: SNMP-Traps
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes  :                0                0 bps
    Output bytes :                0                0 bps
    Input packets:                0                0 pps
    Output packets:              0                0 pps
```

## show interfaces (Virtual Loopback Tunnel)

<b>Syntax</b>	<pre>show interfaces vt-fpc/pic/port &lt;brief   detail   extensive   terse&gt; &lt;descriptions&gt; &lt;media&gt; &lt;snmp-index snmp-index&gt; &lt;statistics&gt;</pre>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display status information about the specified virtual loopback tunnel interface.
<b>Options</b>	<p><b>vt-fpc/pic/port</b>—Display standard information about the specified virtual loopback tunnel interface.</p> <p><b>brief   detail   extensive   terse</b>—(Optional) Display the specified level of output.</p> <p><b>descriptions</b>—(Optional) Display interface description strings.</p> <p><b>media</b>—(Optional) Display media-specific information about network interfaces.</p> <p><b>snmp-index snmp-index</b>—(Optional) Display information for the specified SNMP index of the interface.</p> <p><b>statistics</b>—(Optional) Display static interface statistics.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show interfaces (Virtual Loopback Tunnel) on page 162</a> <a href="#">show interfaces brief (Virtual Loopback Tunnel) on page 163</a> <a href="#">show interfaces detail (Virtual Loopback Tunnel) on page 163</a> <a href="#">show interfaces extensive (Virtual Loopback Tunnel) on page 163</a>
<b>Output Fields</b>	Table 12 on page 160 lists the output fields for the <b>show interfaces</b> (virtual loopback tunnel) command. Output fields are listed in the approximate order in which they appear.

Table 12: Virtual Loopback Tunnel show interfaces Output Fields

Field Name	Field Description	Level of Output
<b>Physical Interface</b>		
<b>Physical interface</b>	Name of the physical interface.	All levels
<b>Enabled</b>	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .	All levels
<b>Interface index</b>	Physical interface's index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	SNMP index number for the physical interface.	<b>detail extensive none</b>

Table 12: Virtual Loopback Tunnel show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Type</b>	Type of interface.	All levels
<b>Link-level type</b>	Encapsulation used on the physical interface.	All levels
<b>MTU</b>	MTU size on the physical interface.	All levels
<b>Speed</b>	Speed at which the interface is running.	All levels
<b>Hold-times</b>	Current interface hold-time up and hold-time down, in milliseconds.	<b>detail extensive</b>
<b>Device flags</b>	Information about the physical device. Possible values are described in the "Device Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Input Rate</b>	Input rate in bits per second (bps) and packets per second (pps).	None specified
<b>Output Rate</b>	Output rate in bps and pps.	None specified
<b>Statistics last cleared</b>	Time when the statistics for the interface were last set to zero.	<b>detail extensive</b>
<b>Traffic statistics</b>	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> <li>• <b>Input bytes, Output bytes</b>—Number of bytes received and transmitted on the interface.</li> <li>• <b>Input packets, Output packets</b>—Number of packets received and transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Logical Interface</b>		
<b>Logical interface</b>	Name of the logical interface.	All levels
<b>Index</b>	Logical interface index number, which reflects its initialization sequence.	<b>detail extensive none</b>
<b>SNMP ifIndex</b>	Logical interface SNMP interface index number.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Flags</b>	Information about the logical interface. Possible values are described in the "Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
<b>Encapsulation</b>	Encapsulation on the logical interface.	All levels
<b>Input packets</b>	Number of packets received on the logical interface.	None specified
<b>Output packets</b>	Number of packets transmitted on the logical interface.	None specified

Table 12: Virtual Loopback Tunnel show interfaces Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>Bandwidth</b>	Bandwidth allotted to the logical interface, in kilobytes per second.	All levels
<b>Traffic statistics</b>	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit 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.</p> <ul style="list-style-type: none"> <li>• <b>Input bytes</b>—Number of bytes received on the interface.</li> <li>• <b>Output bytes</b>—Number of bytes transmitted on the interface.</li> <li>• <b>Input packets</b>—Number of packets received on the interface.</li> <li>• <b>Output packets</b>—Number of packets transmitted on the interface.</li> </ul>	<b>detail extensive</b>
<b>Transit statistics</b>	Statistics for traffic transiting the router. 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.	<b>detail extensive</b>
<b><i>protocol-family</i></b>	Protocol family configured on the logical interface. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>brief</b>
<b>Protocol</b>	Protocol family configured on the logical interface. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>
<b>MTU</b>	Maximum transmission unit size on the logical interface.	<b>detail extensive none</b>
<b>Maximum labels</b>	Maximum number of MPLS labels configured for the MPLS protocol family on the logical interface.	<b>detail extensive none</b>
<b>Generation</b>	Unique number for use by Juniper Networks technical support only.	<b>detail extensive</b>
<b>Route Table</b>	Routing table in which the logical interface address is located. For example, 0 refers to the routing table inet.0.	<b>detail extensive</b>
<b>Flags</b>	Information about protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	<b>detail extensive none</b>

## Sample Output

### show interfaces (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Device flags      : Present Running
  Input rate       : 0 bps (0 pps)
  Output rate      : 0 bps (0 pps)

Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57)
  Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel

```

```

Input packets : 0
Output packets: 0
  Protocol inet, MTU: Unlimited
    Flags: None
  Protocol mpls, MTU: Unlimited, Maximum labels: 3
    Flags: None

```

#### show interfaces brief (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0 brief
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Device flags   : Present Running

Logical interface vt-1/2/0.0
  Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
  inet
  mpls

```

#### show interfaces detail (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0 detail
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40, Generation: 27
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms
  Device flags    : Present Running
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps

Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57) (Generation 17)
  Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel
  Traffic statistics:
    Input bytes   :                0
    Output bytes  :                0
    Input packets :                0
    Output packets:                0
  Transit statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps
  Protocol inet, MTU: Unlimited, Generation: 33, Route table: 0
    Flags: None
  Protocol mpls, MTU: Unlimited, Maximum labels: 3, Generation: 34, Route table:
0
    Flags: None

```

#### show interfaces extensive (Virtual Loopback Tunnel)

```

user@host> show interfaces vt-1/2/0 extensive
Physical interface: vt-1/2/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 40, Generation: 27
  Type: Loopback, Link-level type: Virtual-loopback-tunnel, MTU: Unlimited,
  Speed: 800mbps
  Hold-times      : Up 0 ms, Down 0 ms

```

Device flags : Present Running

Statistics last cleared: Never

Traffic statistics:

Input bytes :	0	0 bps
Output bytes :	0	0 bps
Input packets:	0	0 pps
Output packets:	0	0 pps

Logical interface vt-1/2/0.0 (Index 76) (SNMP ifIndex 57) (Generation 17)

Flags: Point-To-Point 16384 Encapsulation: Virtual-loopback-tunnel

Traffic statistics:

Input bytes :	0
Output bytes :	0
Input packets:	0
Output packets:	0

Transit statistics:

Input bytes :	0	0 bps
Output bytes :	0	0 bps
Input packets:	0	0 pps
Output packets:	0	0 pps

Protocol inet, MTU: Unlimited, Generation: 33, Route table: 0

Flags: None

Protocol mpls, MTU: Unlimited, Maximum labels: 3, Generation: 34, Route table:

0

Flags: None

## show ipsec certificates

<b>Syntax</b>	show ipsec certificates <brief   detail> <crl <i>crl-name</i>   <i>serial-number</i> >
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(Encryption interface on M Series and T Series routers only) Display information about the IPsec certificate database.
<b>Options</b>	<p><b>none</b>—Display standard information about all of the entries in the IPsec certificate database.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>crl <i>crl-name</i>   <i>serial-number</i></b>—(Optional) Display information about the entries on the certificate revocation list (CRL) or for the specified serial number. A CRL is a timestamped list identifying revoked certificates. The CRL is signed by a certificate authority (CA) or CRL issuer and made freely available in a public repository. Each revoked certificate is identified in a CRL by its certificate serial number.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li>• <a href="#">clear ipsec security-associations on page 115</a></li> </ul>
<b>List of Sample Output</b>	<a href="#">show ipsec certificates detail on page 166</a>
<b>Output Fields</b>	<a href="#">Table 13 on page 165</a> lists the output fields for the <b>show ipsec certificates</b> command. Output fields are listed in the approximate order in which they appear.

**Table 13: show ipsec certificates Output Fields**

Field Name	Field Description	Level of Output
<b>Database</b>	Display information about the IPsec certificate database. <ul style="list-style-type: none"> <li>• <b>Total entries</b>—Number of database entries, including entries that are not trusted or that are in the process of being deleted.</li> <li>• <b>Active entries</b>—Number of database entries, excluding entries that are marked as deleted.</li> <li>• <b>Locked entries</b>—Number of statically configured database entries that cannot expire, such as CA certificates that are root or trusted.</li> </ul>	All levels
<b>Subject</b>	Distinguished name for the certificate for <b>C, O, CN</b> , as described in RFC 3280, <i>Internet x.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile</i> .	All levels
<b>ID</b>	Identification number of the database entry. <b>ID</b> is generated by the internal certificate database.	All levels

Table 13: show ipsec certificates Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>References</b>	Reference number the certificate manager has for the particular entry.	<b>detail</b>
<b>Serial</b>	Unique serial number assigned to each certificate by the CA.	All levels
<b>Flags</b>	State of the certificate. <ul style="list-style-type: none"> <li>• <b>Trusted</b>—Passed validity checks.</li> <li>• <b>Not trusted</b>—Failed validity checks.</li> <li>• <b>Root</b>—Entry is locked and may have been learned through IKE or a locally configured CA certificate.</li> <li>• <b>Non-root</b>—Entry is not locked.</li> <li>• <b>Crl-issuer</b>—Entity issues CRLs.</li> <li>• <b>Non-crl-issuer</b>—Entity does not issue CRLs.</li> </ul>	<b>detail</b>
<b>Validity period starts</b>	Start time that the certificate is valid, in the format <i>yyyy mon dd, hh:mm:ss GMT</i> .	<b>detail</b>
<b>Validity period ends</b>	End time that the certificate is valid, in the format <i>yyyy mon dd, hh:mm:ss GMT</i> .	<b>detail</b>
<b>Alternative name information</b>	Auxiliary identity for the certificate: <i>dns-name</i> , <i>email-address</i> , <i>ip-address</i> , or <i>uri</i> (uniform resource identifier).	<b>detail</b>
<b>Issuer</b>	Information about the entity that has signed and issued the CRL as described in RFC 2459, <i>Internet X.509 Public Key Infrastructure Certificate and CRL Profile</i> .	<b>detail</b>

## Sample Output

### show ipsec certificates detail

```

user@host> show ipsec certificates detail
Database: Total entries: 3 Active entries: 4 Locked entries: 1
Subject: C=us, O=x
  ID: 5, References: 0, Serial: 22314868
  Flags: Trusted Non-root Crl-issuer
  Validity period starts: 2003 Mar 1st, 01:20:42 GMT
  Validity period ends: 2003 Mar 31st, 01:50:42 GMT
  Alternative name information:
    IP address: 10.20.210.1
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

Subject: C=us, O=x
  ID: 4, References: 0, Serial: 22315496
  Flags: Trusted Non-root Crl-issuer
  Validity period starts: 2003 Mar 1st, 01:21:45 GMT
  Validity period ends: 2003 Mar 31st, 01:51:45 GMT
  Alternative name information:
    IP address: 10.20.210.20
  Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

Subject: C=FI, O=SSH Company-ABC, CN=Company ABC class 2
  ID: 1, References: 1, Serial: 1538512
  Flags: Trusted Root Non-crl-issuer

```

Validity period starts: 2001 Aug 1st, 07:08:32 GMT  
Validity period ends: 2004 Aug 1st, 07:08:32 GMT  
Alternative name information:  
Email address: certifier-support@ssh.com  
Issuer: C=FI, O=Company-ABC, CN=Company ABC class 2

## show ipsec redundancy

<b>Syntax</b>	<code>show ipsec redundancy (interface &lt;es-fpc/pic/port&gt;   security-associations &lt;sa-name&gt;)</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	(Encryption interface on M Series and T Series routers only) Display information about IPsec redundancy.
<b>Options</b>	<p><b>interface &lt;es-fpc/pic/port&gt;</b>—Display information about all encryption interfaces, or optionally, about a particular encryption interface.</p> <p><b>security-associations &lt;sa-name&gt;</b>—Display information about all remote tunnels, or optionally, about a particular remote tunnel.</p>
<b>Required Privilege Level</b>	view
<b>Related Documentation</b>	<ul style="list-style-type: none"> <li><a href="#">request ipsec switch on page 117</a></li> </ul>
<b>List of Sample Output</b>	<p><a href="#">show ipsec redundancy interface on page 169</a></p> <p><a href="#">show ipsec redundancy security-associations on page 169</a></p>
<b>Output Fields</b>	Table 14 on page 168 lists the output fields for the <b>show ipsec redundancy</b> command. Output fields are listed in the approximate order in which they appear.

Table 14: show ipsec redundancy Output Fields

Field Name	Field Description
<b>Failure counter</b>	Number of times a PIC switched between primary and backup interfaces, or the number of times the tunnel switched between the primary and remote peers since the software has been activated.
<b>Primary interface '</b>	Name of the interface configured to be the primary interface.
<b>Backup interface</b>	Name of the interface configured to be the backup interface.
<b>State</b>	State of the primary or backup interface can be <b>Active</b> , <b>Offline</b> , or <b>Standby</b> . Both ES PICs are initialized to <b>Offline</b> . For primary and remote peers, <b>State</b> can be <b>Active</b> or <b>Standby</b> . Both peers are in a state of <b>Standby</b> by default (there is not yet a connection between the two peers).
<b>Security association</b>	Name of the security association.
<b>Local IP</b>	Local IP address.
<b>Primary remote IP</b>	IP address of the configured primary remote peer.
<b>Backup remote IP</b>	IP address of the configured backup remote peer.

## Sample Output

### show ipsec redundancy interface

```
user@host> show ipsec redundancy interface
Failure counter: 0
Primary interface: es-1/3/0, State: Active
Backup interface : es-1/1/0, State: Standby
```

### show ipsec redundancy security-associations

```
user@host> show ipsec redundancy security-associations sa-dynamic
Security association: sa-dynamic, Failure counter: 0
Local IP: 4.4.4.4
Primary remote IP: 4.4.4.5, State: Standby
Backup remote IP : 3.3.3.3, State: Standby
```

## show ipsec security-associations

<b>Syntax</b>	show ipsec security-associations <brief   detail> <sa-name>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4.
<b>Description</b>	Display information about the IPsec security associations applied to the local or transit traffic stream.
<b>Options</b>	<p><b>none</b>—Display standard information about all IPsec security associations.</p> <p><b>brief   detail</b>—(Optional) Display the specified level of output.</p> <p><b>sa-name</b>—(Optional) Display the specified IPsec security association.</p>
<b>Required Privilege Level</b>	view
<b>List of Sample Output</b>	<a href="#">show ipsec security-associations sa-name on page 172</a> <a href="#">show ipsec security-associations sa-name detail on page 172</a>
<b>Output Fields</b>	Table 15 on page 170 lists the output fields for the <b>show ipsec security-associations</b> command. Output fields are listed in the approximate order in which they appear.

Table 15: show ipsec security-associations Output Fields

Field Name	Field Description	Level of Output
<b>Security association</b>	Name of the security association.	All levels
<b>Interface family</b>	<p>Status of the interface family of the security association. If the interface family field is absent, it is a transport mode security association. The interface family can have one of three options:</p> <ul style="list-style-type: none"> <li>• <b>Up</b>—The security association is referenced in the interface family and the interface family is up.</li> <li>• <b>Down</b>—The security association is referenced in the interface family and the interface family is down.</li> <li>• <b>No reference</b>—The security association is not referenced in the interface family.</li> </ul>	All levels
<b>Local gateway</b>	Gateway address of the local system.	All levels
<b>Remote gateway</b>	Gateway address of the remote system.	All levels
<b>Local identity</b>	Prefix and port number of the local end	All levels
<b>Remote identity</b>	Prefix and port number of the remote end.	All levels
<b>Direction</b>	Direction of the security association: <b>inbound</b> or <b>outbound</b> .	All levels
<b>SPI</b>	Value of the security parameter index.	All levels

Table 15: show ipsec security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
<b>AUX-SPI</b>	Value of the auxiliary security parameter index. <ul style="list-style-type: none"> <li>When the value is <b>AH</b> or <b>ESP</b>, <b>AUX-SPI</b> is always <b>0</b>.</li> <li>When the value is <b>AH+ESP</b>, <b>AUX-SPI</b> is always a positive integer.</li> </ul>	All levels
<b>State</b>	Status of the security association: <ul style="list-style-type: none"> <li><b>Installed</b>—The security association is installed in the security association database. (For transport mode security associations, the value of <b>State</b> must always be <b>Installed</b>.)</li> <li><b>Not installed</b>—The security association is not installed in the security association database.</li> </ul>	<b>detail</b>
<b>Mode</b>	Mode of the security association: <ul style="list-style-type: none"> <li><b>transport</b>—Protects single host-to-host protections.</li> <li><b>tunnel</b>—Protects connections between security gateways.</li> </ul>	All levels
<b>Type</b>	Type of security association: <ul style="list-style-type: none"> <li><b>manual</b>—Security parameters require no negotiation. They are static, and are configured by the user.</li> <li><b>dynamic</b>—Security parameters are negotiated by the IKE protocol. Dynamic security associations are not supported in transport mode.</li> </ul>	All levels
<b>Protocol</b>	Protocol supported: <ul style="list-style-type: none"> <li><b>transport mode</b>—Supports Encapsulation Security Protocol (<b>ESP</b>) or Authentication Header (<b>AH</b>).</li> <li><b>tunnel mode</b>—Supports <b>ESP</b> or <b>AH+ESP</b>.</li> </ul>	All levels
<b>Authentication</b>	Type of authentication used: <b>hmac-md5-96</b> , <b>hmac-sha1-96</b> , or <b>None</b> .	<b>detail</b>
<b>Encryption</b>	Type of encryption used: <b>des-cbc</b> , <b>3des-csc</b> , or <b>None</b> .	<b>detail</b>
<b>Soft lifetime</b> <b>Hard lifetime</b>	( <b>dynamic</b> output only) Each lifetime of a security association has two display options, hard and soft, one of which must be present for a dynamic security association. The <b>hard lifetime</b> specifies the lifetime of the SA. The <b>soft lifetime</b> , which is derived from the hard lifetime, informs the IPsec key management system that the SA is about to expire. This allows the key management system to negotiate a new SA before the hard lifetime expires. <ul style="list-style-type: none"> <li><b>Expires in seconds seconds</b>—Number of seconds left until the security association expires.</li> <li><b>Expires in kilobytes kilobytes</b>—Number of kilobytes left until the security association expires.</li> </ul>	<b>detail</b>
<b>Anti-replay service</b>	State of the service that prevents packets from being replayed: <b>Enabled</b> or <b>Disabled</b> .	<b>detail</b>

Table 15: show ipsec security-associations Output Fields (*continued*)

Field Name	Field Description	Level of Output
Replay window size	Configured size, in packets, of the antireplay service window: 32 or 64. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets. If the replay window size is 0, the antireplay service is disabled.	detail

## Sample Output

### show ipsec security-associations sa-name

```

user@host> show ipsec security-associations sa-cosmic brief
Security association: sa-cosmic, Interface family: Up
Local gateway: 21.21.1.1, Remote gateway: 21.21.2.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction SPI      AUX-SPI    Mode      Type      Protocol
inbound  2908734119  0          tunnel    dynamic   AH
outbound 3494029335  0          tunnel    dynamic   AH

```

### show ipsec security-associations sa-name detail

```

user@host> show ipsec security-associations sa-cosmic detail
Security association: sa-cosmic, Interface family: Up

Local gateway: 21.21.1.1, Remote gateway: 21.21.2.1
Local identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Remote identity: ipv4_subnet(any:0,[0..7]=0.0.0.0/0)
Direction: inbound, SPI: 2908734119, AUX-SPI: 0, State: Installed
Mode: tunnel, Type: dynamic
Protocol: AH, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expired
Hard lifetime: Expires in 120 seconds
Anti-replay service: Disabled

Direction: outbound, SPI: 3494029335, AUX-SPI: 0, State: Installed
Mode: tunnel, Type: dynamic
Protocol: AH, Authentication: hmac-md5-96, Encryption: None
Soft lifetime: Expired
Hard lifetime: Expires in 120 seconds
Anti-replay service: Disabled

```

## show system certificate

<b>Syntax</b>	<code>show system certificate</code> <code>&lt;certificate-id&gt;</code>
<b>Release Information</b>	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
<b>Description</b>	(Encryption interface on M Series, T Series routers, QFX Series, and OCX Series switches only) Display installed certificates signed by the Juniper Networks certificate authority.
<b>Options</b>	<b>none</b> —Display all installed certificates signed by the Juniper Networks certificate authority.  <b>certificate-id</b> —(Optional) Display the details of a particular certificate.
<b>Required Privilege Level</b>	maintenance
<b>List of Sample Output</b>	<a href="#">show system certificate on page 174</a> <a href="#">show system certificate (QFX Series) on page 174</a>
<b>Output Fields</b>	<a href="#">Table 16 on page 173</a> lists the output fields for the <b>show system certificate</b> command. Output fields are listed in the approximate order in which they appear.

**Table 16: show system certificate Output Fields**

Field Name	Field Description
<b>Certificate identifier</b>	Unique identifier associated with a certificate. The certificate identifier is the common name of the subject.
<b>Issuer</b> <b>Subject</b>	Information about the certificate issuer and the distinguished name (DN) of the issuer, respectively: <ul style="list-style-type: none"> <li>• <b>Organization</b>—Name of the owner's organization.</li> <li>• <b>Organizational unit</b>—Name of the owner's department.</li> <li>• <b>Country</b>—Two-character country code in which the owner's system is located.</li> <li>• <b>State</b>—State in the USA in which the owner is using the certificate.</li> <li>• <b>Locality</b>—City in which the owner's system is located.</li> <li>• <b>Common name</b>—Name of the owner of the certificate.</li> <li>• <b>E-mail address</b>—E-mail address of the owner of the certificate.</li> </ul>
<b>Validity</b>	When a certificate is valid.
<b>Signature algorithm</b>	Encryption algorithm applied to the installed certificate.
<b>Public key algorithm</b>	Encryption algorithm applied to the public key.

## Sample Output

### show system certificate

```
user@host> show system certificate
Certificate identifier: Dallas-v3
  Issuer:
    Organization: Juniper Networks, Organizational unit: Juniper CA,
    Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
    E-mail address:ca@device.net
  Subject:
    Organization: Juniper Networks, Organizational unit: Juniper CA,
    Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
    E-mail address:ca@device.net
  Validity:
    Not before: Mar 13 03:23:25 2004 GMT
    Not after: Mar 24 03:23:25 2014 GMT
  Signature algorithm: sha1WithRSAEncryption
  Public key algorithm: dsaEncryption
```

### show system certificate (QFX Series)

```
user@host> show system certificate
Certificate identifier: Dallas-v3
  Issuer:
    Organization: Juniper Networks, Organizational unit: Juniper CA,
    Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas CA,
    E-mail address:ca@device.net
  Subject:
    Organization: Juniper Networks, Organizational unit: Juniper CA,
    Country: US, State: CA, Locality: Sunnyvale, Common name: Dallas-v3,
    E-mail address:ca@device.net
  Validity:
    Not before: Mar 13 03:23:25 2004 GMT
    Not after: Mar 24 03:23:25 2014 GMT
  Signature algorithm: sha1WithRSAEncryption
  Public key algorithm: dsaEncryption
```

## PART 4

# Index

- [Index on page 177](#)



# Index

## Symbols

#, comments in configuration statements.....	xiv
( ), in syntax descriptions.....	xiv
< >, in syntax descriptions.....	xiv
[ ], in configuration statements.....	xiv
{ }, in configuration statements.....	xiv
(pipe), in syntax descriptions.....	xiv

## A

address statement	
encryption.....	84
usage guidelines.....	69
allow-fragmentation statement.....	85
usage guidelines.....	21
apply-groups-except statement.....	85

## B

backup-destination statement.....	86
usage guidelines.....	78
backup-interface statement.....	86
usage guidelines.....	77
braces, in configuration statements.....	xiv
brackets	
angle, in syntax descriptions.....	xiv
square, in configuration statements.....	xiv

## C

certificates	
installed, displaying.....	173
key pairs, generating.....	121
provided by Juniper Networks, adding.....	122
signed certificate, obtaining.....	118
unsigned certificate, obtaining.....	120
clear ike security-associations command.....	114
clear ipsec security-associations command.....	115
clear-dont-fragment-bit statement	
usage guidelines.....	12, 20
comments, in configuration statements.....	xiv
conventions	
text and syntax.....	xiii

copy-tos-to-outer-ip-header statement.....	87
usage guidelines.....	21
core-facing statement.....	87
CoS	
for tunnels	
GRE TOS bits.....	21
curly braces, in configuration statements.....	xiv
customer support.....	xv
contacting JTAC.....	xv

## D

data-fill-with-zeros statement.....	95
destination statement.....	88
encryption	
usage guidelines.....	69, 78
tunnel	
usage guidelines.....	17, 59
destination-networks statement	
tunnel.....	90
usage guidelines.....	63
digital certificates See certificates	
do-not-fragment statement	
tunnel.....	91
usage guidelines.....	21
documentation	
comments on.....	xv
dynamic tunnels.....	92
destination.....	90
source.....	106
dynamic-tunnels statement.....	92
usage guidelines.....	63

## E

encryption interface.....	69
applying inbound filter.....	75
example configuration.....	75
applying outbound filter.....	74
example configuration.....	73, 74
configuring inbound filter.....	74
example configuration.....	75
configuring MTU.....	70
encryption interfaces	
status information, displaying.....	127
ES interfaces	
example configuration.....	70
ES PIC	
apply inbound filter.....	75
PIC redundancy.....	77

redundancy		IPsec	
example configuration.....	77	ES PIC.....	69
tunnel redundancy.....	78	example configuration	
es-options statement.....	93	inbound traffic.....	75
usage guidelines.....	77	outbound traffic.....	73
		traffic.....	71
<b>F</b>		IPsec services	
family bridge statement.....	96	encryption services interfaces	
family statement		backup and primary, switching	
encryption.....	94	interfaces.....	117
usage guidelines.....	69	backup and primary, switching	
filter statement		services.....	117
encryption.....	97	certificate database, displaying.....	165
usage guidelines.....	75	IKE security associations, clearing.....	114
font conventions.....	xiii	IKE security associations, displaying.....	123
fragmentation		IPSec security associations, clearing.....	115
GRE tunnels.....	20	IPSec security associations,	
		displaying.....	170
<b>G</b>		redundancy information, displaying.....	168
GRE interfaces		ipsec-sa statement	
status information, displaying.....	133	encryption.....	98
GRE tunnels		usage guidelines.....	69
fragmentation.....	20	IPv6	
key number.....	19	transition	
		configured tunnel.....	15
<b>H</b>		IPv6-over-IPv4 tunnel	
hold-time statement		example configuration.....	15
GRE tunnel interface.....	97	standards supported.....	15
<b>I</b>		<b>K</b>	
IKE		keepalive-time statement	
encryption services interfaces		GRE tunnel interface.....	99
security associations, clearing.....	114	key pair for digital certificate, generating.....	121
security associations, displaying.....	123	key statement	
interface statement		tunnel.....	100
encryption		usage guidelines.....	19
usage guidelines.....	69	<b>L</b>	
interfaces statement		logical tunnel interfaces	
encryption.....	98	status information, displaying.....	146
usage guidelines.....	69	logical tunnels.....	25
tunnel		example configuration.....	26
usage guidelines.....	17	loopback tunnels.....	57
IP-over-IP interfaces		<b>M</b>	
status information, displaying.....	142	manuals	
		comments on.....	xv
		multicast tunnel interfaces	
		status information, displaying.....	151

multicast tunnels.....	23
multicast-only statement.....	100
usage guidelines.....	23

## P

parentheses, in syntax descriptions.....	xiv
peer-certificate-type statement.....	101
peer-unit statement	
tunnel.....	101
usage guidelines.....	25
PIM	
tunnels.....	55
PIM de-encapsulation interfaces	
status information, displaying.....	156
PIM encapsulation interfaces	
status information, displaying.....	156

## R

reassemble-packets statement.....	102
usage guidelines.....	21
redundancy group	
logical interfaces in.....	102
redundancy-group	
logical tunnel.....	103
redundant logical tunnels	
overview.....	28
redundant logical tunnels	
configuring.....	31
request ipsec switch command.....	117
request security certificate (signed) command.....	118
request security certificate (unsigned)	
command.....	120
request security key-pair command.....	121
request system certificate add command.....	122
RFC 2890.....	19
routing-instance statement	
tunnel.....	104
usage guidelines.....	59
routing-instances statement.....	104
routing-options statement.....	105

## S

security certificate See certificates	
show ike security-associations command.....	123
show interfaces (Encryption) command.....	127
show interfaces (GRE) command.....	133
show interfaces (IP-over-IP) command.....	142
show interfaces (Logical Tunnel) command.....	146
show interfaces (Multicast Tunnel) command.....	151

show interfaces (PIM) command.....	156
show interfaces (Virtual Loopback Tunnel)	
command.....	160
show ipsec certificates command.....	165
show ipsec redundancy command.....	168
show ipsec security-associations command.....	170
show system certificate command.....	173
show system statistics command.....	122
source statement	
encryption.....	105
tunnel.....	106
usage guidelines.....	78
source-address statement	
tunnel.....	106
tunnel services	
usage guidelines.....	63
support, technical See technical support	
syntax conventions.....	xiii

## T

technical support	
contacting JTAC.....	xv
traffic.....	71
inbound (decryption).....	75
IPsec, configuring.....	71
outbound (encryption).....	73
ttl statement	
tunnel.....	107
tunnel interfaces	
configuration statements.....	17, 23, 57
dynamic tunnels.....	63
example configuration.....	22
logical tunnels.....	25
loopback tunnels.....	57
multicast tunnels.....	23
PIM tunnels.....	55
unicast tunnels.....	17
tunnel services interfaces.....	164
tunnel statement.....	109
encryption.....	108
usage guidelines.....	69
redundancy	
usage guidelines.....	78
unicast	
usage guidelines.....	17

tunnels	
definition.....	3
GRE	
fragmentation of.....	20
key number.....	19
interface types.....	3
IPv6-over-IPv4.....	15
<b>U</b>	
unicast tunnels.....	17
unit statement	
encryption.....	110
usage guidelines.....	69
tunnel.....	111
usage guidelines.....	17
<b>V</b>	
virtual loopback tunnel	
configuration guidelines.....	59
VRF table lookup	
example configuration.....	59
virtual loopback tunnel interfaces	
status information, displaying.....	160