JUNIPEC
NETWORKS

Junos® OS

Logical Systems Feature Guide for Security Devices

Release

15.1X49-D70

Modified: 2016-11-21

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® OS Logical Systems Feature Guide for Security Devices
15.1X49-D70
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

Table of Contents

List of Figures

List of Tables

About the Documentation

Documentation and Release Notes

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

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

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

Supported Platforms

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

- SRX5600
- SRX5800
- SRX5400

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xv 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 xvi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	user@host> **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric *metric*>;
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast | multicast (*string1* | *string2* | *string3*)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [*community-ids*]

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

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

GUI Conventions

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

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

Documentation Feedback

We encourage you to provide feedback, 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. 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.

Overview

CHAPTER 1

Introduction to Logical Systems

Understanding Logical Systems for SRX Series Services Gateways

Logical systems for SRX Series devices enable you to partition a single device into secure contexts. Each logical system has its own discrete administrative domain, logical interfaces, routing instances, security firewall and other security features. By transforming an SRX Series device into a multitenant logical systems device, you can give various departments, organizations, customers, and partners—depending on your environment—private use of portions of its resources and a private view of the device. Using logical systems, you can share system and underlying physical machine resources among discrete user logical systems and the master logical system.

The top part of shows the three main configuration components of a logical system. The lower part of the figure shows a single device with a master logical system and discrete user logical systems.

Figure 1: Understanding Logical Systems

Logical system components

SRX Series device running logical systems
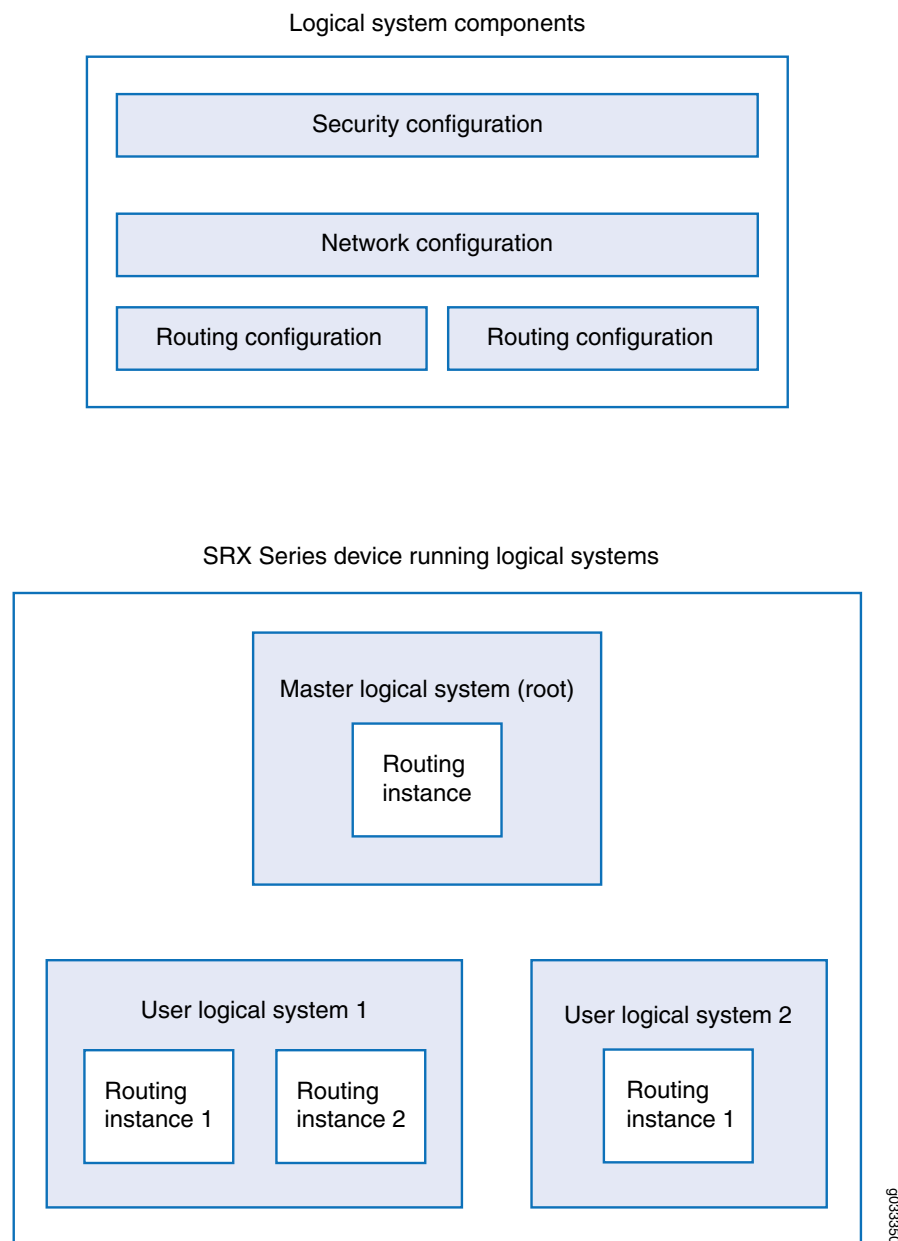
Logical systems on SRX Series devices offer many benefits, allowing you to:

- Curtail costs. Using logical systems, you can reduce the number of physical devices required for your company. Because you can consolidate services for various groups of users on a single device, you reduce both hardware costs and power expenditure.

- Create many logical systems on a single device and provision resources and services for them quickly. Because services are converged, it is easier for the master, or root, administrator to manage a single device configured for logical systems than it is to manage many discrete devices.

You can deploy an SRX Series device running logical systems in many environments, in particular, in the enterprise and in the data center.

- In the enterprise, you can create and provision logical systems for various departments and groups.

  You can configure logical systems to enable communication among groups sharing the device. When you create logical systems for various departments on the same device, users can communicate with one another without traffic leaving the device if you have configured an interconnect logical system to serve as an internal switch. For example, members of the product design group, the marketing department, and the accounting department sharing an SRX Series Services Gateway running logical systems can communicate with one another just as they could if separate devices were deployed for their departments. You can configure logical systems to interconnect through *logical tunnel (lt-0/0/0)* internal interfaces. The lt-0/0/0 interfaces on the interconnect logical system connect to an lt-0/0/0 interface that you configure for each logical system. The interconnect logical system switches traffic between logical systems. The SRX Series device running logical systems provides for high, fast interaction among all logical systems created on the device when an interconnect logical system is used.

  Logical systems on the same device can also communicate with one another directly through ports on the device, as if they were separate devices. Although this method allows for direct connections between logical systems, it consumes more resources—you must configure interfaces and an external switch—and therefore it is more costly.

- In the data center, as a service provider, you can deploy an SRX Series device running logical systems to offer your customers secure and private user logical systems and discrete use of the device's resources.

  For example, one corporation might require 10 user logical systems and another might require 20. Because logical systems are secure, private, and self-contained, data belonging to one logical system cannot be viewed by administrators or users of other logical systems. That is, employees of one corporation cannot view the logical systems of another corporation.

Logical systems include both master and user logical systems and their administrators. The roles and responsibilities of the master administrator and those of a user logical system administrator differ greatly. This differentiation of privileges and responsibilities is considered role-based administration and control.

> NOTE: To use the internal switch, which is optional, you must also configure an interconnect logical system. The interconnect logical system does not require an administrator.

Related Documentation

- Understanding the Master Logical System and the Master Administrator Role on page 19
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Understanding the Fundamentals and Constraints of Logical Systems

This topic covers basic information about logical systems features and limitations.

- By default, logical systems delivers a master logical system, which exists at the root level. You can purchase licenses for logical systems that you intend to create with the total not exceeding 32.

- You can configure up to 32 security profiles, from 1 through 32, with ID 0 reserved for the internally configured default security profile. When the maximum number of security profiles is reached, if you want to add a new security profile, you must first delete one or more existing security profiles, commit the configuration, and then create the new security profile and commit it. You cannot add a new security profile and remove an existing one within a single configuration commit.

  If you want to add more than one new security profile, the same rule is true. You must first delete the equivalent number of of existing security profiles, commit the configuration, and then create the new security profiles and commit them.

- You can configure one or more master administrators to oversee administration of the device and the logical systems they configure.

  As master administrator for an SRX Series Services Gateway running logical systems, you have root control over the device, its resources, and the logical systems that you create. You allocate security, networking, and routing resources to user logical systems. You can configure one logical system to serve as an interconnect logical system virtual private LAN service (VPLS) switch. The interconnect logical system, which is not mandatory, does not require security resources. However, if you configure an interconnect logical system, you must bind a dummy security profile to it. The master administrator configures it and all lt-0/0/0 interfaces for it.

- A user logical system can have one or more administrators, referred to as user logical system administrators. The master administrator creates login accounts for these administrators and assigns them to a user logical system. Currently, the master administrator must configure all user logical system administrators. The first assigned user logical administrator cannot configure additional user logical system administrators for his logical system. As a user logical system administrator, you can configure the resources assigned to your user logical system, including logical interfaces assigned by the master administrator, routing instances and their routes, and security components. You can display configuration information only for your logical system.

- A logical system can include more than one routing instance based on available system resources.

- You cannot configure class of service on lt-0/0/0 interfaces.

- The trace and debug features are supported at the root level only.

- Commit rollback is supported at the root level only.

- Quality of service (QoS) classification across interconnected logical systems does not work.

- The master administrator can configure Application Layer Gateways (ALGs) at the root level. The configuration is inherited by all user logical systems. It cannot be configured discretely for user logical systems.

- The master administrator can configure IDP policies at the root level and then apply an IDP policy to a user logical system.

- Only the master administrator can create user accounts and login IDs for users for all logical systems. The master administrator creates these user accounts at the root level and assigns them to the appropriate user logical systems.

- The same name cannot be used in two separate logical systems. For example, if logical-system1 includes a user with Bob configured as the username, then other logical systems on the device cannot include a user with the username Bob.

- Configuration for users for all logical systems and all user logical systems administrators must be performed at the root level by the master administrator. A user logical system administrator cannot create other user logical system administrators or user accounts for their logical systems.

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3
- Understanding the Master Logical System and the Master Administrator Role on page 19
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Understanding Licenses for Logical Systems on SRX Series Devices

This topic provides licensing information for SRX Series devices running logical systems. For general licensing information, such as how to install a license, see the Junos OS Software Installation and Upgrade Guide.

By default, a device running logical systems delivers a master logical system at the root level. You can purchase licenses for other logical systems that you intend to create. If you intend to configure an interconnect logical system to use as a switch, it also requires a license.

Complications arise if the number of logical systems that you configure exceeds the number of licenses that you have purchased. The system will allow you to configure additional logical systems. However, when you attempt to commit their configurations, the system issues a warning message similar to the following: **Warning: 2 more license(s) are needed, logical system won't work without license!**. The message indicates the number of logical systems without licenses. We recommend that you do not configure more logical systems than the number of licenses you have purchased.

If you configure more logical systems than the number of licenses that you have purchased, the additional logical systems will not be activated until a license is available. The system will drop packets destined to them. They are inactive.

When a logical system is deleted, its license is freed up. That license is assigned to an inactive logical system, and the logical system is activated.

You can use the **show system license status logical-system all** command on the command-line interface (CLI) to determine which logical systems are active.

`user@host>` **show system license status logical-system all**

```
logical system name      license status
root-logical-system      enabled
LSYS2             enabled
LSYS0             enabled
LSYS11            enabled
LSYS12            enabled
LSYS23            enabled
LSYS10            enabled
LSYS13            enabled
LSYS18            enabled
```

When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Related
Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3

- Understanding the Master Logical System and the Master Administrator Role on page 19

- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Understanding the Interconnect Logical System and Logical Tunnel Interfaces

This topic covers the interconnect logical system that serves as an internal virtual private LAN service (VPLS) switch connecting one logical system on the device to another. The topic also explains how logical tunnel (lt-0/0/0) interfaces are used to connect logical systems through the interconnect logical system.

A device running logical systems can use an internal VPLS switch to pass traffic without it leaving the device. The interconnect logical system switches traffic across logical systems that use it. Although a virtual switch is used typically, it is not mandatory. If you choose to use a virtual switch, you must configure the interconnect logical system. There can be only one interconnect logical system on a device.

For communication between logical systems on the device to occur, you must configure an lt-0/0/0 interface on each logical system that will use the internal switch, and you must associate it with its peer lt-0/0/0 interface on the interconnect logical system, effectively creating a logical tunnel between them. You define a peer relationship at each end of the tunnel when you configure the logical system's lt-0/0/0 interfaces.

You might want all logical systems on the device to be able to communicate with one another without using an external switch. Alternatively, you might want some logical systems to connect across the internal switch but not all of them.

The interconnect logical system does not require security resources assigned to it through a security profile. However, you must assign a dummy security profile containing no resources to the interconnect logical system. Otherwise you will not be able to successfully commit the configuration for it.

> **WARNING:** If you configure an lt-0/0/0 interface in any user logical system or the master logical system and you do not configure an interconnect logical system containing a peer lt-0/0/0 interface for it, the commit will fail.

An SRX Series device running logical systems can be used in a chassis cluster. Each node has the same configuration, including the interconnect logical system.

When you use SRX Series devices running logical systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Related Documentation

- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 140

- Understanding Logical Systems for SRX Series Services Gateways on page 3

- Understanding Logical Systems in the Context of Chassis Cluster on page 167

Understanding Flow in Logical Systems for SRX Series Devices

This topic explains how packets are processed in flow sessions on SRX Series devices running logical systems. It describes how an SRX Series device running logical systems handles pass-through traffic in a single logical system and between logical systems. It also covers self-traffic as self-initiated traffic within a logical system and self-traffic terminated on another logical system. Before addressing logical systems, the topic provides basic information about the SRX Series architecture in with respect to packet processing and sessions. Finally, it addresses sessions and how to change session characteristics.

The concepts explained in this example rely on the topology shown in Figure 2 on page 10.

Figure 2: Logical Systems, Their Virtual Routers, and Their Interfaces
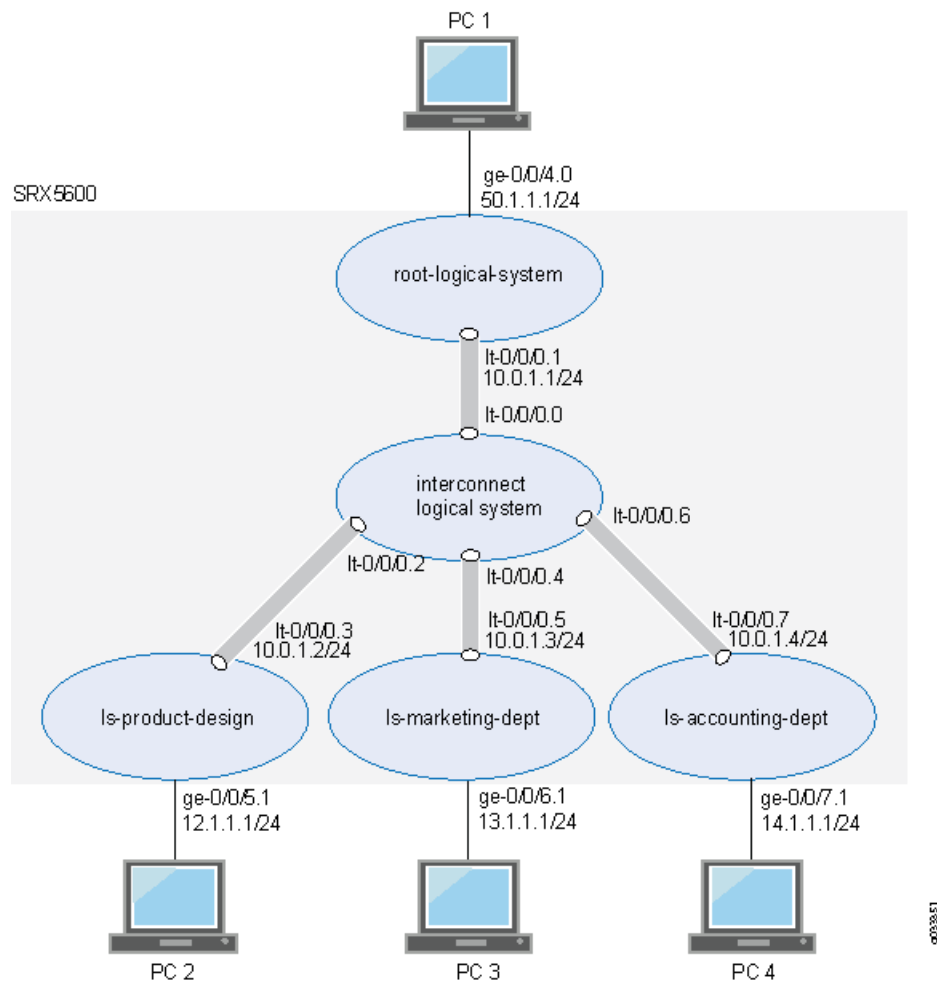
- Understanding Junos OS SRX Series Services Gateways Architecture on page 10
- Session Creation for Devices Running Logical Systems on page 12
- Understanding Flow on Logical Systems on page 12
- Understanding Packet Classification on page 12
- Handling Pass-Through Traffic for Logical Systems on page 13
- Handling Self-Traffic on page 14
- Understanding Session and Gate Limitation Control on page 16
- Understanding Sessions on page 16
- About Configuring Sessions on page 16

Understanding Junos OS SRX Series Services Gateways Architecture

Junos OS is a distributed parallel processing high throughput and high performance system. The distributed parallel processing architecture of the services gateways includes multiple processors to manage sessions and run security and other services processing.

This architecture provides greater flexibility and allows for high throughput and fast performance.

The SRX5000 line devices include I/O cards (IOC) and Services Processing Cards (SPCs) that each contain processing units that process a packet as it traverses the device. A Network Processing Unit (NPU) runs on an IOC. An IOC has one or more NPUs. One or more Services Processing Units (SPUs) run on an SPC.

These processing units have different responsibilities. All flow-based services for a packet are executed on a single SPU. Otherwise, however, the lines are not clearly divided in regard to the kinds of services that run on these processors. (For details on flow-based processing, see *Juniper Networks Devices Processing Overview*.)

For example:

- An NPU processes packets discretely. It performs sanity checks and applies some screens that are configured for the interface, such as denial-of-service (DoS) screens, to the packet.

- An SPU manages the session for the packet flow and applies security features and other services to the packet. It also applies packet-based stateless firewall filters, classifiers, and traffic shapers to the packet.

- The system uses one processor as a central point to take care of arbitration and allocation of resources and distribute sessions in an intelligent way. The central point assigns an SPU to be used for a particular session when the first packet of its flow is processed.

These discrete, cooperating parts of the system, including the central point, each store the information identifying whether a session exists for a stream of packets and the information against which a packet is matched to determine if it belongs to an existing session.

This architecture allows the device to distribute processing of all sessions across multiple SPUs. It also allows an NPU to determine if a session exists for a packet, to check the packet, and to apply screens to it. How a packet is handled depends on whether it is the first packet of a flow.

Flow-based packet processing treats related packets, or a stream of packets, in the same way. Packet treatment depends on characteristics that are established for the first packet of the packet stream when the flow session is established. Most packet processing occurs within a flow. For the distributed processing architecture of the services gateway, some packet-based processing, such as traffic shaping, occurs on the NPU. Some packet-based processing, such as application of classifiers to a packet, occurs on the SPU.

Configuration settings that determine the fate of a packet—such as the security policy that applies to it, Aplication Layer Gateway (ALG)s configured for it, if NAT should be applied to translate the packet's source and/or destination IP address—are assessed for the first packet of a flow.

Session Creation for Devices Running Logical Systems

Session establishment for SRX Series devices running logical systems differs in minor ways from that of SRX series devices not running logical systems. Despite the complexities that logical systems introduce, traffic is handled in a manner similar to how it is handled on SRX Series devices not running logical systems. Flow-based packet processing, which is stateful, requires the creation of sessions. In considering flow based processing and session establishment for logical systems, it helps to think of each logical system on the device as a discrete device with respect to session establishment.

A session is created, based on routing and other classification information, to store information and allocate resources for a flow. Basically, a session is established when traffic enters a logical system interface, route lookup is performed to identify the next hop interface, and policy lookup is performed.

Optionally, logical systems enable you to configure an internal software switch. This virtual private LAN switch (VPLS) is implemented as an interconnect logical system. It enables both transit traffic and traffic terminated at a logical system to pass between logical systems. To enable traffic to pass between logical systems, logical tunnel (lt-0/0/0) interfaces across the interconnect logical system are used.

Communication between logical systems across the interconnect logical system requires establishment of two sessions: one for traffic that enters a logical system and exits its lt-0/0/0 interface, and one for traffic that enters the lt-0/0/0 interface of another logical system and either exits the device through one of its physical interface or is destined for it.

> **NOTE:** Packet sequence occurs at the ingress and the egress interfaces. Packets traveling between logical systems might not be processed in the order in which they were received on the physical interface.

Understanding Flow on Logical Systems

To understand how traffic is handled for logical systems, it is helpful to consider each logical system as a discrete device.

> **NOTE:** Traffic is processed for the master logical system in the same way as it is for user logical systems on the device.

Understanding Packet Classification

Packet classification is assessed the same way for SRX Series devices running with or without logical systems. Filters and class-of-service features are typically associated with an interface to influence which packets are allowed to transit the system and to apply special actions to packets as needed. (Within a flow, some packet-based processing also takes place on an SPU.)

Packet classification is based on the incoming interface and performed at the ingress point. Traffic for a dedicated interface is classified to the logical system that contains that interface. Within the context of a flow, packet classification is based on both the physical interface and the logical interface.

Handling Pass-Through Traffic for Logical Systems

For SRX Series devices not running logical systems, pass-through traffic is traffic that enters and exits a device. You can think of pass-through traffic for logical systems similarly, but as having a larger dimension as a result of the nature of a multitenant device. For SRX Series devices running logical systems, pass-through traffic can exist within a logical system or between logical systems.

- Pass-Through Traffic Within a Logical System on page 13
- Pass-Through Traffic Between Logical Systems on page 13

Pass-Through Traffic Within a Logical System

For pass-through traffic within a logical system, traffic comes in on an interface belonging to one of the logical system's virtual routing instances, and it is sent to another of its virtual routing instances. To exit the device, the traffic is sent out an interface belonging to the second virtual routing instance. The traffic does not transit between logical systems but rather enters and exits the device in a single logical system. Pass-through traffic within a logical system is transmitted according to the routing tables in each of its routing instances.

Consider how pass-through traffic is handled within a logical system given the topology shown in Figure 2 on page 10.

- When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.

- Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1 with pd-vr2 identified as the next hop.

- A second route lookup is performed in pd-vr2 to identify the egress interface to use—in this case— ge-0/0/8.

- The packet is sent out ge-0/0/8 to the network.

- The security policy lookup is performed in ls-product-design, and one session is established.

Pass-Through Traffic Between Logical Systems

Pass-through traffic between logical systems is complicated by fact that each logical system has an ingress and an egress interface that the traffic must transit. It is as if traffic were coming into and going out from two devices.

Two sessions must be established for pass-through traffic between logical systems. (Note that policy lookup is performed in both logical systems).

- On the incoming logical system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).

- On the egress logical system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second logical system) and its egress interface (a physical interface).

Consider how pass-through traffic is handled across logical systems in the topology shown in Figure 2 on page 10.

- A session is established in the incoming logical system.

  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.

  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.

  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, which is the ingress interface in the ls-marketing-dept.

  - A session is established between ge-0/0/5 and lt-0/0/0.3.

- A session is established in the outgoing logical system.

  - The packet is injected into the flow again from lt-0/0/0.5, and the logical system context identified as ls-marketing-dept is derived from the interface.

  - Packet processing continues in the ls-marketing-dept logical system.

  - To identify the egress interface, route lookup for the packet is performed in the mk-vr1 routing instances.

  - The outgoing interface is identified as ge-0/0/6, and the packet is transmitted from the interface to the network.

Handling Self-Traffic

Self-traffic is traffic that originates in a logical system on the device and is either sent out to the network from that logical system or is terminated on another logical system on the device.

Self-Initiated Traffic

Self-initiated traffic is generated from a source logical system context and forwarded directly to the network from the logical system interface.

The following process occurs:

- When a packet is generated in a logical system, a process for handling the traffic is started in the logical system.

- Route lookup is performed to identify the egress interface, and a session is established.

- The logical system performs a policy lookup and processes the traffic accordingly.

- If required, a management session is set up.

Consider how self-initiated traffic is handled across logical systems given the topology shown in Figure 2 on page 10.

- A packet is generated in the ls-product-design logical system, and a process for handling the traffic is started in the logical system.

- Route lookup performed in pd-vr2 to identifies the egress interface as ge-0/0/8.

- A session is established.

- The packet is transmitted to the network from ge-0/0/8.

Traffic Terminated on a Logical System

When a packet enters the device on an interface belonging to a logical system and the packet is destined for another logical system on the device, the packet is forwarded between the logical systems in the same manner as is pass-through traffic. However, route lookup in the second logical system identifies the local egress interface as the packet destination. Consequently the packet is terminated on the second logical system as self-traffic.

- For terminated self-traffic, two policy lookups are performed, and two sessions are established.

  - On the incoming logical system, one session is set up between the ingress interface (a physical interface) and its egress interface (an lt-0/0/0 interface).

  - On the destination logical system, another session is set up between the ingress interface (the lt-0/0/0 interface of the second logical system) and the local interface.

Consider how terminated self-traffic is handled across logical systems in the topology shown in Figure 2 on page 10.

- A session is established in the incoming logical system.

  - When a packet arrives on interface ge-0/0/5, it is identified as belonging to the ls-product-design logical system.

  - Because ge-0/0/5 belongs to the pd-vr1 routing instance, route lookup is performed in pd-vr1.

  - As a result of the lookup, the egress interface for the packet is identified as lt-0/0/0.3 with the next hop identified as lt-0/0/0.5, the ingress interface in the ls-marketing-dept.

  - A session is established between ge-0/0/5 and lt-0/0/0.3.

- A management session is established in the destination logical system.

  - The packet is injected into the flow again from lt-0/0/0.5, and the logical system context identified as ls-marketing-dept is derived from the interface.

  - Packet processing continues in the ls-marketing-dept logical system.

- Route lookup for the packet is performed in the mk-vr1 routing instance. The packet is terminated in the destination logical system as self-traffic.

- A management session is established.

Understanding Session and Gate Limitation Control

The logical systems flow module provides session and gate limitation to ensure that these resources are shared fairly among the logical systems. Resources allocation and limitations for each logical system are specified in the security profile bound to the logical system.

- For session limiting, the system checks the first packet of a session against the maximum number of sessions configured for the logical system. If the maximum is reached, the device drops the packet and logs the event.

- For gate limiting, the device checks the first packet of a session against the maximum number of gates configured for the logical system. If the maximum number of gates for a logical system is reached, the device rejects the gate open request and logs the event.

Understanding Sessions

Sessions are created based on routing and other classification information to store information and allocate resources for a flow. You can change some characteristics of sessions, such as when a session is terminated. For example, you might want to ensure that a session table is never entirely full to protect against an attacker's attempt to flood the table and thereby prevent legitimate users from starting sessions.

About Configuring Sessions

Depending on the protocol and service, a session is programmed with a timeout value. For example, the default timeout for TCP is 1800 seconds. The default timeout for UDP is 60 seconds. When a flow is terminated, it is marked as invalid, and its timeout is reduced to 10 seconds. If no traffic uses the If no traffic uses the session before the service timeout, the session is aged out and freed to a common resource pool for reuse.

You can affect the life of a session in the following ways:

- Age out sessions, based on how full the session table is.

- Set an explicit timeout for aging out TCP sessions.

- Configure a TCP session to be invalidated when it receives a TCP RST (reset) message.

- You can configure sessions to accommodate other systems as follows:

  - Disable TCP packet security checks.

  - Change the maximum segment size.

Related Documentation

- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 8

CHAPTER 2

Understanding Master Logical Systems

Understanding the Master Logical System and the Master Administrator Role

When, as a master administrator, you initialize an SRX Series device running logical systems, a master logical system is created at the root level. You can log in to the device as root and change the root password.

By default, all system resources are assigned to the master logical system, and the master administrator allocates them to the user logical systems.

As master administrator, you manage the device and all its logical systems. You also manage the master logical system and configure its assigned resources. There can be more than one master administrator managing a device running logical systems.

- The master administrator's role and main responsibilities include:

  - Creating user logical systems and configuring their administrators. You can create one or more user logical system administrators for each user logical system.

  - Creating login accounts for users for all logical systems and assigning them to the appropriate logical systems.

  - Configuring an interconnect logical system if you want to allow communication between logical systems on the device. The interconnect logical system acts as an internal switch. It does not require an administrator.

    To configure an interconnect logical system, you configure lt-0/0/0 interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for establishment of tunnels.

  - Configuring security profiles to provision portions of the system's security resources to user logical systems and the master logical system.

    Only the master administrator can create, change, and delete security profiles and bind them to logical systems.

> **NOTE:** A user logical system administrator can configure interface, routing, and security resources allocated to his logical system.

- Creating logical interfaces to assign to user logical systems. (The user logical system administrator configures logical interfaces assigned to his logical system.)

- Viewing and managing user logical systems, as required, and deleting user logical systems. When a user logical system is deleted, its allocated reserved resources are released for use by other logical systems.

- Configuring IDP, AppTrack, application identification, and application firewall features. The master administrator can also use trace and debug at the root level, and he can perform commit rollbacks. The master administrator manages the master logical system and configures all the features that a user logical system administrator can configure for his or her own logical systems including routing instances, static routes, dynamic routing protocols, zones, security policies, screens, and firewall authentication.

Related Documentation

- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

- Understanding Logical Systems for SRX Series Services Gateways on page 3

- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 140

SRX Series Logical System Master Administrator Configuration Tasks Overview

This topic identities and describes the master administrator's tasks in the order in which they are performed.

An SRX Series device running logical systems is managed by a master administrator. The master administrator has the same capabilities as the root administrator of an SRX Series device not running logical systems. However, the master administrator's role and responsibilities extend beyond those of other SRX Series device administrators because an SRX Series device running logical systems is partitioned into discrete logical systems, each with its own resources, configuration, and management concerns. The master administrator is responsible for creating these user logical systems and provisioning them with resources.

For an overview of the master administrator's role and responsibilities, see "Understanding the Master Logical System and the Master Administrator Role" on page 19.

As the master administrator, you perform the following tasks to configure an SRX Series device running logical systems:

1. Configure a root password. Initially the master administrator logs in to the device as the root user without needing to specify a password. After you log in to the device, you must define a root password for later use.

   See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 39 for configuration information.

2. Create user logical systems and their administrators and users. Optionally, create an interconnect logical system.

   For each user logical system that you want to configure on the device, you must create a logical system, define one or more administrators for it, and add users to it.

   The master administrator configures login accounts for user logical system administrators and users and associates them with the user logical system. A user logical system can have more than one administrator; the master administrator must define and add all user logical system administrators and add them to their user logical systems.

   The master administrator adds users to user logical systems on behalf of the user logical system administrator. For example, if you have created a user logical system for the product design department, you must create user accounts for the users who belong to that department and associate them with the user logical system. The user logical system administrator does not have the ability to do this. Rather, the user logical administrator tells you the user accounts that you must create and add for his logical system.

   If you intend to use an internal virtual private LAN service (VPLS) switch to allow logical systems to communicate with one another, you must create an interconnect logical system. An interconnect logical system does not require an administrator.

   - For configuration information, see "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40

   - For information on user logical system administrators, see "Understanding User Logical Systems and the User Logical System Administrator Role" on page 25.

   - For information on the interconnect logical system, see "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 8.

3. Configure one or more security profiles. Security profiles assign security resources to logical systems. You can assign a single security profile to more than one logical system if you intend to allocate the same kinds and amounts of resources to them.

   - For configuration information, see "Example: Configuring Logical Systems Security Profiles (Master Administrators Only)" on page 56.

   - For information on security profiles, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 51.

4. Configure interfaces, routing instances, and static routes for logical systems, as appropriate.

   - If you plan to use an interconnect logical system, configure its logical tunnel interfaces and add them to its virtual routing instance.

- Configure interfaces for the master logical system. Optionally, create its logical tunnel interface to allow it to communicate with other logical systems on the device. Create a virtual routing instance for the master logical system and add its interfaces and static routes to it. Also configure logical interfaces for user logical systems with VLAN tagging.

> NOTE: The master administrator tells the user logical system administrators which interfaces are assigned to their logical systems. It is the user logical system administrator's responsibility to configure their interfaces.

- Optionally, configure logical tunnel interfaces for any user logical systems that you want to allow to communicate with one another using the internal VPLS switch.

- For configuration information, see "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)" on page 140.

- For information about the interconnect logical system and logical tunnel (lt-0/0/0) interfaces, see "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 8.

5. Enable CPU utilization control and configure the CPU control target and reserved CPU quotas for logical systems. See "Example: Configuring CPU Utilization (Master Administrators Only)" on page 263.

6. Optionally, configure dynamic routing protocols for the master logical system. See "Example: Configuring OSPF Routing Protocol for the Master Logical System" on page 148

7. Configure zones, security policies, and security features for the master logical system. See "Example: Configuring Security Features for the Master Logical System" on page 69.

8. Configure IDP for the master logical system. See "Example: Configuring an IDP Policy for the Master Logical System" on page 79.

9. Configure application firewall services on the master logical system. See "Understanding Logical System Application Firewall Services" on page 85 and "Example: Configuring Application Firewall Services for a Master Logical System" on page 86.

10. Configure a route-based VPN to secure traffic between a logical system and a remote site. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)" on page 92.

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

Understanding User Logical Systems

User Logical System Configuration Overview

When the master administrator creates a user logical system, he assigns a user logical system administrator to manage it. A user logical system can have multiple user logical system administrators.

As a user logical system administrator, you can access and view resources in your user logical system but not those of other user logical systems or the master logical system. You can configure resources allocated to your user logical system, but you cannot modify the numbers of allocated resources.

The following procedure lists the tasks that the user logical system administrator performs to configure resources in the user logical system:

1. Log in to the user logical system with the login and password configured by the master administrator:

   a. Telnet or SSH to the management IP address configured on the device. Log in to the user logical system with the administrator login and password provided by the master administrator.

      You enter a UNIX shell in the user logical system configured by the master administrator.

   b. The presence of the > prompt indicates the CLI has started. The prompt is preceded by a string that contains your username, the hostname of the router, and the name of the user logical system. When the CLI starts, you are at the top level in operational mode. You enter configuration mode by entering the **configure** operational mode command. The CLI prompt changes from user@host: *logical-system* > to user@host: *logical-system* #.

      To exit the CLI and return to the UNIX shell, enter the **quit** command.

2. Configure the logical interfaces assigned to the user logical system by the master administrator. Configure one or more routing instances and the routing protocols and options within each instance. See "Example: Configuring Interfaces and Routing Instances for a User Logical System" on page 158.

3. Configure security resources for the user logical system:

   a. Create zones for the user logical system and bind the logical interfaces to the zones. Address books can be created that are attached to zones for use in policies. See "Example: Configuring Zones for a User Logical System" on page 99.

   b. Configure screen options at the zone level. See "Example: Configuring Screen Options for a User Logical System" on page 102.

   c. Configure security policies between zones in the user logical system. See "Example: Configuring Security Policies in a User Logical System" on page 106.

   Custom applications or application sets can be created for specific types of traffic. To create a custom application, use the **application** configuration statement at the [**edit applications**] hierarchy level. To create an application set, use the **application-set** configuration statement at the [**edit applications**] hierarchy level.

   d. Configure firewall authentication. The master administrator creates access profiles in the master logical system. See "Example: Configuring Access Profiles (Master Administrators Only)" on page 67.

   The user logical system administrator then configures a security policy that specifies firewall authentication for matching traffic and configures the type of authentication (pass-through or Web authentication), default access profile, and success banner. See "Example: Configuring Firewall Authentication for a User Logical System" on page 111.

   e. Configure a route-based VPN tunnel to secure traffic between a user logical system and a remote site. The master administrator assigns a secure tunnel interface to the user logical system and configures IKE and IPsec SAs for the VPN tunnel. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)" on page 92.

   The user logical system administrator then configures a route-based VPN tunnel. See "Example: Configuring a Route-Based VPN Tunnel in a User Logical System" on page 133.

   f. Configure Network Address Translation (NAT). See "Example: Configuring Network Address Translation for a User Logical System" on page 154.

   g. Enable IDP. The master administrator configures IDP policies at the root level and specifies an IDP policy in the security profile that is bound to a logical system. See "Example: Configuring an IDP Policy for a User Logical System" on page 119.

   The user logical system administrator then enables IDP in a security policy. See "Example: Enabling IDP in a User Logical System Security Policy" on page 121.

   h. Display or clear application system cache (ASC) entries. See "Understanding Logical System Application Identification Services" on page 84.

i. Configure application firewall services on a user logical system. See "Understanding Logical System Application Firewall Services" on page 85 and "Example: Configuring Application Firewall Services for a User Logical System" on page 124.

j. Configure the AppTrack application tracking tool. See "Example: Configuring AppTrack for a User Logical System" on page 129.

Related Documentation

- Example: Configuring User Logical Systems on page 26
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Understanding User Logical Systems and the User Logical System Administrator Role

Logical systems allow a master administrator to partition an SRX Series device into discrete contexts called user logical systems. User logical systems are self-contained, private contexts, separate both from one another and from the master logical system. A user logical system has its own security, networking, logical interfaces, routing configurations, and one or more user logical system administrators.

When the master administrator creates a user logical system, he assigns one or more user logical system administrators to manage it. A user logical system administrator has a view of the device that is limited to his logical system. Although a user logical system is managed by a user logical system administrator, the master administrator has a global view of the device and access to all user logical systems. If necessary, the master administrator can manage any user logical system on the device.

The role and responsibilities of a user logical system administrator differ from those of the master administrator. As a user logical system administrator, you can access, configure, and view the configuration for your user logical system resources, but not those of other user logical systems or the master logical system.

As a user logical system administrator, you can:

- Configure zones, address books, security policies, user lists, custom services, and so forth, for your user logical system environment, based on the resources allocated to it.

  For example, if the master administrator allocates 40 zones to your user logical system, you can configure and administer those zones, but you cannot change the allocated number.

- Configure routing instances and assign allotted interfaces to them. Create static routes and add them to your routing instances. Configure routing protocols.

- Configure, enable, and monitor application firewall policy on your user logical system.

- Configure AppTrack.

- View all assigned logical interfaces and configure their attributes. The attributes that you configure for logical interfaces for your user logical system cannot be seen by other user logical system administrators.

- Run operational commands for your user logical system.

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3

- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 140

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 56

Example: Configuring User Logical Systems

This example shows the configuration of interfaces, routing instances, zones, and security policies for user logical systems.

- Requirements on page 26
- Overview on page 26
- Configuration on page 28
- Verification on page 36

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Be sure you know which logical interfaces and optionally, which logical tunnel interface (and its IP address) are allocated to your user logical system by the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

Overview

This example configures the ls-marketing-dept and ls-accounting-dept user logical systems shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example configures the parameters described in Table 3 on page 27 and Table 4 on page 27.

Table 3: ls-marketing-dept Logical System Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/6.1	• IP address 13.1.1.1/24 • VLAN ID 800
Routing instance	mk-vr1	• Instance type: virtual router • Includes interfaces ge-0/0/6.1 and lt-0/0/0.5 • Static routes: 　◦ 12.1.1.0/24 next-hop 10.0.1.2 　◦ 14.1.1.0/24 next-hop 10.0.1.4 　◦ 12.12.1.0/24 next-hop 10.0.1.1
Zones	ls-marketing-trust	Bind to interface ge-0/0/6.1.
	ls-marketing-untrust	Bind to interface lt-0/0/0.5
Address books	marketing-internal	• Address marketers: 13.1.1.0/24 • Attach to zone ls-marketing-trust
	marketing-external	• Address design: 12.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: design, accounting • Attach to zone ls-marketing-untrust
Policies	permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-marketing-trust • To zone: ls-marketing-untrust • Source address: marketers • Destination address: otherlsys • Application: any
	permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-marketing-untrust • To zone: ls-marketing-trust • Source address: otherlsys • Destination address: marketers • Application: any

Table 4: ls-accounting-dept Logical System Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/7.1	• IP address 14.1.1.1/24 • VLAN ID 900

Table 4: ls-accounting-dept Logical System Configuration *(continued)*

Feature	Name	Configuration Parameters

Routing instance	acct-vr1	• Instance type: virtual router
		• Includes interfaces ge-0/0/7.1 and lt-0/0/0.7
		• Static routes:
		  ◦ 12.1.1.0/24 next-hop 10.0.1.2
		  ◦ 13.1.1.0/24 next-hop 10.0.1.3
		  ◦ 12.12.1.0/24 next-hop 10.0.1.1

Zones	ls-accounting-trust	Bind to interface ge-0/0/7.1.
	ls-accounting-untrust	Bind to interface lt-0/0/0.7
Address books	accounting-internal	• Address accounting: 14.1.1.0/24 • Attach to zone ls-accounting-trust
	accounting-external	• Address design: 12.1.1.0/24 • Address marketing: 13.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: design, marketing • Attach to zone ls-accounting-untrust
Policies	permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-accounting-trust • To zone: ls-accounting-untrust • Source address: accounting • Destination address: otherlsys • Application: any
	permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-accounting-untrust • To zone: ls-accounting-trust • Source address: otherlsys • Destination address: accounting • Application: any

Configuration

Configuring the ls-marketing-dept User Logical System

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

```
set interfaces ge-0/0/6 unit 1 family inet address 13.1.1.1/24
set interfaces ge-0/0/6 unit 1 vlan-id 800
set routing-instances mk-vr1 instance-type virtual-router
set routing-instances mk-vr1 interface ge-0/0/6.1
set routing-instances mk-vr1 interface lt-0/0/0.5
set routing-instances mk-vr1 routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances mk-vr1 routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
set routing-instances mk-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
set security zones security-zone ls-marketing-trust interfaces ge-0/0/6.1
set security zones security-zone ls-marketing-untrust interfaces lt-0/0/0.5
set security address-book marketing-external address design 12.1.1.0/24
set security address-book marketing-external address accounting 14.1.1.0/24
set security address-book marketing-external address others 12.12.1.0/24
set security address-book marketing-external address-set otherlsys address design
set security address-book marketing-external address-set otherlsys address accounting
set security address-book marketing-external attach zone ls-marketing-untrust
set security address-book marketing-internal address marketers 13.1.1.0/24
set security address-book marketing-internal attach zone ls-marketing-trust
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
    permit-all-to-otherlsys match source-address marketers
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
    permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
    permit-all-to-otherlsys match application any
set security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust policy
    permit-all-to-otherlsys then permit
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
    permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
    permit-all-from-otherlsys match destination-address marketers
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
    permit-all-from-otherlsys match application any
set security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust policy
    permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

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

To configure a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

    ```
    lsmarketingadmin1@host:ls-marketing-dept> configure
    lsmarketingadmin1@host:ls-marketing-dept#
    ```

2. Configure the logical interface for a user logical system.

    ```
    [edit interfaces]
    lsmarketingadmin1@host:ls-marketing-dept# set ge-0/0/6 unit 1 family inet address
        13.1.1.1/24
    lsmarketingadmin1@host:ls-marketing-dept# set ge-0/0/6 unit 1 vlan-id 800
    ```

3. Configure the routing instance and assign interfaces.

    ```
    [edit routing-instances]
    ```

lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 instance-type virtual-router
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 interface ge-0/0/6.1
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 interface lt-0/0/0.5

4. Configure static routes.

[edit routing-instances]
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
12.12.1.0/24 next-hop 10.0.1.2
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
14.1.1.0/24 next-hop 10.0.1.4
lsmarketingadmin1@host:ls-marketing-dept# set mk-vr1 routing-options static route
12.12.1.0/24 next-hop 10.0.1.1

5. Configure security zones and assign interfaces to each zone.

[edit security zones]
lsmarketingadmin1@host:ls-marketing-dept# set security-zone ls-marketing-trust
interfaces ge-0/0/6.1
lsmarketingadmin1@host:ls-marketing-dept# set security-zone ls-marketing-untrust
interfaces lt-0/0/0.5

6. Create address book entries.

[edit security]
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-internal
address marketers 13.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
address design 12.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
address accounting 14.1.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
address others 12.12.1.0/24
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
address-set otherlsys address design
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
address-set otherlsys address accounting

7. Attach address books to zones.

[edit security]
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-internal
attach zone ls-marketing-trust
lsmarketingadmin1@host:ls-marketing-dept# set address-book marketing-external
attach zone ls-marketing-untrust

8. Configure a security policy that permits traffic from the ls-marketing-trust zone to
the ls-marketing-untrust zone.

[edit security policies from-zone ls-marketing-trust to-zone ls-marketing-untrust]
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys
match source-address marketers
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys
match destination-address otherlsys
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys
match application any
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-to-otherlsys then
permit

9. Configure a security policy that permits traffic from the ls-marketing-untrust zone to the ls-marketing-trust zone.

```
[edit security policies from-zone ls-marketing-untrust to-zone ls-marketing-trust]
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
    match source-address otherlsys
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
    match destination-address marketers
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
    match application any
lsmarketingadmin1@host:ls-marketing-dept# set policy permit-all-from-otherlsys
    then permit
```

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

```
lsmarketingadmin1@host:ls-marketing-dept# show routing instances
  mk-vr1 {
    instance-type virtual-router;
    interface ge-0/0/6.1;
    interface lt-0/0/0.5;
    routing-options {
      static {
        route 12.1.1.0/24 next-hop 10.0.1.2;
        route 14.1.1.0/24 next-hop 10.0.1.4;
        route 12.12.1.0/24 next-hop 10.0.1.1;
      }
    }
  }
lsmarketingadmin1@host:ls-marketing-dept# show security
  address-book {
    marketing-external {
      address product-designers 12.1.1.0/24;
      address accounting 14.1.1.0/24;
      address others 12.12.1.0/24;
      address-set otherlsys {
        address product-designers;
        address accounting;
      }
      attach {
        zone ls-marketing-untrust;
      }
    }
    marketing-internal {
      address marketers 13.1.1.0/24;
      attach {
        zone ls-marketing-trust;
      }
    }
  }
    policies {
      from-zone ls-marketing-trust to-zone ls-marketing-untrust {
        policy permit-all-to-otherlsys {
          match {
```

```
                    source-address marketers;
                    destination-address otherlsys;
                    application any;
                  }
                  then {
                    permit;
                  }
                }
              }
              from-zone ls-marketing-untrust to-zone ls-marketing-trust {
                policy permit-all-from-otherlsys {
                  match {
                    source-address otherlsys;
                    destination-address marketers;
                    application any;
                  }
                  then {
                    permit;
                  }
                }
              }
            }
          }
          zones {
            security-zone ls-marketing-trust {
              interfaces {
                ge-0/0/6.1;
              }
            }
            security-zone ls-marketing-untrust {
              interfaces {
                lt-0/0/0.5;
              }
            }
          }
        }
```

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

Configuring the ls-accounting-dept User Logical System

CLI Quick
Configuration

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

```
set interfaces ge-0/0/7 unit 1 family inet address 14.1.1.1/24
set interfaces ge-0/0/7 unit 1 vlan-id 900
set routing-instances acct-vr1 instance-type virtual-router
set routing-instances acct-vr1 interface ge-0/0/7.1
set routing-instances acct-vr1 interface lt-0/0/0.7
set routing-instances acct-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
set routing-instances acct-vr1 routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances acct-vr1 routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set security address-book accounting-internal address accounting 14.1.1.0/24
set security address-book accounting-internal attach zone ls-accounting-trust
set security address-book accounting-external address design 12.1.1.0/24
```

```
set security address-book accounting-external address marketing 13.1.1.0/24
set security address-book accounting-external address others 12.12.1.0/24
set security address-book accounting-external address-set otherlsys address design
set security address-book accounting-external address-set otherlsys address marketing
set security address-book accounting-external attach zone ls-accounting-untrust
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
    permit-all-to-otherlsys match source-address accounting
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
    permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
    permit-all-to-otherlsys match application any
set security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust policy
    permit-all-to-otherlsys then permit
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
    permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
    permit-all-from-otherlsys match destination-address accounting
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
    permit-all-from-otherlsys match application any
set security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust policy
    permit-all-from-otherlsys then permit
set security zones security-zone ls-accounting-trust interfaces ge-0/0/7.1
set security zones security-zone ls-accounting-untrust interfaces lt-0/0/0.7
```

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

To configure a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

    ```
    lsaccountingadmin1@host:ls-accounting-dept> configure
    lsaccountingadmin1@host:ls-accounting-dept#
    ```

2. Configure the logical interface for a user logical system.

    ```
    [edit interfaces]
    lsaccountingadmin1@host:ls-accounting-dept# set ge-0/0/7 unit 1 family inet
        address 14.1.1.1/24
    lsaccountingadmin1@host:ls-accounting-dept# set ge-0/0/7 unit 1 vlan-id 900
    ```

3. Configure the routing instance and assign interfaces.

    ```
    [edit routing-instances]
    lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 instance-type
        virtual-router
    lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 interface ge-0/0/7.1
    lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 interface lt-0/0/0.7
    ```

4. Configure static routes.

    ```
    [edit routing-instances]
    lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 routing-options static
        route 12.1.1.0/24 next-hop 10.0.1.2
    lsaccountingadmin1@host:ls-accounting-deptt# set acct-vr1 routing-options static
        route 13.1.1.0/24 next-hop 10.0.1.3
    ```

lsaccountingadmin1@host:ls-accounting-dept# set acct-vr1 routing-options static
    route 12.12.1.0/24 next-hop 10.0.1.1

5. Configure security zones and assign interfaces to each zone.

    [edit security zones]
    lsaccountingadmin1@host:ls-accounting-dept# set security-zone ls-accounting-trust
        interfaces ge-0/0/7.1
    lsaccountingadmin1@host:ls-accounting-dept# set security-zone
        ls-accounting-untrust interfaces lt-0/0/0.7

6. Create address book entries.

    [edit security]
    lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-internal
        address accounting 14.1.1.0/24
    lsaccountingadmin1@host:ls-accounting-dept# set address-book
        accounting-external address design 12.1.1.0/24
    lsaccountingadmin1@host:ls-accounting-dept# set address-book
        accounting-external address marketing 13.1.1.0/24
    lsaccountingadmin1@host:ls-accounting-dept# set address-book
        accounting-external address others 12.12.1.0/24
    lsaccountingadmin1@host:ls-accounting-dept# set address-book
        accounting-external address-set otherlsys address design
    lsaccountingadmin1@host:ls-accounting-dept# set address-book
        accounting-external address-set otherlsys address marketing

7. Attach address books to zones.

    [edit security]
    lsaccountingadmin1@host:ls-accounting-dept# set address-book accounting-internal
        attach zone ls-accounting-trust
    lsaccountingadmin1@host:ls-accounting-dept# set address-book
        accounting-external attach zone ls-accounting-untrust

8. Configure a security policy that permits traffic from the ls-accounting-trust zone to
    the ls-accounting-untrust zone.

    [edit security policies from-zone ls-accounting-trust to-zone ls-accounting-untrust]
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
        match source-address accounting
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
        match destination-address otherlsys
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
        match application any
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-to-otherlsys
        then permit

9. Configure a security policy that permits traffic from the ls-accounting-untrust zone
    to the ls-accounting-trust zone.

    [edit security policies from-zone ls-accounting-untrust to-zone ls-accounting-trust]
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
        match source-address otherlsys
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
        match destination-address accounting
    lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
        match application any

lsaccountingadmin1@host:ls-accounting-dept# set policy permit-all-from-otherlsys
then permit

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

```
lsaccountingadmin1@host:ls-accounting-dept# show routing-instances
acct-vr1 {
    instance-type virtual-router;
    interface ge-0/0/7.1;
    interface lt-0/0/0.7;
    routing-options {
        static {
            route 12.12.1.0/24 next-hop 10.0.1.1;
            route 12.1.1.0/24 next-hop 10.0.1.2;
            route 13.1.1.0/24 next-hop 10.0.1.3;
        }
    }
}
lsaccountingadmin1@host:ls-accounting-dept# show security
address-book {
    accounting-internal {
        address accounting 14.1.1.0/24;
        attach {
            zone ls-accounting-trust;
        }
    }
    accounting-external {
        address design 12.1.1.0/24;
        address marketing 13.1.1.0/24;
        address others 12.12.1.0/24;
        address-set otherlsys {
            address design;
            address marketing;
        }
        attach {
            zone ls-accounting-untrust;
        }
    }
}
    policies {
        from-zone ls-accounting-trust to-zone ls-accounting-untrust {
            policy permit-all-to-otherlsys {
                match {
                    source-address accounting;
                    destination-address otherlsys;
                    application any;
                }
                then {
                    permit;
                }
            }
        }
        from-zone ls-accounting-untrust to-zone ls-accounting-trust {
```

```
                                policy permit-all-from-otherlsys {
                                  match {
                                    source-address otherlsys;
                                    destination-address accounting;
                                    application any;
                                  }
                                  then {
                                    permit;
                                  }
                                }
                              }
                            }
                            zones {
                              security-zone ls-accounting-trust {
                                interfaces {
                                  ge-0/0/7.1;
                                }
                              }
                              security-zone ls-accounting-untrust {
                                interfaces {
                                  lt-0/0/0.7;
                                }
                              }
                            }
                          }
```

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

Verification

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

Verifying Policy Configuration

Purpose   Verify information about policies and rules.

Action   From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related
Documentation

PART 2

Getting Started for Master Administrators

- Configuring Device for Master Logical Systems on page 39

Configuring Device for Master Logical Systems

Example: Configuring a Root Password for the Device (Master Administrators Only)

Requirements

Before you begin, read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 to understand how this task fits into the overall configuration process.

The example uses an SRX5600 device running Junos OS with logical systems.

Overview

The Junos OS software is installed on the router before it is delivered from the factory. When you power on your router, it is ready for you to configure. Initially you log in as *root* user without using a password.

After you log in, you can configure a password for the root user, or, in logical systems terms, the master administrator. The master administrator has root privileges over the device.

Configuration

Configuring the Root Password

Step-by-Step Procedure
- Configure a root password for the device.

  user@host# set system root-authentication Talk22rt6

Related Documentation
- Understanding the Master Logical System and the Master Administrator Role on page 19
- Understanding Logical Systems for SRX Series Services Gateways on page 3

Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)

This example shows how to create user logical systems and assign administrators to them. It shows how to add users to a user logical system. And the example shows how to create an interconnect logical system, which is optional.

> **NOTE:** Only the master administrator can create user login accounts for administrators and users. If a user logical system administrator wants to add users to his logical system, he must convey the information to the master administrator, who will add the users.

- Requirements on page 40
- Overview on page 40
- Configuration on page 42
- Verification on page 46

Requirements

The example uses an SRX5600 device running Junos OS with logical systems.

Overview

Before you begin, read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 to understand how this task fits into the overall configuration process.

This example is for a company that includes product design, marketing, and accounting departments. The company wants to curtail hardware and energy costs, but not at the risk of exposing data across departments or to the Internet.

Each department has its own security requirements in regard both to other departments and to the Internet. To meet its requirements for cost control without forfeiting security, the company deploys the SRX5600 device. The master administrator configures three user logical systems giving each department a logical device that is private and fully secured.

This topic covers how to:

- Create user logical systems and an interconnect logical system that is used as an internal VPLS switch to allow traffic to pass from one logical system to another.

- Create administrators for user logical systems other than the interconnect logical system. A user logical system can have more than one administrator. The interconnect logical system does not require an administrator.

- Add users to a user logical system.

> **(i)** NOTE: This example shows how to configure only two users—lsdesignuser1 and lsdesignuser2. In reality, every user logical system will include many users that would require configurations similar to those shown in this example.

Figure 3 on page 41 shows an SRX5600 device deployed and configured for logical systems. The configuration examples reflect this deployment.

Figure 3: SRX Series Device Configured for Logical Systems

Configuration

- Configuring User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System on page 42

Configuring User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System

CLI Quick Configuration

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

```
set logical-systems ls-product-design
set system login class ls-design-admin logical-system ls-product-design
set system login class ls-design-admin permissions all
set system login user lsdesignadmin1 full-name lsdesignadmin1
set system login user lsdesignadmin1 class ls-design-admin
set system login user lsdesignadmin1 authentication encrypted-password "$ABC123"
set system login class ls-design-user logical-system ls-product-design
set system login class ls-design-user permissions view
set system login user lsdesignuser1 full-name lsdesignuser1
set system login user lsdesignuser1 class ls-design-user
set system login user lsdesignuser1 authentication encrypted-password "$ABC123"
set system login user lsdesignuser2 full-name lsdesignuser2
set system login user lsdesignuser2 class ls-design-user
set system login user lsdesignuser2 authentication encrypted-password "$ABC123"
set logical-systems ls-marketing-dept
set system login class ls-marketing-admin logical-system ls-marketing-dept
set system login class ls-marketing-admin permissions all
set system login user lsmarketingadmin1 class ls-marketing-admin
set system login user lsmarketingadmin1 full-name lsmarketingadmin1
set system login user lsmarketingadmin1 authentication encrypted-password "$ABC123"
set system login user lsmarketingadmin2 full-name lsmarketingadmin2
set system login user lsmarketingadmin2 class ls-marketing-admin
set system login user lsmarketingadmin2 authentication encrypted-password "$ABC123"
set logical-systems ls-accounting-dept
set system login class ls-accounting-admin logical-system ls-accounting-dept
set system login class ls-accounting-admin permissions all
set system login user lsaccountingadmin1 full-name lsaccountingadmin1
set system login user lsaccountingadmin1 class ls-accounting-admin
set system login user lsaccountingadmin1 authentication encrypted-password "$ABC123"
set logical-systems interconnect-logical-system
```

Step-by-Step Procedure

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

1. Create the first user logical system and define its administrator.

   a. Create the user logical system.

      ```
      [edit]
      user@host# set logical-systems ls-product-design
      ```

b. Assign the user login class to the user logical system.

[edit system]
user@host# **set login class ls-design-admin logical-system ls-product-design**

c. Create the login class to give the user logical system administrator full permission over the user logical system.

[edit system]
user@host# **set login class ls-design-admin permissions all**

d. Assign a full name to the user logical system administrator.

[edit system]
user@host# **set login user lsdesignadmin1 full-name lsdesignadmin1**

e. Associate the login class with the user logical system administrator to allow the administrator to log in to the user logical system.

[edit system]
user@host# **set login user lsdesignadmin1 class ls-design-admin**

f. Create a user login password for the user logical system administrator.

[edit system]
user@host# **set login user lsdesignadmin1 authentication plain-text-password**
New password: **Talk1234**
Retype new password: **Talk1234**

2. Configure the first user for the logical system.

a. Configure the user login class and assign it to the user logical system.

[edit system]
user@host# **set login class ls-design-user logical-system ls-product-design**

b. To give the first user the ability to see the logical system's resources and settings but not change them, assign **view** as the permission to the login class.

[edit system]
user@host# **set login class ls-design-user permissions view**

c. Assign a full name to the logical system user.

[edit system]
user@host# **set login user lsdesignuser1 full-name lsdesignuser1**

d. Associate the login class with the user to allow the user to log in to the user logical system.

user@host# **set login user lsdesignuser1 class ls-design-user**

e. Create a user login password for the user.

[edit system]
user@host# **set login user lsdesignuser1 authentication plain-text-password**
New password: **Talk4234**

Retype new password: **Talk4234**

3. Create the second user for logical system ls-product-design.

   a. Assign a full name to the user.

      [edit system]
      user@host# **set login user lsdesignuser2 full-name lsdesignuser2**

   b. Associate the user with the login class to allow the user to log in to the user logical system.

      user@host# **set login user lsdesignuser2 class ls-design-user**

   c. Create a user login password.

      [edit system]
      user@host# **set login user lsdesignuser2 authentication plain-text-password**
      New password: **Talk9234**
      Retype new password: **Talk9234**

4. Create the second user logical system and define its administrator.

   a. Create the user logical system.

      [edit]
      user@host# **set logical-systems ls-marketing-dept**

   b. Configure the user login class and assign it to the user logical system.

      [edit system]
      user@host# **set login class ls-marketing-admin logical-system ls-marketing-dept**

   c. To give the user logical system administrator control over the user logical system, assign **all** as the permissions to the login class.

      [edit system]
      user@host# **set login class ls-marketing-admin permissions all**

   d. Assign a full name to the user logical system administrator.

      [edit system]
      user@host# **set login user lsmarketingadmin1 full-name lsmarketingadmin1**

   e. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

      [edit system]
      user@host# **set login user lsmarketingadmin1 class ls-marketing-admin**

   f. Create a user login password for the user logical system administrator.

      [edit system]
      user@host# **set login user lsmarketingadmin1 authentication plain-text-password**
      New password: **Talk2345**
      Retype new password: **Talk2345**

5. Create a second user logical system administrator for the ls-marketing-dept logical system.

   a. Assign a full name to the user logical system administrator.

[edit system]
user@host# **set login user lsmarketingadmin2 full-name lsmarketingadmin2**

b. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

[edit system]
user@host# **set login lsmarketingadmin2 class ls-marketing-admin**

c. Create a user login password for the user logical system administrator.

[edit system]
user@host# **set login user lsmarketingadmin2 authentication plain-text-password**
New password: **Talk6345**
Retype new password: **Talk6345**

6. Create the third user logical system and define its administrator.

a. Create the user logical system.

[edit]
user@host# **set logical-systems ls-accounting-dept**

b. Configure the user login class and assign it to the user logical system.

[edit system]
user@host# **set login class ls-accounting-admin logical-system**
   ls-accounting-dept

c. To give the user logical system administrator control over the user logical system, assign permissions to the login class.

[edit system]
user@host# **set login class ls-accounting-admin permissions all**

d. Assign a full name to the user logical system administrator.

[edit system]
user@host# **set login user lsaccountingadmin1 full-name lsaccountingadmin1**

e. Associate the user logical system administrator with the login class to allow the administrator to log in to the user logical system.

[edit system]
user@host# **set login user lsaccountingadmin1 class ls-accounting-admin**

f. Create a login password for the user logical system administrator.

[edit system]
user@host# **set login user lsaccountingadmin1 authentication**
   plain-text-password
New password: **Talk5678**
Retype new password: **Talk5678**

7. Configure an interconnect logical system to allow logical systems to pass traffic from one to another.

user@host# **set logical-systems interconnect-logical-system**

Results   From configuration mode, confirm your configuration by entering the **show logical-systems** command to verify that the logical systems were created. Also enter the **show system login class** command for each class that you defined.

To ensure that the logical systems administrators were created, enter the **show system login user** command.

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

```
user@host# show logical-systems ?
interconnect-logical-system;
ls-accounting-dept;
ls-marketing-dept;
ls-product-design;

user@host# show system login class ls-design-admin
logical-system ls-product-design;
permissions all;

user@host# show system login class ls-design-user
logical-system ls-product-design
permissions view;

user@host show system login class ls-marketing-admin
logical-system ls-marketing-dept;
permissions all;

user@host show system login class ls-accounting-admin
logical-system ls-accounting-dept;
permissions all;

user@host show system login user ?
lsaccountingadmin1 lsaccountingadmin1
lsdesignadmin1 lsdesignadmin1
lsdesignuser2 lsdesignuser2
lsmarketingadmin1 lsmarketingadmin1
lsmarketingadmin2 lsmarketingadmin2
```

Verification

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

Verifying User Logical Systems and Login Configurations from the Master Logical System

Purpose   Verify that the user logical systems exist and that you, as the master administrator, can enter them from root. Return from a user logical system to the master logical system.

Action   From operational mode, enter the following command:

```
root@host> set cli logical-system ls-product-design
```

Logical system:ls-product-design
root@host:ls-product-design>

root@host:ls-product-design> **clear cli logical-system**
Cleared default logical system
root@host>

root@host> **set cli logical-system ls-marketing-dept**
Logical system:ls-marketing-dept
root@host:ls-marketing-dept>

root@host:ls-marketing-dept> **clear cli logical-system**
Cleared default logical system
root@host>

root@host> **set cli logical-system ls-accounting-dept**
Logical system:ls-accounting-dept
root@host:ls-accounting-dept>

root@host:ls-accounting-dept> **clear cli logical-system**
Cleared default logical system
root@host>

Verifying User Logical Systems and Login Configurations Using Telnet

Purpose　Verify that the user logical systems you created exist and that the administrators' login IDs and passwords that you created are correct.

Action　Use Telnet to log in to each user logical system as its user administrator would do.

1. Run Telnet specifying the IP address of your SRX Series device. For example:

   telnet 10.11.11.19

2. Enter the login ID and password for the administrator for one of the user logical systems that you created. After you log in, the prompt shows the administrator name. Notice how this result differs from the result produced when you log in to the user logical system from the master logical system at root. Repeat this procedure for all of your user logical systems.

   login: **lsdesignadmin1**
   Password: **Talk1234**
   lsdesignadmin1@host: ls-product-design>

Related Documentation
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 56

- Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only) on page 140

PART 3

Configuring Security Features

Configuring Master Logical System Security Profiles

Understanding Logical System Security Profiles (Master Administrators Only)

Logical systems allow you to virtually divide a supported SRX Series device into multiple devices, isolating one from another, securing them from intrusion and attacks, and protecting them from faulty conditions outside their own contexts. To protect logical systems, security resources are configured in a manner similar to how they are configured for a discrete device. However, as the master administrator, you must allocate the kinds and amounts of security resources to logical systems. The logical system administrator allocates resources for his own logical system.

An SRX Series device running logical systems can be partitioned into user logical systems, an interconnect logical system, if desired, and the default master logical system. When the system is initialized, the master logical system is created at the root level. All system resources are assigned to it, effectively creating a default master logical system security profile. To distribute security resources across logical systems, the master administrator creates security profiles that specify the kinds and amounts of resources to be allocated to a logical system that the security profile is bound to. Only the master administrator can configure security profiles and bind them to logical systems. The user logical system administrator configures these resources for his or her logical system.

Logical systems are defined largely by the resources allocated to them, including security components, interfaces, routing instances, static routes, and dynamic routing protocols. When the master administrator configures a user logical system, he binds a security profile to it. Any attempt to commit a configuration for a user logical system without a security profile bound to it will fail.

This topic includes the following sections:

Logical Systems Security Profiles

As master administrator, you can configure a single security profile to assign resources to a specific logical system, use the same security profile for more than one logical system, or use a mix of both methods. You can configure up to 32 security profiles on an SRX Series device running logical systems. When you reach the limit, you must delete a security profile and commit the configuration change before you can create and commit another security profile. In many cases fewer security profiles are needed because you might bind a single security profile to more than one logical system.

Security profiles allow you to:

- Share the device's resources, including policies, zones, addresses and address books, flow sessions, and various forms of NAT, among all logical systems appropriately. You can dedicate various amounts of a resource to the logical systems and allow them to compete for use of the free resources.

  Security profiles protect against one logical system exhausting a resource that is required at the same time by other logical systems. Security profiles protect critical system resources and maintain a fair level of performance among user logical systems when the device is experiencing heavy traffic flow. They defend against one user logical system dominating the use of resources and depriving other user logical systems of them.

- Configure the device in a scalable way to allow for future creation of additional user logical systems.

You must delete a logical system's security profile before you delete that logical system.

How the System Assesses Resources Assignment and Use Across Logical Systems

To provision a logical system with security resources, you, as a master administrator, configure a security profile that specifies for each resource:

- A reserved quota that guarantees that the specified resource amount is always available to the logical system.

- A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems must compete for global resources.

If a reserved quota is not configured for a resource, the default value is 0. If a maximum allowed quota is not configured for a resource, the default value is the global system quota for the resource (global system quotas are platform-dependent). The master administrator must configure appropriate maximum allowed quota values in the security profiles so the maximum resource usage of a specific logical system does not negatively impact other logical systems configured on the device. The master administrator must configure the appropriate maximum-allowed quota values in the security profiles so that the maximum resource usage of a specific logical system does not negatively impact other logical systems configured on the device.

The system maintains a count of all allocated resources that are reserved, used, and made available again when a logical system is deleted. This count determines whether resources are available to use for new logical systems or to increase the amount of the resources allocated to existing logical systems through their security profiles.

When a user logical system is deleted, its reserved resource allocations are released for use by other logical systems.

Resources configured in security profiles are characterized as static modular resources or dynamic resources. For static resources, we recommend setting a maximum quota for a resource equal or close to the amount specified as its reserved quota, to allow for scalable configuration of logical systems. A high maximum quota for a resource might give a logical system greater flexibility through access to a larger amount of that resource, but it would constrain the amount available to allocate to a new user logical system.

The difference between reserved and maximum allowed amounts for a dynamic resource is not important because dynamic resources are aged out and do not deplete the pool available for assignment to other logical systems.

The following resources can be specified in a security profile:

- Security policies, including schedulers

- Security zones

- Addresses and address books for security policies

- Application firewall rule sets

- Application firewall rules

- Firewall authentication

- Flow sessions and gates

- NAT, including:

  - Cone NAT bindings

  - NAT destination rule

  - NAT destination pool

  - NAT IP address in source pool without Port Address Translation (PAT)

> **NOTE:** IPv6 addresses in IPv6 source pools without PAT are not included in security profiles.

- NAT IP address in source pool with PAT
- NAT port overloading
- NAT source pool
- NAT source rule
- NAT static rule

> **NOTE:** All resources except flow sessions are static.

You can modify a logical system security profile dynamically while the security profile is assigned to other logical systems. However, to ensure that the system resource quota is not exceeded, the system takes the following actions:

- If a static quota is changed, system daemons that maintain logical system counts for resources specified in security profiles revalidate the security profile. This check identifies the number of resources assigned across all logical systems to determine whether the allocated resources, including their increased amounts, are available.

  These quota checks are the same quota checks that the system performs when you add a new user logical system and bind a security profile to it. The are also performed when you bind a different security profile from the security profile that is presently assigned to it to an existing user logical system (or the master logical system).

- If a dynamic quota is changed, no check is performed, but the new quota is imposed on future resource usage.

Cases: Assessments of Reserved Resources Assigned Through Security Profiles

To understand how the system assesses allocation of reserved resources through security profiles, consider the following three cases that address allocation of one resource, zones. To keep the example simple, 10 zones are allocated in security-profile-1: 4 reserved zones and 6 maximum zones. This example assumes that the full maximum amount specified—six zones—is available for the user logical systems. The system maximum number of zones is 10.

These cases address configuration across logical systems. They test to see whether a configuration will succeed or fail when it is committed based on allocation of zones.

shows the security profiles and their zone allocations.

Table 5: Security Profiles Used for Reserved Resource Assessments

Two Security Profiles Used in the Configuration Cases
security-profile-1 • zones reserved quota = 4 • zones maximum quota = 6 NOTE: Later the master administrator dynamically increases the reserved zone count specified in this profile.
master-logical-system-profile • zones maximum quota = 10 • no reserved quota

Table 6 on page 55 shows three cases that illustrate how the system assesses reserved resources for zones across logical systems based on security profile configurations.

- The configuration for the first case succeeds because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 8, which is less than the system maximum resource quota.

- The configuration for the second case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.

- The configuration for the third case fails because the cumulative reserved resource quota for zones configured in the security profiles bound to all logical systems is 12, which is greater than the system maximum resource quota.

Table 6: Reserved Resource Allocation Assessment Across Logical Systems

Reserved Resource Quota Checks Across Logical Systems
Example 1: Succeeds This configuration is within bounds: 4+4+0=8, maximum capacity =10. Security Profiles Used • The security profile security-profile-1 is bound to two user logical systems: user-logical-system-1 and user-logical-system-2. • The master-logical-system-profile profile is used exclusively for the master logical system. • user-logical-system-1 = 4 reserved zones. • user-logical-system-2 = 4 reserved zones. • master-logical-system = 0 reserved zones.

Table 6: Reserved Resource Allocation Assessment Across Logical Systems *(continued)*

Reserved Resource Quota Checks Across Logical Systems
Example 2: Fails This configuration is out of bounds: 4+4+4=12, maximum capacity =10. • user-logical-system-1 = 4 reserved zones. • user-logical-system-2 = 4 reserved zones. • master-logical-system = 0 reserved zones. • new-user-logical-system = 4 reserved zones. Security Profiles • The security profile security-profile-1 is bound to two user logical systems: user-logical-system-1 and user-logical-system-2. • The master-logical-system-profile is bound to the master logical system and used exclusively for it. • The master administrator configures a new user logical system called new-user-logical-system and binds security-profile-1 to it.
Example 3: Fails This configuration is out of bounds: 6+6=12, maximum capacity =10. The master administrator modifies the reserved zones quota in security-profile-1, increasing the count to 6. • user-logical-system-1 = 6 reserved zones. • user-logical-system-2 = 6 reserved zones. • master-logical-system = 0 reserved zones.

Related Documentation

-

-

-

Example: Configuring Logical Systems Security Profiles (Master Administrators Only)

This example shows how a master administrator configures three logical system security profiles to assign to user logical systems and the master logical system to provision them with security resources.

-

-

-

-

Requirements

The example uses an SRX5600 device running Junos OS with logical systems.

Before you begin, read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 to understand how this task fits into the overall configuration process.

Overview

This example shows how to configure security profiles for the following logical systems:

- The root-logical-system logical system. The security profile master-profile is assigned to the master, or root, logical system.

- The ls-product-design logical system. The security profile ls-design-profile is assigned to the logical system.

- The ls-marketing-dept logical system. The security profile ls-accnt-mrkt-profile is assigned to the logical system.

- The ls-accounting-dept logical system. The security profile ls-accnt-mrkt-profile is assigned to the logical system.

- The interconnect-logical-system, if you use one. You must assign a dummy, or null, security profile to it.

This configuration relies on the deployment shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

Configuration

- Configuring Logical System Security Profiles on page 57

Configuring Logical System Security Profiles

CLI Quick Configuration

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

```
set system security-profile master-profile policy maximum 65
set system security-profile master-profile policy reserved 60
set system security-profile master-profile zone maximum 22
set system security-profile master-profile zone reserved 17
set system security-profile master-profile flow-session maximum 3000
set system security-profile master-profile flow-session reserved 2100
set system security-profile master-profile nat-nopat-address maximum 115
set system security-profile master-profile nat-nopat-address reserved 100
set system security-profile master-profile nat-static-rule maximum 125
set system security-profile master-profile nat-static-rule reserved 100
set system security-profile master-profile idp
set system security-profile master-profile logical-system root-logical-system
set system security-profile ls-accnt-mrkt-profile policy maximum 65
```

```
set system security-profile ls-accnt-mrkt-profile policy reserved 60
set system security-profile ls-accnt-mrkt-profile zone maximum 22
set system security-profile ls-accnt-mrkt-profile zone reserved 17
set system security-profile ls-accnt-mrkt-profile flow-session maximum 2500
set system security-profile ls-accnt-mrkt-profile flow-session reserved 2000
set system security-profile ls-accnt-mrkt-profile nat-nopat-address maximum 125
set system security-profile ls-accnt-mrkt-profile nat-nopat-address reserved 100
set system security-profile ls-accnt-mrkt-profile nat-static-rule maximum 125
set system security-profile ls-accnt-mrkt-profile nat-static-rule reserved 100
set system security-profile ls-accnt-mrkt-profile logical-system ls-marketing-dept
set system security-profile ls-accnt-mrkt-profile logical-system ls-accounting-dept
set system security-profile ls-design-profile policy maximum 50
set system security-profile ls-design-profile policy reserved 40
set system security-profile ls-design-profile zone maximum 10
set system security-profile ls-design-profile zone reserved 5
set system security-profile ls-design-profile flow-session maximum 2500
set system security-profile ls-design-profile flow-session reserved 2000
set system security-profile ls-design-profile nat-nopat-address maximum 120
set system security-profile ls-design-profile nat-nopat-address reserved 100
set system security-profile ls-design-profile logical-system ls-product-design
set system security-profile interconnect-profile logical-system
    interconnect-logical-system
```

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

Create three security profiles.

1. Create the first security profile.

   a. Specify the number of maximum and reserved policies.

      ```
      [edit system security-profile]
      user@host# set master-profile policy maximum 65 reserved 60
      ```

   b. Specify the number of maximum and reserved zones.

      ```
      [edit system security-profile]
      user@host# set master-profile zone maximum 22 reserved 17
      ```

   c. Specify the number of maximum and reserved sessions.

      ```
      [edit system security-profile]
      user@host# set master-profile flow-session maximum 3000 reserved 2100
      ```

   d. Specify the number of maximum and reserved source NAT no-PAT addresses and static NAT rules.

      ```
      [edit system security-profile]
      user@host# set master-profile nat-nopat-address maximum 115 reserved 100
      user@host# set master-profile nat-static-rule maximum 125 reserved 100
      ```

   e. Enable intrusion detection and prevention (IDP). You can enable IDP only for the master (root) logical system.

      ```
      [edit system security-profile]
      user@host# set idp
      ```

f. Bind the security profile to the logical system.

    [edit system security-profile]
    user@host# set master-profile logical-system root-logical-system

2. Create the second security profile.

    a. Specify the number of maximum and reserved policies.

        [edit system security-profile]
        user@host# set ls-accnt-mrkt-profile policy maximum 65 reserved 60

    b. Specify the number of maximum and reserved zones.

        [edit system security-profile]
        user@host# set ls-accnt-mrkt-profile zone maximum 22 reserved 17

    c. Specify the number of maximum and reserved sessions.

        [edit system security-profile]
        user@host# set ls-accnt-mrkt-profile flow-session maximum 2500 reserved
            2000

    d. Specify the number of maximum and reserved source NAT no-PAT addresses.

        [edit system security-profile]
        user@host# set ls-accnt-mrkt-profile nat-nopat-address maximum 125 reserved
            100

    e. Specify the number of maximum and reserved static NAT rules.

        [edit system security-profile]
        user@host# set ls-accnt-mrkt-profile nat-static-rule maximum 125 reserved 100

    f. Bind the security profile to two logical systems.

        [edit system]
        user@host# set security-profile ls-accnt-mrkt-profile logical-system
            ls-marketing-dept
        user@host# set security-profile ls-accnt-mrkt-profile logical-system
            ls-accounting-dept

3. Create the third security profile.

    a. Specify the number of maximum and reserved policies.

        [edit system security-profile]
        user@host# set ls-design-profile policy maximum 50 reserved 40

    b. Specify the number of maximum and reserved zones.

        [edit system security-profile]
        user@host# set ls-design-profile zone maximum 10 reserved 5

    c. Specify the number of maximum and reserved sessions.

        [edit system security-profile]
        user@host# set ls-design-profile flow-session maximum 2500 reserved 2000

    d. Specify the number of maximum and reserved source NAT no-PAT addresses.

        [edit system security-profile]

user@host# set ls-design-profile nat-nopat-address maximum 120 reserved 100

4. Bind the security profile to a logical system.

   user@host# set system security-profile ls-design-profile logical-system
   ls-product-design

5. Bind a null security profile to the interconnect logical system.

   user@host# set system security-profile interconnect-profile logical-system
   interconnect-logical-system

Results   From configuration mode, confirm your configuration by entering the **show system security-profile** command to see all security profiles configured.

To see individual security profiles, enter the **show system security-profile master-profile**, the **show system security-profile ls-accnt-mrkt-profile** and, the **show system security-profile ls-design-profile** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host# show system security-profile
  interconnect-profile {
    logical-system interconnect-logical-system;
  }
    ls-accnt-mrkt-profile {
      policy {
        maximum 65;
        reserved 60;
      }
      zone {
        maximum 22;
        reserved 17;
      }
      flow-session {
        maximum 2500;
        reserved 2000;
      }
      nat-nopat-address {
        maximum 125;
        reserved 100;
      }
      nat-static-rule {
        maximum 125;
        reserved 100;
      }
      logical-system [ ls-marketing-dept ls-accounting-dept ];
    }
    ls-design-profile {
      policy {
        maximum 50;
        reserved 40;
      }
      zone {
        maximum 10;
        reserved 5;
      }
```

```
                    flow-session {
                      maximum 2500;
                      reserved 2000;
                    }
                    nat-nopat-address {
                      maximum 120;
                      reserved 100;
                    }
                    nat-static-rule {
                      maximum 125;
                      reserved 100;
                    }
                    logical-system ls-product-design;
                }
                master-profile {
                    policy {
                      maximum 65;
                      reserved 60;
                    }
                    zone {
                      maximum 22;
                      reserved 17;
                    }
                    flow-session {
                      maximum 3000;
                      reserved 2100;
                    }
                    nat-nopat-address {
                      maximum 115;
                      reserved 100;
                    }
                    nat-static-rule {
                      maximum 125;
                      reserved 100;
                    }
                    root-logical-system;
                }

            user@host# show system security-profile master-profile
                policy {
                  maximum 65;
                  reserved 60;
                }
                zone {
                  maximum 22;
                  reserved 17;
                }
                flow-session {
                  maximum 3000;
                  reserved 2100;
                }
                nat-nopat-address {
                  maximum 115;
                  reserved 100;
                }
                nat-static-rule {
```

```
                maximum 125;
                reserved 100;
            }
        root-logical-system;

user@host# show system security-profile ls-accnt-mrkt-profile
    policy {
        maximum 65;
        reserved 60;
    }
    zone {
        maximum 22;
        reserved 17;
    }
    flow-session {
        maximum 2500;
        reserved 2000;
    }
    nat-nopat-address {
        maximum 125;
        reserved 100;
    }
    nat-static-rule {
        maximum 125;
        reserved 100;
    }
    logical-system [ ls-accounting-dept ls-marketing-dept ];

user@host# show system security-profile ls-design-profile
    policy {
        maximum 50;
        reserved 40;
    }
    zone {
        maximum 10;
        reserved 5;
    }
    flow-session {
        maximum 2500;
        reserved 2000;
    }
    nat-nopat-address {
        maximum 120;
        reserved 100;
    }
    nat-static-rule {
        maximum 125;
        reserved 100;
    }
    logical-system ls-product-design;
```

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

Verification

To confirm that the security resources that you allocated for logical systems have been assigned to them, follow this procedure for each logical system and for all its resources.

- Verifying That Security Profile Resources Are Effectively Allocated for Logical Systems on page 63

Verifying That Security Profile Resources Are Effectively Allocated for Logical Systems

Purpose   Verify security resources for each logical system. Follow this process for all configured logical systems.

Action   1. Use Telnet to log in to each user logical system as its user logical system administrator.

Run Telnet, specifying the IP address of your SRX Series device. For example:

telnet 10.11.11.19

2. Enter the login ID and password for one of the user logical systems that you created

login: **lsmarketingadmin1**
password: **Talk2345**
lsmarketingadmin1@host:ls-marketing-dept>

3. Enter the following statement to Identify the resources configured for the profile.

lsmarketingadmin1@host:ls-marketing-dept> **show system security-profile ?**

4. Enter the following command at the resulting prompt. Do this for each feature configured for the profile.

lsmarketingadmin1@host:ls-marketing-dept> **show system security-profile zone detail**
logical system name : ls-marketing-dept
security profile name : ls-accnt-mrkt-profile
used amount : 0
reserved amount : 17
maximum quota : 22

Related Documentation
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- Understanding the Master Logical System and the Master Administrator Role on page 19

- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Configuring Master Logical System Security Features

Understanding Logical System Firewall Authentication

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Junos OS enables administrators to restrict and permit firewall users to access protected resources (different zones) behind a firewall based on their source IP address and other credentials.

The master administrator is responsible for configuring access profiles in the master logical system. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored. Access profiles configured at the master logical system are available to all user logical systems.

The master administrator configures the maximum and reserved numbers of firewall authentications for each user logical system. The user logical system administrator can then create firewall authentications in the user logical system. From a user logical system, the user logical system administrator can use the **show system security-profile auth-entry**

command to view the number of authentication resources allocated to the user logical system.

To configure the access profile, the master administrator uses the **profile** configuration statement at the [**edit access**] hierarchy level in the master logical system. The access profile can also include the order of authentication methods, LDAP or RADIUS server options, and session options.

The user logical system administrator can then associate the access profile with a security policy in the user logical system. The user logical system administrator also specifies the type of authentication:

- With pass-through authentication, a host or a user from one zone tries to access resources on another zone using an FTP, a Telnet, or an HTTP client. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.

- With Web authentication, users use HTTP to connect to an IP address on the device that is enabled for Web authentication and are prompted for the username and password. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

The user logical system administrator configures the following properties for firewall authentication in the user logical system:

- Security policy that specifies firewall authentication for matching traffic. Firewall authentication is specified with the **firewall-authentication** configuration statement at the [**edit security policies from-zone** *zone-name* **to-zone** *zone-name* **policy** *policy-name* **then permit**] hierarchy level.

  Users or user groups in an access profile who are allowed access by the policy can optionally be specified with the client-match configuration statement. (If no users or user groups are specified, any user who is successfully authenticated is allowed access.)

  For pass-through authentication, the access profile can optionally be specified and Web redirect (redirecting the client system to a webpage for authentication) can be enabled.

- Type of authentication (pass-through or Web authentication), default access profile, and success banner for the FTP, Telnet, or HTTP session. These properties are configured with the **firewall-authentication** configuration statement at the [**edit access**] hierarchy level.

- Host inbound traffic. Protocols, services, or both are allowed to access the logical system. The types of traffic are configured with the **host-inbound-traffic** configuration statement at the [**edit security zones security-zone** *zone-name*] or [**edit security zones security-zone** *zone-name* **interfaces** *interface-name*] hierarchy levels.

From a user logical system, the user logical system administrator can use the **show security firewall-authentication users** or **show security firewall-authentication history** commands to view the information about firewall users and history for the user logical system. From the master logical system, the master administrator can use the same

commands to view information for the master logical system, a specific user logical system, or all logical systems.

Example: Configuring Access Profiles (Master Administrators Only)

The master administrator is responsible for configuring access profiles in the master logical system. This example shows how to configure access profiles.

-
-
-

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

- Read *Firewall User Authentication Overview*.

Overview

This example configures an access profile for LDAP authentication for logical system users. This example creates the access profile described in Table 7 on page 67.

> *i* NOTE: The master administrator creates the access profile.

Table 7: Access Profile Configuration

Name	Configuration Parameters
ldap1	- LDAP is used as the first (and only) authentication method. - Base distinguished name: - Organizational unit name (OU): people - Domain components (DC): example, com - A user's LDAP distinguished name is assembled through the use of a common name identifier, username, and base distinguished name. The common name identifier is user ID (UID). - The LDAP server address is 10.155.26.104 and is reached through port 389.

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

> *i* NOTE: You must be logged in as the master administrator.

```
set access profile ldap1 authentication-order ldap
set access profile ldap1 ldap-options base-distinguished-name
    ou=people,dc=example,dc=com
set access profile ldap1 ldap-options assemble common-name uid
set access profile ldap1 ldap-server 10.155.26.104 port 389
```

Step-by-Step Procedure

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

To configure an access profile in the master logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Configure an access profile and set the authentication order.

   ```
   [edit access profile ldap1]
   admin@host# set authentication-order ldap
   ```

3. Configure LDAP options.

   ```
   [edit access profile ldap1]
   admin@host# set ldap-options base-distinguished-name
       ou=people,dc=example,dc=com
   admin@host# set ldap-options assemble common-name uid
   ```

4. Configure the LDAP server.

   ```
   [edit access profile ldap1]
   admin@host# set ldap-server 10.155.26.104 port 389
   ```

Results

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

```
admin@host# show access profile ldap1
authentication-order ldap;
ldap-options {
    base-distinguished-name ou=people,dc=example,dc=com;
    assemble {
```

```
            common-name uid;
        }
    }
    ldap-server {
        10.155.26.104 port 389;
    }
```

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

Related Documentation

- Example: Configuring Firewall Authentication for a User Logical System on page 111

- Understanding Logical System Firewall Authentication on page 65

- User Logical System Configuration Overview on page 23

Example: Configuring Security Features for the Master Logical System

This example shows how to configure security features, such as zones, policies, and firewall authentication, for the master logical system.

- Requirements on page 69

- Overview on page 69

- Configuration on page 70

- Verification on page 74

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 39.

- Use the **show system security-profile** command to see the resources allocated to the master logical system.

- Configure logical interfaces for the master logical system. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)" on page 140.

- Configure the access profile ldap1 in the master logical system. The ldap1 access profile is used for Web authentication of firewall users. See "Example: Configuring Access Profiles (Master Administrators Only)" on page 67.

Overview

In this example, you configure security features for the master logical system, called root-logical-system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40. This example configures the security features described in Table 8 on page 70.

Table 8: root-logical-system Security Feature Configuration

Feature	Name	Configuration Parameter
Zones	ls-root-trust	Bind to interface ge-0/0/4.0.
	ls-root-untrust	Bind to interface lt-0/0/0.1
Address books	root-internal	• Address masters: 12.12.1.0/24 • Attach to zone ls-root-trust
	root-external	• Address design: 12.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address marketing: 13.1.1.0/24 • Address set userlsys: design, accounting, marketing • Attach to zone ls-root-untrust
Security policies	permit-to-userlsys	Permit the following traffic: • From zone: ls-root-trust • To zone: ls-root-untrust • Source address: masters • Destination address: userlsys • Application: any
	permit-authorized-users	Permit the following traffic: • From zone: ls-root-untrust • To zone: ls-root-trust • Source address: userlsys • Destination address: masters • Application: junos-http, junos-https
Firewall authentication		• Web authentication • Authentication success banner "WEB AUTH LOGIN SUCCESS" • Default access profile ldap1
HTTP daemon		Activate on interface ge-0/0/4.0

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

```
set security address-book root-internal address masters 12.12.1.0/24
set security address-book root-internal attach zone ls-root-trust
set security address-book root-external address design 12.1.1.0/24
set security address-book root-external address accounting 14.1.1.0/24
set security address-book root-external address marketing 13.1.1.0/24
set security address-book root-external address-set userlsys address design
```

```
set security address-book root-external address-set userlsys address accounting
set security address-book root-external address-set userlsys address marketing
set security address-book root-external attach zone ls-root-untrust
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
    permit-to-userlsys match source-address masters
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
    permit-to-userlsys match destination-address userlsys
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
    permit-to-userlsys match application any
set security policies from-zone ls-root-trust to-zone ls-root-untrust policy
    permit-to-userlsys then permit
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
    permit-authorized-users match source-address userlsys
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
    permit-authorized-users match destination-address masters
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
    permit-authorized-users match application junos-http
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
    permit-authorized-users match application junos-https
set security policies from-zone ls-root-untrust to-zone ls-root-trust policy
    permit-authorized-users then permit firewall-authentication web-authentication
set security zones security-zone ls-root-trust interfaces ge-0/0/4.0
set security zones security-zone ls-root-untrust interfaces lt-0/0/0.1
set system services web-management http interface ge-0/0/4.0
set access firewall-authentication web-authentication default-profile ldap1
set access firewall-authentication web-authentication banner success "WEB AUTH
    LOGIN SUCCESS"
```

Step-by-Step Procedure

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

To configure zones and policies for the master logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Create security zones and assign interfaces to each zone.

   ```
   [edit security zones]
   admin@host# set security-zone ls-root-trust interfaces ge-0/0/4.0
   admin@host# set security-zone ls-root-untrust interfaces lt-0/0/0.1
   ```

3. Create address book entries.

   ```
   [edit security]
   admin@host# set address-book root-internal address masters 12.12.1.0/24
   admin@host# set address-book root-external address design 12.1.1.0/24
   admin@host# set address-book root-external address accounting 14.1.1.0/24
   admin@host# set address-book root-external address marketing 13.1.1.0/24
   admin@host# set address-book root-external address-set userlsys address design
   admin@host# set address-book root-external address-set userlsys address
       accounting
   ```

admin@host# **set address-book root-external address-set userlsys address marketing**

4. Attach address books to zones.

[edit security]
admin@host# **set address-book root-internal attach zone ls-root-trust**
admin@host# **set address-book root-external attach zone ls-root-untrust**

5. Configure a security policy that permits traffic from the ls-root-trust zone to the ls-root-untrust zone.

[edit security policies from-zone ls-root-trust to-zone ls-root-untrust]
admin@host# **set policy permit-to-userlsys match source-address masters**
admin@host# **set policy permit-to-userlsys match destination-address userlsys**
admin@host# **set policy permit-to-userlsys match application any**
admin@host# **set policy permit-to-userlsys then permit**

6. Configure a security policy that authenticates traffic from the ls-root-untrust zone to the ls-root-trust zone.

[edit security policies from-zone ls-root-untrust to-zone ls-root-trust]
admin@host# **set policy permit-authorized-users match source-address userlsys**
admin@host# **set policy permit-authorized-users match destination-address masters**
admin@host# **set policy permit-authorized-users match application junos-http**
admin@host# **set policy permit-authorized-users match application junos-https**
admin@host# **set policy permit-authorized-users then permit firewall-authentication web-authentication**

7. Configure the Web authentication access profile and define a success banner.

[edit access]
admin@host# **set firewall-authentication web-authentication default-profile ldap1**
admin@host# **set firewall-authentication web-authentication banner success "WEB AUTH LOGIN SUCCESS"**

8. Activate the HTTP daemon on the device.

[edit system]
admin@host# **set services web-management http interface ge-0/0/4.0**

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

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show security
...
address-book {
  root-internal {
    address masters 12.12.1.0/24;
    attach {
      zone ls-root-trust;
    }
```

```
            }
            root-external {
               address design 12.1.1.0/24;
               address accounting 14.1.1.0/24;
               address marketing 13.1.1.0/24;
               address-set userlsys {
                  address design;
                  address accounting;
                  address marketing;
               }
               attach {
                  zone ls-root-untrust;
               }
            }
         }
         policies {
            from-zone ls-root-trust to-zone ls-root-untrust {
               policy permit-to-userlsys {
                  match {
                     source-address masters;
                     destination-address userlsys;
                     application any;
                  }
                  then {
                     permit;
                  }
               }
            }
            from-zone ls-root-untrust to-zone ls-root-trust {
               policy permit-authorized-users {
                  match {
                     source-address userlsys;
                     destination-address masters;
                     application [ junos-http junos-https ];
                  }
                  then {
                     permit {
                        firewall-authentication {
                           web-authentication;
                        }
                     }
                  }
               }
            }
         }
         zones {
            security-zone ls-root-trust {
               interfaces {
                  ge-0/0/4.0;
               }
            }
            security-zone ls-root-untrust {
               interfaces {
                  lt-0/0/0.1;
               }
            }
```

```
    }
  [edit]
  admin@host# show access
  ...
  firewall-authentication {
    web-authentication {
      default-profile ldap1;
      banner {
        success "WEB AUTH LOGIN SUCCESS";
      }
    }
  }
  [edit]
  admin@host# show system services
  web-management {
    http {
      interface ge-0/0/4.0;
    }
  }
```

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

Verification

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

- Verifying Policy Configuration on page 74

Verifying Policy Configuration

Purpose Verify information about policies and rules.

Action From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation
- Understanding Logical System Zones on page 97
- Understanding Logical System Security Policies on page 104
- Understanding Logical System Firewall Authentication on page 65

IDP in Logical Systems Overview

A Junos OS Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a logical system.

This topic includes the following sections:

- IDP Policies on page 75
- IDP Installation and Licensing for Logical Systems on page 75

IDP Policies

The master administrator configures IDP policies at the root level. Configuring an IDP policy for logical systems is similar to configuring an IDP policy on a device that is not configured for logical systems. This can include the configuration of custom attack objects.

NOTE: User logical system administrators cannot create or modify IDP policies for their user logical systems. Only the master administrator can create IDP policies and bind them to user logical systems through a logical systems security profile.

NOTE: The user logical system administrator can create security zones in the user logical system and assign interfaces to each security zone. Zones that are specific to user logical systems cannot be referenced in IDP policies configured by the master administrator. The master administrator can reference zones in the master logical system in an IDP policy configured for the master logical system.

The master administrator then specifies an IDP policy in the security profile that is bound to a logical system. To enable IDP in a logical system, the master administrator or user logical system administrator configures a security policy that defines the traffic to be inspected and specifies the **permit application-services idp** action.

Although the master administrator can configure multiple IDP policies, a logical system can have only one active IDP policy at a time. For user logical systems, the master administrator can either bind the same IDP policy to multiple user logical systems or bind a unique IDP policy to each user logical system. To specify the active IDP policy for the master logical system, the master administrator can *either* reference the IDP policy in the security profile that is bound to the master logical system or use the **active-policy** configuration statement at the [**edit security idp**] hierarchy level.

NOTE: A commit error is generated if an IDP policy is both configured in the security profile that is bound to the master logical system and specified with the **active-policy** configuration statement. Use only one method to specify the active IDP policy for the master logical system.

IDP Installation and Licensing for Logical Systems

A single IDP security package is installed for all logical systems on the device. The download and install options can only be executed at the root level. The same version of the IDP attack database is shared by all logical systems.

An idp-sig license must be installed at the root level. Once IDP is enabled at the root level, it can be used with any logical system on the device.

- Understanding IDP Features in Logical Systems on page 76

- Example: Configuring an IDP Policy for a User Logical System on page 119

- Example: Configuring an IDP Policy for the Master Logical System on page 79

- User Logical System Configuration Overview on page 23

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- *IDP Policies Overview*

Understanding IDP Features in Logical Systems

This topic includes the following sections:

- Rulebases on page 76
- Protocol Decoders on page 76
- SSL Inspection on page 77
- Inline Tap Mode on page 77
- Multi-Detectors on page 77
- Logging and Monitoring on page 77

Rulebases

A single IDP policy can contain only one instance of any type of rulebase. The following IDP rulebases are supported for logical systems:

- The Intrusion prevention system (IPS) rulebase uses attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies.

- The application-level distributed denial-of-service (DDoS) rulebase defines parameters to protect servers such as DNS or HTTP. The application-level DDoS rulebase defines the source match condition for traffic that should be monitored and takes an action, such as drop the connection, drop the packet, or no action. It can also perform actions against future connections that use the same IP address.

> NOTE: Status monitoring for IPS and application-level DDoS is global to the device and not on a per logical system basis.

Protocol Decoders

The Junos IDP module ships with a set of preconfigured protocol decoders. These protocol decoders have default settings for various protocol-specific contextual checks that they perform. The IDP protocol decoder configuration is global and applies to all logical systems. Only the master administrator at the root level can modify the settings at the [**edit security idp sensor-configuration**] hierarchy level.

SSL Inspection

IDP SSL inspection uses the Secure Sockets Layer (SSL) protocol suite to enable inspection of HTTP traffic encrypted in SSL.

SSL inspection configuration is global and applies to all logical systems on a device. SSL inspection can only be configured by the master administrator at the root level with the **ssl-inspection** configuration statement at the [**edit security idp sensor-configuration**] hierarchy level.

Inline Tap Mode

The inline tap mode feature provides passive, inline detection of Application Layer threats for traffic matching security policies that have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module. This allows the packets to get to the next service module without waiting for IDP processing results.

Inline tap mode is enabled or disabled for all logical systems at the root level by the master administrator. To enable inline tap mode, use the **inline-tap** configuration statement at the [**edit security forwarding-process application-services maximize-idp-sessions**] hierarchy level. Delete the inline tap mode configuration to switch the device back to regular mode.

> *i* NOTE: The device must be restarted when switching to inline tap mode or back to regular mode.

Multi-Detectors

When a new IDP security package is received, it contains attack definitions and a detector. After a new policy is loaded, it is also associated with a detector. If the policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

The version of the detector is common to all logical systems.

Logging and Monitoring

Status monitoring options are available to the master administrator only. All status monitoring options under the **show security idp** and **clear security idp** CLI operational commands present global information, but not on a per logical system basis.

> *i* NOTE: SNMP monitoring for IDP is not supported on logical systems.

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled.

The logical systems identification is added to the following types of IDP traffic processing logs:

- Attack logs. The following example shows an attack log for the ls-product-design logical system:

```
Oct 12 17:33:32 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In
ls-product-design at 1286930013, SIG Attack log <4.0.0.1/34327->5.0.0.1/21>
 for TCP protocol and service SERVICE_IDP application NONE by rule 1 of
rulebase IPS in policy Recommended. attack: repeat=0, action=IGNORE,
threat-severity=MEDIUM, name=FTP:USER:ROOT, NAT <0.0.0.0:0->0.0.0.0:0>,
time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0,
intf:ls-product-design-untrust:ge-0/0/0.0->ls-product-design-trust:ge-0/0/1.0,
 packet-log-id: 65535 and misc-message -
```

- IP action logs. The following example shows an IP action log for the ls-product-design logical system:

```
Oct 13 16:56:04 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In
ls-product-design at 1287014163, TRAFFIC Attack log
<25.0.0.1/34802->15.0.0.1/21> for TCP protocol and service SERVICE_NONE
application NONE by rule 1 of rulebase IPS in policy Recommended. attack:
repeat=0, action=TRAFFIC_IPACTION_NOTIFY, threat-severity=INFO, name=_, NAT
 <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0,
 outpackets=0,
intf:ls-product-design-trust:ge-0/0/1.0->ls-product-design-untrust:plt0.3,
 packet-log-id: 0 and misc-message -
```

- Application DDoS logs. The following example shows an application DDoS log for the ls-product-design logical system:

```
Oct 11 16:29:57 8.0.0.254 RT_IDP: IDP_APPDDOS_APP_ATTACK_EVENT_LS: DDOS
Attack in ls-product-design at 1286839797 on my-http,
<ls-product-design-untrust:ge-0/0/0.0:4.0.0.1:33738->ls-product-design-trust:ge-0/0/1.0:5.0.0.1:80>
 for TCP protocol and service HTTP by rule 1 of rulebase DDOS in policy
Recommended. attack: repeats 0 action DROP threat-severity INFO,
connection-hit-rate 0,  context-name http-url-parsed, hit-rate 6,
value-hit-rate 6 time-scope PEER time-count 2  time-period 10 secs, context
 value:  ascii: /abc.html hex: 2f 61 62 63 2e 68 74 6d 6c
```

Related Documentation

- *Understanding IDP Policy Rule Bases*
- *Understanding IDP Protocol Decoders*
- *IDP SSL Overview*
- *Understanding IDP Inline Tap Mode*
- *Understanding Multiple IDP Detector Support*
- *Understanding IDP Logging*

Example: Configuring an IDP Policy for the Master Logical System

This example shows how to configure an IDP policy in a master logical system.

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

- Read "IDP in Logical Systems Overview" on page 74.

- Use the **show system security-profile** command to see the resources allocated to the master logical system.

Overview

In this example you configure a custom attack that is used in an IDP policy. The IDP policy is specified in a security profile that is applied to the master logical system. IDP is then enabled in a security policy configured in the master logical system.

You configure the features described in Table 9 on page 79.

Table 9: IDP Configuration for the Master Logical System

Feature	Name	Configuration Parameters
Custom attack	http-bf	- Severity critical - Detect three attacks between source and destination addresses of sessions. - Stateful signature attack type with the following characteristics: - location http-url-parsed - pattern .*juniper.* - client to server traffic
IPS rulebase policy	root-idp-policy	Match: - application default - http-bf custom attacks Action: - drop-connection - notification log-attacks

Table 9: IDP Configuration for the Master Logical System *(continued)*

Feature	Name	Configuration Parameters
Logical system security profile	master-profile (previously configured and applied to root-logical-system)	Add IDP policy root-idp-policy.
Security policy	enable-idp	Enable IDP in a security policy that matches any traffic from the lsys-root-untrust zone to the lsys-root-trust zone.

> **NOTE:** A logical system can have only one active IDP policy at a time. To specify the active IDP policy for the master logical system, the master administrator can reference the IDP policy in the security profile that is bound to the master logical system as shown in this example. Alternatively, the master administrator can use the **active-policy** configuration statement at the [edit security idp] hierarchy level.
>
> A commit error is generated if an IDP policy is both configured in the security profile that is bound to the master logical system and specified with the **active-policy** configuration statement. Use only one method to specify the active IDP policy for the master logical system.

Configuration

- Configuring a Custom Attack on page 80
- Configuring an IDP Policy for the Master Logical System on page 82
- Enabling IDP in a Security Policy on page 83

Configuring a Custom Attack

CLI Quick Configuration

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

```
set security idp custom-attack http-bf severity critical
set security idp custom-attack http-bf time-binding count 3
set security idp custom-attack http-bf time-binding scope peer
set security idp custom-attack http-bf attack-type signature context http-url-parsed
set security idp custom-attack http-bf attack-type signature pattern .*juniper.*
set security idp custom-attack http-bf attack-type signature direction client-to-server
```

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

To configure a custom attack object:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   [edit]
   admin@host> configure
   admin@host#
   ```

2. Create the custom attack object and set the severity level.

   ```
   [edit security idp]
   admin@host# set custom-attack http-bf severity critical
   ```

3. Configure attack detection parameters.

   ```
   [edit security idp]
   admin@host# set custom-attack http-bf time-binding count 3
   admin@host# set custom-attack http-bf time-binding scope peer
   ```

4. Configure stateful signature parameters.

   ```
   [edit security idp]
   admin@host# set custom-attack http-bf attack-type signature context
       http-url-parsed
   admin@host# set custom-attack http-bf attack-type signature pattern .*juniper.*
   admin@host# set custom-attack http-bf attack-type signature direction
       client-to-server
   ```

Results   From configuration mode, confirm your configuration by entering the **show security idp custom-attack http-bf** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp custom-attack http-bf
severity critical;
  time-binding {
    count 3;
    scope peer;
  }
    attack-type {
      signature {
        context http-url-parsed;
        pattern .*juniper.*;
        direction client-to-server;
      }
    }
```

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

Configuring an IDP Policy for the Master Logical System

CLI Quick Configuration

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

```
set security idp idp-policy root-idp-policy rulebase-ips rule 1 match application default
set security idp idp-policy root-idp-policy rulebase-ips rule 1 match attacks custom-attacks
    http-bf
set security idp idp-policy root-idp-policy rulebase-ips rule 1 then action drop-connection
set security idp idp-policy root-idp-policy rulebase-ips rule 1 then notification log-attacks
set system security-profile master-profile idp-policy lsys1-idp-policy
```

Step-by-Step Procedure

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

To configure an IDP policy:

1. Create the IDP policy and configure match conditions.

    ```
    [edit security idp]
    admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 match application
        default
    admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 match attacks
        custom-attacks http-bf
    ```

2. Configure actions for the IDP policy.

    ```
    [edit security idp]
    admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 then action
        drop-connection
    admin@host# set idp-policy root-idp-policy rulebase-ips rule 1 then notification
        log-attacks
    ```

3. Add the IDP policy to the security profile.

    ```
    [edit system security-profile master-profile]
    admin@host# set idp-policy lsys1-idp-policy
    ```

Results

From configuration mode, confirm your configuration by entering the **show security idp idp-policy root-idp-policy** and **show system security-profile master-profile** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
admin@host# show security idp idp-policy root-idp-policy
  rulebase-ips {
    rule 1 {
      match {
        application default;
        attacks {
          custom-attacks http-bf;
        }
```

```
              }
              then {
                action {
                  drop-connection;
                }
                notification {
                  log-attacks;
                }
              }
            }
          }
        }
      admin@host# show system security-profile master-profile
      ...
      idp-policy lsys1-idp-policy;
```

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

Enabling IDP in a Security Policy

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

> **set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp match source-address any**
> **set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp match destination-address any**
> **set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp match application any**
> **set security policies from-zone lsys-root-untrust to-zone lsys-root-trust policy enable-idp then permit application-services idp**

Step-by-Step Procedure   To enable IDP in a security policy:

1.   Create the security policy and configure match conditions.

   [edit security policies from-zone lsys-root-untrust to-zone lsys-root-trust]
   admin@host# **set policy enable-idp match source-address any**
   admin@host# **set policy enable-idp match destination-address any**
   admin@host# **set policy enable-idp match application any**

2.   Enable IDP.

   [edit security policies from-zone lsys-root-untrust to-zone lsys-root-trust]
   admin@host# **set policy enable-idp then permit application-services idp**

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

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
  admin@host# show security policies
  from-zone lsys-root-untrust to-zone lsys-root-trust {
    policy enable-idp {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit {
          application-services {
            idp;
          }
        }
      }
    }
  }
...
```

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

Verification

Verifying Attack Matches

Purpose   Verify that attacks are being matched in network traffic.

Action   From operational mode, enter the **show security idp attack table** command.

```
admin@host> show security idp attack table
  IDP attack statistics:
    Attack name                             #Hits
    http-bf                                     1
```

Related   • IDP in Logical Systems Overview on page 74
Documentation
              • SRX Series Logical System Master Administrator Configuration Tasks Overview on
                page 20

Understanding Logical System Application Identification Services

Predefined and custom application signatures identify an application by matching patterns in the first few packets of a session. Identifying applications provides the following benefits:

- Allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports.

- Improves performance by narrowing the scope of attack signatures for applications without decoders.

- Enables you to create detailed reports using AppTrack on applications passing through the device.

With logical systems, predefined and custom application signatures are global resources that are shared by all logical systems. The master administrator is responsible for downloading and installing predefined Juniper Networks application signatures and creating custom application and nested application signatures to identify applications that are not part of the predefined database.

Application identification is enabled by default.

The application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. Each user logical system has its own ASC. A user logical system administrator can display the ASC entries for their logical system with the **show services application-identification application-system-cache** command. A user logical system administrator can use the **clear services application-identification application-system-cache** command to clear the ASC entries for their logical system.

The master administrator can display or clear ASC entries for any logical system. The master administrator can also display or clear global counters with the **show services application-identification counter** and **clear services application-identification counter** commands.

Related
Documentation

- *Understanding the Junos OS Application Identification Database*

- *Example: Scheduling the Application Signature Package Updates*

- *Example: Configuring Junos OS Application Identification Custom Application Signatures*

- *Understanding IDP Application Identification*

- *Understanding the Application System Cache*

- *Verifying Application System Cache Statistics*

Understanding Logical System Application Firewall Services

An application firewall enables administrators of logical systems to create security policies for traffic based on application identification defined by application signatures. The application firewall provides additional security protection against dynamic-application traffic that might not be adequately controlled by standard network firewall policies. The application firewall controls information transmission by allowing or blocking traffic originating from particular applications.

To configure an application firewall, you define a rule set that contains rules specifying the action to be taken on identified dynamic applications. The rule set is configured independently and assigned to a security policy. Each rule set contains at least two rules, a matched rule (consisting of match criteria and action) and a default rule.

- A matched rule defines the action to be taken on matching traffic. When traffic matches an application and other criteria specified in the rule, the traffic is allowed or blocked based on the action specified in the rule.

- A default rule is applied when traffic does not match any other rule in the rule set.

The master administrator can download a predefined application signature database from the Juniper Networks Security Engineering website or can define application signatures using the Junos OS configuration CLI. For more information about application identification and application signatures, see *AppSecure Services Feature Guide for Security Devices*.

Configuring an application firewall on a logical system is the same process as configuring an application firewall on a device that is not configured with logical systems. However, the application firewall applies only to the logical system for which it is configured. The master administrator can configure, enable, and monitor application firewalls on the master logical system and all user logical systems on a device. User logical system administrators can configure, enable, and monitor application firewalls only on the user logical systems for which they have access.

Related Documentation
- Example: Configuring Application Firewall Services for a Master Logical System on page 86
- Example: Configuring Application Firewall Services for a User Logical System on page 124

Example: Configuring Application Firewall Services for a Master Logical System

This example describes how to configure application firewall services on the master, or root, logical system by a master administrator. Only the master administrator can configure, manage, and view configuration of the master logical system, in addition to all user logical systems.

After configuring application firewall rule sets and rules, the master administrator adds the application firewall rule set information to the security policy on the master logical system.

For information about configuring an application firewall within a security policy, see *Application Firewall Overview*.

- Requirements on page 86
- Overview on page 87
- Configuration on page 87
- Verification on page 89

Requirements

Before you begin:

- Verify that all interfaces, routing instances, and security zones have been configured on the master logical system.

  See "Example: Configuring Security Features for the Master Logical System" on page 69.

- Verify that application firewall resources (appfw-rule-set and appfw-rule) have been allocated in a security profile and bound to the master logical system through the [**system security-profile**] command. For application firewall resources, a security profile configuration allows 0 to 10,000 rule sets and 0 to 10,000 rules.

> **NOTE:** The master administrator allocates various global system resources through a security profile configuration which is then bound to the various logical systems on the device. The master administrator owns this function and configures the security profile for all user logical systems as well as the master logical system.
>
> For more information, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 51.

- Log in to the master logical system as the master administrator.

  For information about master administrator role functions, see "Understanding the Master Logical System and the Master Administrator Role" on page 19.

Overview

In this example you create application firewall services on the master logical system, called root-logical-system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example creates the following application firewall configuration:

- Rule set, root-rs1, with rules r1 and r2. When r1 is matched, telnet traffic is allowed through the firewall. When r2 is matched, web traffic is allowed through the firewall.

- Rule set, root-rs2, with rule r1. When r1 is matched, example2 traffic is blocked by the firewall.

All rule sets require a default rule, which specifies whether to permit or deny traffic that is not specified in any rules of a rule set. The default-rule action (permit or deny) must be the opposite from the action that is specified for the other rule(s) in the rule set.

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

```
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    rule r1 match dynamic-application junos:telnet
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    rule r1 then permit
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    rule r2 match dynamic-application-group junos:web
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    rule r2 then permit
set logical-systems root-logical-system security application-firewall rule-sets root-rs1
    default-rule deny
set logical-systems root-logical-system security application-firewall rule-sets root-rs2
    rule r1 match dynamic-application junos:facebook
```

```
set logical-systems root-logical-system security application-firewall rule-sets root-rs2
    rule r1 then deny
set logical-systems root-logical-system security application-firewall rule-sets root-rs2
    default-rule permit
```

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

To configure application firewall for a master logical system:

1. Log in to the master logical system as the master administrator. See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 39 and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Configure an application firewall rule set for root-logical-system.

   ```
   [edit ]
   admin@host# set logical-systems security application-firewall rule-sets root-rs1
   ```

3. Configure a rule for this rule set and specify which dynamic applications and dynamic application groups the rule should match.

   ```
   [edit]
   admin@host# set logical-systems security application-firewall rule-sets root-rs1
       rule r1 match dynamic-application telnet then permit
   ```

4. Configure the default rule for this rule set and specify the action to take when the identified dynamic application is not specified in any rules of the rule set.

   ```
   [edit]
   admin@host# set logical-systems security application-firewall rule-sets root-rs1
       default-rule deny
   ```

5. Repeat these steps to configure another rule set, root-rs2, if desired.

Results   From configuration mode, confirm your configuration by entering the **show security application-firewall rule-sets** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show security application-firewall rule-sets all
...
  application-firewall {
    rule-sets root-rs1 {
      rule r1 {
        match {
          dynamic-application [junos:telnet];
        }
        then {
```

```
            permit;
          }
        }
        default-rule {
          deny;
        }
      }
      rule-sets root-rs1 {
        rule r2 {
          match {
            dynamic-application-group [junos:web];
          }
          then {
            permit;
          }
        }
      rule-sets root-rs2 {
        rule r1 {
          match {
            dynamic-application [junos:FACEBOOK];
          }
          then {
            deny;
          }
        }
        default-rule {
          permit;
        }
      }
    }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying Application Firewall Configuration on page 89

Verifying Application Firewall Configuration

Purpose   View the application firewall configuration on the master logical system.

Action   From operational mode, enter the **show security application-firewall rule-set logical-system root-logical-system rule-set all** command.

```
admin@host> show security application-firewall rule-set logical-system root-logical-system
rule-set all

Rule-set: root-rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:Telnet
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
```

```
Number of sessions with appid pending: 2

Rule-set: root-rs1
    Logical system: root-logical-system
    Rule: r2
        Dynamic Applications: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Related Documentation

- SRX Series Logical System Master Administrator Configuration Tasks Overview on page 20

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- Understanding Logical System Application Firewall Services on page 85

- Example: Configuring Security Features for the Master Logical System on page 69

Understanding Logical System Application Tracking Services

AppTrack is an application tracking tool that provides statistics for analyzing bandwidth usage of your network. When enabled, AppTrack collects byte, packet, and duration statistics for application flows in the specified zone. By default, when each session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends it to the host device. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility.

AppTrack can be enabled and configured within any logical system. Configuring AppTrack in a logical system is the same as configuring AppTrack on a device that is not configured for logical systems. An AppTrack configuration only applies to the logical system in which it is configured. The name of the logical system is added to AppTrack logs. The master administrator can configure AppTrack for any logical system while a user logical system administrator can only configure AppTrack for the logical system that they are logged in to.

> NOTE: The system log configuration is global on the device and must be configured by the master administrator. The user logical system administrator cannot configure system logging for a logical system.

Counters keep track of the number of log messages sent and logs that have failed. AppTrack counters are global to the device. The master administrator as well as user logical system administrators can view AppTrack counters with the **show security application-tracking counters** command.

Related Documentation

- *Understanding AppTrack*

- *Example: Configuring AppTrack*

- Example: Configuring AppTrack for a User Logical System on page 129

Understanding Route-Based VPN Tunnels in Logical Systems

A VPN connection can secure traffic that passes between a logical system and a remote site across a WAN. With route-based VPNs, you configure one or more security policies in a logical system to regulate the traffic flowing through a single IP Security (IPsec) tunnel. For each IPsec tunnel, there is one set of IKE and IPsec security associations (SAs) that must be configured at the root level by the master administrator.

> NOTE: Only route-based VPNs are supported for logical systems. Policy-based VPNs are not supported.

In addition to configuring IKE and IPsec SAs for each VPN, the master administrator must also assign a secure tunnel (st0) interface to a user logical system. An st0 interface can only be assigned to a single user logical system. However, multiple user logical systems can each be assigned their own st0 interface.

> NOTE: The st0 unit 0 interface should not be assigned to a logical system, as an SA cannot be set up for this interface.

The user logical system administrator can configure the IP address and other attributes of the st0 interface assigned to the user logical system. The user logical system administrator cannot delete an st0 interface assigned to their user logical system.

For route-based VPNs, a security policy refers to a destination address and not a specific VPN tunnel. For cleartext traffic in a user logical system to be sent to the VPN tunnel for encapsulation, the user logical system administrator must make the following configurations:

- Security policy that permits traffic to a specified destination.

- Static route to the destination with the st0 interface as the next hop.

When Junos OS looks up routes in the user logical system to find the interface to use to send traffic to the destination address, it finds a static route through the st0 interface. Traffic is routed to the VPN tunnel as long as the security policy action is permit.

The master logical system and a user logical system can share a route-based VPN tunnel. An st0 interface assigned to a user logical system can also be used by the master logical system. For the master logical system, the master administrator configures a security policy that permits traffic to the remote destination and a static route to the remote destination with the st0 interface as the next hop.

VPN monitoring is configured by the master administrator in the master logical system. For the VPN monitor source interface, the master administrator must specify the st0 interface; a physical interface for a user logical system cannot be specified.

Related Documentation

- *Understanding Route-Based IPsec VPNs*

- User Logical System Configuration Overview on page 23

- Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only) on page 92

- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 133

Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)

The master administrator is responsible for assigning an st0 interface to a user logical system and configuring IKE and IPsec SAs at the root level for each VPN tunnel. This example shows how to assign an st0 interface to a user logical system and configure IKE and IPsec SA parameters.

- Requirements on page 92
- Overview on page 92
- Configuration on page 93
- Verification on page 96

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See ""Understanding the Master Logical System and the Master Administrator Role" on page 19.

- Read *Understanding Route-Based IPsec VPNs*.

Overview

In this example you configure a VPN tunnel for the ls-product-design user logical system. This example configures the VPN tunnel parameters described in Table 10 on page 92.

Table 10: Logical System VPN Tunnel Configuration

Feature	Name	Configuration Parameters
Tunnel interface	st0 unit 1	Assigned to ls-product-design logical system

Table 10: Logical System VPN Tunnel Configuration *(continued)*

Feature	Name	Configuration Parameters
IKE proposal	ike-phase1-proposal	• Preshared keys authentication • Diffie-Hellman group 2 • sha1 authentication algorithm • aes-128-cbc encryption algorithm
IKE policy		• Main mode • References IKE proposal ike-phase1-proposal • ASCII preshared key 395psksecr3t
IKE gateway	ike-gw	• External interface ge-0/0/3.0 • References IKE policy ike-phase1-policy • Address 2.2.2.2
IPsec proposal	ipsec-phase2-proposal	• ESP protocol • hmac-sha1-96 authentication algorithm • aes-128-cbc encryption algorithm
IPsec policy	vpn-policy1	• References ipsec-phase2-proposal • perfect-forward-secrecy keys group2
VPN	ike-vpn	• bind-interface st0.1 • References ike-gw gateway • References vpn-policy1 policy
VPN monitoring		For ike-vpn VPN: • source-interface st0.1 • destination-ip 4.0.0.1

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

```
set logical-systems ls-product-design interfaces st0 unit 1
set security ike proposal ike-phase1-proposal authentication-method pre-shared-keys
set security ike proposal ike-phase1-proposal dh-group group2
set security ike proposal ike-phase1-proposal authentication-algorithm sha1
set security ike proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
set security ike policy ike-phase1-policy mode main
set security ike policy ike-phase1-policy proposals ike-phase1-proposal
set security ike policy ike-phase1-policy pre-shared-key ascii-text "$ABC123"
set security ike gateway ike-gw ike-policy ike-phase1-policy
set security ike gateway ike-gw address 2.2.2.2
set security ike gateway ike-gw external-interface ge-0/0/3.0
set security ipsec proposal ipsec-phase2-proposal protocol esp
set security ipsec proposal ipsec-phase2-proposal authentication-algorithm hmac-sha1-96
```

```
set security ipsec proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
set security ipsec policy vpn-policy1 perfect-forward-secrecy keys group2
set security ipsec policy vpn-policy1 proposals ipsec-phase2-proposal
set security ipsec vpn ike-vpn bind-interface st0.1
set security ipsec vpn ike-vpn vpn-monitor source-interface st0.1
set security ipsec vpn ike-vpn vpn-monitor destination-ip 4.0.0.1
set security ipsec vpn ike-vpn ike gateway ike-gw
set security ipsec vpn ike-vpn ike ipsec-policy vpn-policy1
```

Step-by-Step Procedure

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

To assign a VPN tunnel interface to a user logical system and configure IKE and IPsec SAs:

1. Log in to the master logical system as the master administrator and enter configuration mode.

    ```
    [edit]
    admin@host> configure
    admin@host#
    ```

2. Assign a VPN tunnel interface.

    ```
    [edit logical-systems ls-product-design]
    admin@host# set interfaces st0 unit 1
    ```

3. Configure an IKE proposal.

    ```
    [edit security ike]
    admin@host# set proposal ike-phase1-proposal authentication-method
        pre-shared-keys
    admin@host# set proposal ike-phase1-proposal dh-group group2
    admin@host# set proposal ike-phase1-proposal authentication-algorithm sha1
    admin@host# set proposal ike-phase1-proposal encryption-algorithm aes-128-cbc
    ```

4. Configure an IKE policy.

    ```
    [edit security ike]
    admin@host# set policy ike-phase1-policy mode main
    admin@host# set policy ike-phase1-policy proposals ike-phase1-proposal
    admin@host# set policy ike-phase1-policy pre-shared-key ascii-text 395psksecr3t
    ```

5. Configure an IKE gateway.

    ```
    [edit security ike]
    admin@host# set gateway ike-gw external-interface ge-0/0/3.0
    admin@host# set gateway ike-gw ike-policy ike-phase1-policy
    admin@host# set gateway ike-gw address 2.2.2.2
    ```

6. Configure an IPsec proposal.

    ```
    [edit security ipsec]
    admin@host# set proposal ipsec-phase2-proposal protocol esp
    admin@host# set proposal ipsec-phase2-proposal authentication-algorithm
        hmac-sha1-96
    admin@host# set proposal ipsec-phase2-proposal encryption-algorithm aes-128-cbc
    ```

7. Configure an IPsec policy.

```
[edit security ipsec]
admin@host# set policy vpn-policy1 proposals ipsec-phase2-proposal
admin@host# set policy vpn-policy1 perfect-forward-secrecy keys group2
```

8. Configure the VPN.

```
[edit security ipsec]
admin@host# set vpn ike-vpn bind-interface st0.1
admin@host# set vpn ike-vpn ike gateway ike-gw
admin@host# set vpn ike-vpn ike ipsec-policy vpn-policy1
```

9. Configure VPN monitoring.

```
[edit security ipsec]
admin@host# set vpn ike-vpn vpn-monitor source-interface st0.1
admin@host# set vpn ike-vpn vpn-monitor destination-ip 4.0.0.1
```

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

```
[edit]
admin@host# show interfaces
  st0 {
    unit 1;
  }
[edit]
admin@host# show security ike
  proposal ike-phase1-proposal {
    authentication-method pre-shared-keys;
    dh-group group2;
    authentication-algorithm sha1;
    encryption-algorithm aes-128-cbc;
  }
    policy ike-phase1-policy {
      mode main;
      proposals ike-phase1-proposal;
      pre-shared-key ascii-text "$ABC123"; ## SECRET-DATA
    }
    gateway ike-gw {
      ike-policy ike-phase1-policy;
      address 2.2.2.2;
      external-interface ge-0/0/3.0;
    }
[edit]
admin@host# show security ipsec
  proposal ipsec-phase2-proposal {
    protocol esp;
    authentication-algorithm hmac-sha1-96;
    encryption-algorithm aes-128-cbc;
  }
    policy vpn-policy1 {
      perfect-forward-secrecy {
        keys group2;
      }
      proposals ipsec-phase2-proposal;
```

```
    }
    vpn ike-vpn {
       bind-interface st0.1;
       vpn-monitor {
          source-interface st0.1;
          destination-ip 4.0.0.1;
       }
       ike {
          gateway ike-gw;
          ipsec-policy vpn-policy1;
       }
    }
```

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

Verification

Confirm that the configuration is working properly.

Verifying the Configuration

Purpose   Verify that the IKE and IPsec SA configuration is correct.

Action   From operational mode, enter the **show security ike** and **show security ipsec** commands.

Related
Documentation

- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 133

- Understanding Route-Based VPN Tunnels in Logical Systems on page 91

- User Logical System Configuration Overview on page 23

Configuring User Logical System Security Features

Understanding Logical System Zones

Security zones are logical entities to which one or more interfaces are bound. Security zones can be configured on the master logical system by the master administrator or on user logical systems by the user logical system administrator. On a logical system, the administrator can configure multiple security zones, dividing the network into network segments to which various security options can be applied.

The master administrator configures the maximum and reserved numbers of security zones for each user logical system. The user logical system administrator can then create

security zones in the user logical system and assign interfaces to each security zone. From a user logical system, the user logical system administrator can use the **show system security-profile zones** command to view the number of security zones allocated to the user logical system and the **show interfaces** command to view the interfaces allocated to the user logical system.

> ⓘ NOTE: The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of security zones applied to the master logical system. The number of zones configured in the master logical system count toward the maximum number of zones available on the device.

The master and user administrator can configure the following properties of a security zone in a logical system:

- Interfaces that are part of a security zone.

- Screen options—For every security zone, you can enable a set of predefined screen options that detect and block various kinds of traffic that the device determines as potentially harmful.

- TCP-Reset—When this feature is enabled, the system sends a TCP segment with the RESET flag set when traffic arrives that does not match an existing session and does not have the synchronize flag set.

- Host inbound traffic—This feature specifies the kinds of traffic that can reach the device from systems that are directly connected to its interfaces. You can configure these parameters at the zone level, in which case they affect all interfaces of the zone, or at the interface level. (Interface configuration overrides that of the zone.)

There are no preconfigured security zones in the master logical system or user logical system.

The management functional zone (MGT) can only be configured for the master logical system. There is only one management interface per device and that interface is allocated to the master logical system.

The **all** interface can only be assigned to a zone in the master logical system by the master administrator.

The user logical system administrator can configure and view all attributes for a security zone in a user logical system. All attributes of a security zone in a user logical system are also visible to the master administrator.

Related
Documentation
- Example: Configuring Zones for a User Logical System on page 99

- User Logical System Configuration Overview on page 23

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- Understanding Logical System Interfaces and Routing Instances on page 139

• *Security Zones and Interfaces Overview*

Example: Configuring Zones for a User Logical System

This example shows how to configure zones for a user logical system.

Requirements

Before you begin:

• Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 23.

• Use the **show system security-profile zones** command to see the zone resources allocated to the logical system.

• Logical interfaces for the user logical system must be configured. See "Example: Configuring Interfaces and Routing Instances for a User Logical System" on page 158.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example creates the zones and address books described in Table 11 on page 99.

Table 11: User Logical System Zone and Address Book Configuration

Feature	Name	Configuration Parameters
Zones	ls-product-design-trust	• Bind to interface ge-0/0/5.1. • TCP reset enabled.
	ls-product-design-untrust	• Bind to interface lt-0/0/0.3.
Address books	product-design-internal	• Address product-designers: 12.1.1.0/24 • Attach to zone ls-product-design-trust
	product-design-external	• Address marketing: 13.1.1.0/24 • Address accounting: 14.1.1.0/24 • Address others: 12.12.1.0/24 • Address set otherlsys: marketing, accounting • Attach to zone ls-product-design-untrust

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

```
set security address-book product-design-internal address product-designers 12.1.1.0/24
set security address-book product-design-internal attach zone ls-product-design-trust
set security address-book product-design-external address marketing 13.1.1.0/24
set security address-book product-design-external address accounting 14.1.1.0/24
set security address-book product-design-external address others 12.12.1.0/24
set security address-book product-design-external address-set otherlsys address
    marketing
set security address-book product-design-external address-set otherlsys address
    accounting
set security address-book product-design-external attach zone ls-product-design-untrust
set security zones security-zone ls-product-design-trust tcp-rst
set security zones security-zone ls-product-design-trust interfaces ge-0/0/5.1
set security zones security-zone ls-product-design-untrust interfaces lt-0/0/0.3
```

Step-by-Step Procedure

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

To configure zones in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security zone and assign it to an interface.

   ```
   [edit security zones]
   lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-trust
       interfaces ge-0/0/5.1
   ```

3. Configure the TCP-Reset parameter for the zone.

   ```
   [edit security zones security-zone ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set tcp-rst
   ```

4. Configure a security zone and assign it to an interface.

   ```
   [edit security zones]
   lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-untrust
       interfaces lt-0/0/0.3
   ```

5. Create global address book entries.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set address-book product-design-internal
       address product-designers 12.1.1.0/24
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external
       address marketing 13.1.1.0/24
   ```

> lsdesignadmin1@host:ls-product-design# set address-book product-design-external address accounting 14.1.1.0/24
> lsdesignadmin1@host:ls-product-design# set address-book product-design-external address others 12.12.1.0/24
> lsdesignadmin1@host:ls-product-design# set address-book product-design-external address-set otherlsys address marketing
> lsdesignadmin1@host:ls-product-design# set address-book product-design-external address-set otherlsys address accounting

6. Attach address books to zones.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set address-book product-design-internal attach zone ls-product-design-trust
   lsdesignadmin1@host:ls-product-design# set address-book product-design-external attach zone ls-product-design-untrust
   ```

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

```
lsdesignadmin1@host:ls-product-design# show security
  address-book {
    product-design-internal {
      address product-designers 12.1.1.0/24;
      attach {
        zone ls-product-design-trust;
      }
    }
    product-design-external {
      address marketing 13.1.1.0/24;
      address accounting 14.1.1.0/24;
      address others 12.12.1.0/24;
      address-set otherlsys {
        address marketing;
        address accounting;
      }
      attach {
        zone ls-product-design-untrust;
      }
    }
  }
    zones {
      security-zone ls-product-design-trust {
        tcp-rst;
        interfaces {
          ge-0/0/5.1;
        }
      }
      security-zone ls-product-design-untrust {
        interfaces {
          lt-0/0/0.3;
        }
      }
    }
```

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

- Understanding Logical System Zones on page 97
- User Logical System Configuration Overview on page 23

Understanding Logical System Screen Options

Junos OS screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone. Junos OS then applies firewall policies, which can contain content filtering and IDP components, to the traffic that passes the screen filters.

All screen options available on the device are available in each logical system. Each user logical system administrator can configure screen options for their user logical system. The master administrator can configure screen options for the master logical system as well as all user logical systems.

The user logical system administrator can configure and view all screen options in a user logical system. All screen options in a user logical system are visible to the master administrator.

- Example: Configuring Screen Options for a User Logical System on page 102
- User Logical System Configuration Overview on page 23
- *Attack Detection and Prevention Overview*

Example: Configuring Screen Options for a User Logical System

This example shows how to configure screen options for a user logical system.

- Requirements on page 102
- Overview on page 103
- Configuration on page 103

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Configure zones for the user logical system. See "Example: Configuring Zones for a User Logical System" on page 99.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

You can limit the number of concurrent sessions to the same destination IP address in a user logical system. Setting a destination-based session limit can ensure that Junos OS allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. When the number of concurrent connection requests to an IP address surpasses the limit, Junos OS blocks further connection attempts to that IP address. This example creates the screen options described in Table 12 on page 103.

Table 12: User Logical System Screen Options Configuration

Name	Configuration Parameters
limit-destination-sessions	• Limits concurrent connection requests to destination IPs to 80. • Applied to ls-product-design-untrust zone.

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

```
set security screen ids-option limit-destination-sessions limit-session destination-ip-based
    80
set security zones security-zone ls-product-design-untrust screen
    limit-destination-sessions
```

Step-by-Step Procedure

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

To configure destination-based session limits in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a screen option for a destination-based session limit.

   ```
   [edit security]
   lsdesignadmin1@host:ls-product-design# set screen ids-option
       limit-destination-sessions limit-session destination-ip-based 80
   ```

3. Set the security zone for the screen option.

   ```
   [edit security]
   ```

lsdesignadmin1@host:ls-product-design# **set zones security-zone
ls-product-design-untrust screen limit-destination-sessions**

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

For brevity, this **show** command output includes only the configuration that is relevant
to this example. Any other configuration on the system has been replaced with ellipses
(...).

```
lsdesignadmin1@host:ls-product-design# show security screen
ids-option limit-destination-sessions {
  limit-session {
    destination-ip-based 80;
  }
}
lsdesignadmin1@host:ls-product-design# show security zones
security-zone ls-product-design-trust {
  ...
}
  security-zone ls-product-design-untrust {
    screen limit-destination-sessions;
    ...
  }
```

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

Related
Documentation

Understanding Logical System Security Policies

Security Policies in Logical Systems

Security policies enforce rules for what traffic can pass through the firewall and actions
that need to take place on the traffic as it passes through the firewall. From the
perspective of security policies, traffic enters one security zone and exits another security
zone.

By default, a logical system denies all traffic in all directions, including intra-zone and
inter-zone directions. Through the creation of security policies, the logical system
administrator can control the traffic flow from zone to zone by defining the kinds of traffic
permitted to pass from specified sources to specified destinations.

Security policies can be configured in the master logical system and in user logical systems. Configuring a security policy in a logical system is the same as configuring a security policy on a device that is not configured for logical systems. Any security policies, policy rules, address books, applications and application sets, and schedulers created within a logical system are only applicable to that logical system. Only predefined applications and application sets, such as **junos-ftp**, can be shared between logical systems.

> NOTE: In a logical system, you cannot specify **global** as either the from-zone or the to-zone in a security policy.

The user logical system administrator can configure and view all attributes for security policies in a user logical system. All attributes of a security policy in a user logical system are also visible to the master administrator.

Application Timeouts

The application timeout value set for an application determines the session timeout. Application timeout behavior is the same in a logical system as at the root level. However, user logical system administrators can use predefined applications in security policies but cannot modify the timeout value of predefined applications. This is because the predefined applications are shared by the master logical system and all user logical systems, so the user logical system administrator is not allowed to change its behavior. Application timeout values are stored in the application entry database and in the corresponding logical system TCP and UDP port-based timeout tables.

If the application that is matched for the traffic has a timeout value, that timeout value is used. Otherwise, the lookup proceeds in the following order until an application timeout value is found:

1. The logical system TCP and UDP port-based timeout table is searched for a timeout value.

2. The root TCP and UDP port-based timeout table is searched for a timeout value.

3. The protocol-based default timeout table is searched for a timeout value.

Security Policy Allocation

The master administrator configures the maximum and reserved numbers of security policies for each user logical system. The user logical system administrator can then create security policies in the user logical system. From a user logical system, the user logical system administrator can use the **show system security-profile policy** command to view the number of security policies allocated to the user logical system.

> ⓘ **NOTE:** The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of security policies applied to the master logical system. The number of policies configured in the master logical system count toward the maximum number of policies available on the device.

Related Documentation

- Example: Configuring Security Policies in a User Logical System on page 106
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51
- User Logical System Configuration Overview on page 23
- *Security Policies Overview*
- *Understanding Policy Application Timeout Configuration and Lookup*

Example: Configuring Security Policies in a User Logical System

This example shows how to configure security policies for a user logical system.

- Requirements on page 106
- Overview on page 106
- Configuration on page 107
- Verification on page 109

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.
- Use the **show system security-profiles policy** command to see the security policy resources allocated to the logical system.
- Configure zones and address books. See "Example: Configuring Zones for a User Logical System" on page 99.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example configures the security policies described in Table 13 on page 107.

Table 13: User Logical System Security Policies Configuration

Name	Configuration Parameters
permit-all-to-otherlsys	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: product-designers • Destination address: otherlsys • Application: any
permit-all-from-otherlsys	Permit the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-designers • Application: any

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

```
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys match source-address product-designers
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys match destination-address otherlsys
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys match application any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy permit-all-to-otherlsys then permit
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys match source-address otherlsys
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys match destination-address product-designers
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys match application any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

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

To configure security policies in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   lsdesignadmin1@host:ls-product-design> **configure**

lsdesignadmin1@host:ls-product-design#

2. Configure a security policy that permits traffic from the ls-product-design-trust zone to the ls-product-design-untrust zone.

    [edit security policies from-zone ls-product-design-trust to-zone
        ls-product-design-untrust]
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-to-otherlsys match
        source-address product-designers**
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-to-otherlsys match
        destination-address otherlsys**
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-to-otherlsys match
        application any**
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-to-otherlsys then
        permit**

3. Configure a security policy that permits traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

    [edit security policies from-zone ls-product-design-untrust to-zone
        ls-product-design-trust]
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-from-otherlsys match
        source-address otherlsys**
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-from-otherlsys match
        destination-address product-designers**
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-from-otherlsys match
        application any**
    lsdesignadmin1@host:ls-product-design# **set policy permit-all-from-otherlsys then
        permit**

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

    lsdesignadmin1@host:ls-product-design# **show security policies**
    from-zone ls-product-design-trust to-zone ls-product-design-untrust {
        policy permit-all-to-otherlsys {
            match {
                source-address product-designers;
                destination-address otherlsys;
                application any;
            }
            then {
                permit;
            }
        }
    }
        from-zone ls-product-design-untrust to-zone ls-product-design-trust {
            policy permit-all-from-otherlsys {
                match {
                    source-address otherlsys;
                    destination-address product-designers;
                    application any;
                }
                then {
                    permit;
            }

```
            }
          }
        }
      }
```

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

Verification

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

- Verifying Policy Configuration on page 109

Verifying Policy Configuration

Purpose   Verify information about policies and rules.

Action   From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation
- Understanding Logical System Security Policies on page 104
- User Logical System Configuration Overview on page 23
- *Troubleshooting Security Policies*

Understanding Logical System Firewall Authentication

A firewall user is a network user who must provide a username and password for authentication when initiating a connection across the firewall. Junos OS enables administrators to restrict and permit firewall users to access protected resources (different zones) behind a firewall based on their source IP address and other credentials.

The master administrator is responsible for configuring access profiles in the master logical system. Access profiles store usernames and passwords of users or point to external authentication servers where such information is stored. Access profiles configured at the master logical system are available to all user logical systems.

The master administrator configures the maximum and reserved numbers of firewall authentications for each user logical system. The user logical system administrator can then create firewall authentications in the user logical system. From a user logical system, the user logical system administrator can use the **show system security-profile auth-entry** command to view the number of authentication resources allocated to the user logical system.

To configure the access profile, the master administrator uses the **profile** configuration statement at the [**edit access**] hierarchy level in the master logical system. The access profile can also include the order of authentication methods, LDAP or RADIUS server options, and session options.

The user logical system administrator can then associate the access profile with a security policy in the user logical system. The user logical system administrator also specifies the type of authentication:

- With pass-through authentication, a host or a user from one zone tries to access resources on another zone using an FTP, a Telnet, or an HTTP client. The device uses FTP, Telnet, or HTTP to collect username and password information, and subsequent traffic from the user or host is allowed or denied based on the result of this authentication.

- With Web authentication, users use HTTP to connect to an IP address on the device that is enabled for Web authentication and are prompted for the username and password. Subsequent traffic from the user or host to the protected resource is allowed or denied based on the result of this authentication.

The user logical system administrator configures the following properties for firewall authentication in the user logical system:

- Security policy that specifies firewall authentication for matching traffic. Firewall authentication is specified with the **firewall-authentication** configuration statement at the [**edit security policies from-zone** *zone-name* **to-zone** *zone-name* **policy** *policy-name* **then permit**] hierarchy level.

  Users or user groups in an access profile who are allowed access by the policy can optionally be specified with the client-match configuration statement. (If no users or user groups are specified, any user who is successfully authenticated is allowed access.)

  For pass-through authentication, the access profile can optionally be specified and Web redirect (redirecting the client system to a webpage for authentication) can be enabled.

- Type of authentication (pass-through or Web authentication), default access profile, and success banner for the FTP, Telnet, or HTTP session. These properties are configured with the **firewall-authentication** configuration statement at the [**edit access**] hierarchy level.

- Host inbound traffic. Protocols, services, or both are allowed to access the logical system. The types of traffic are configured with the **host-inbound-traffic** configuration statement at the [**edit security zones security-zone** *zone-name*] or [**edit security zones security-zone** *zone-name* **interfaces** *interface-name*] hierarchy levels.

From a user logical system, the user logical system administrator can use the **show security firewall-authentication users** or **show security firewall-authentication history** commands to view the information about firewall users and history for the user logical system. From the master logical system, the master administrator can use the same commands to view information for the master logical system, a specific user logical system, or all logical systems.

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- *Firewall User Authentication Overview*

Example: Configuring Firewall Authentication for a User Logical System

This example shows how to configure firewall authentication for a user logical system.

- Requirements on page 111
- Overview on page 111
- Configuration on page 112
- Verification on page 114

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Use the **show system security-profiles auth-entry** command to see the firewall authentication entries allocated to the logical system.

- Access profiles must be configured in the master logical system by the master administrator. See "Example: Configuring Access Profiles (Master Administrators Only)" on page 67.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

In this example, users in the ls-marketing-dept and ls-accounting-dept logical systems are required to authenticate when initiating certain connections to the product designers subnet. This example configures the firewall authentication described in Table 14 on page 112.

> *i*   NOTE: This example uses the access profile configured in "Example: Configuring Access Profiles (Master Administrators Only)" on page 67and address book entries configured in "Example: Configuring Zones for a User Logical System" on page 99.

Table 14: User Logical System Firewall Authentication Configuration

Feature	Name	Configuration Parameters
Security policy	permit-authorized-users NOTE: Policy lookup is performed in the order that the policies are configured. The first policy that matches the traffic is used. If you have previously configured a policy that permits traffic for the same from-zone, to-zone, source address, and destination address but with application **any**, the policy configured in this example would never be matched. (See "Example: Configuring Security Policies in a User Logical System" on page 106.) Therefore, this policy should be reordered so that it is checked first.	Permit firewall authentication for the following traffic: • From zone: ls-product-design-untrust • To zone: ls-product-design-trust • Source address: otherlsys • Destination address: product-engineers • Application: junos-h323 The ldap1 access profile is used for pass-through authentication.
Firewall authentication		• Pass-through authentication • HTTP login prompt "welcome" • Default access profile ldap1

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

```
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users match source-address otherlsys
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users match destination-address product-designers
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users match application junos-h323
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy permit-authorized-users then permit firewall-authentication pass-through
    access-profile ldap1
set access firewall-authentication pass-through default-profile ldap1
set access firewall-authentication pass-through http banner login "welcome"
```

Step-by-Step Procedure

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

To configure firewall authentication in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   lsdesignadmin1@host:ls-product-design> **configure**
   lsdesignadmin1@host:ls-product-design#

2. Configure a security policy that permits firewall authentication.

   [edit security policies from-zone ls-product-design-untrust to-zone
       ls-product-design-trust]

lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match
source-address otherlsys
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match
destination -address product-designers
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users match
application junos-h323
lsdesignadmin1@host:ls-product-design# set policy permit-authorized-users then
permit firewall-authentication pass-through access-profile ldap1

3. Reorder the security policies.

[edit]
lsdesignadmin1@host:ls-product-design# insert security policies from-zone
ls-product-design-untrust to-zone ls-product-design-trust policy
permit-authorized-users before policy permit-all-from-otherlsys

4. Configure firewall authentication.

[edit access firewall-authentication]
lsdesignadmin1@host:ls-product-design# set pass-through http banner login
"welcome"
lsdesignadmin1@host:ls-product-design# set pass-through default-profile ldap1

Results   From configuration mode, confirm your configuration by entering the **show security policies**
and **show access firewall-authentication** commands. If the output does not display the
intended configuration, repeat the configuration instructions in this example to correct
it.

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
  policy permit-all-to-otherlsys {
    match {
      source-address product-designers;
      destination-address otherlsys;
      application any;
    }
    then {
      permit;
    }
  }
}
  from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy permit-authorized-users {
      match {
        source-address otherlsys;
        destination-address product-designers;
        application junos-h323;
      }
      then {
        permit {
          firewall-authentication {
            pass-through {
              access-profile ldap1;
            }
          }
        }
```

```
          }
        }
        policy permit-all-from-otherlsys {
          match {
            source-address otherlsys;
            destination-address product-designers;
            application any;
          }
          then {
            permit;
          }
        }
      }
    }
lsdesignadmin1@host:ls-product-design# show access firewall-authentication
pass-through {
  default-profile ldap1;
  http {
    banner {
      login welcome;
    }
  }
}
```

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

Verification

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

- Verifying Firewall User Authentication and Monitoring Users and IP Addresses on page 114

Verifying Firewall User Authentication and Monitoring Users and IP Addresses

Purpose   Display firewall authentication user history and verify the number of firewall users who successfully authenticated and firewall users who failed to log in.

Action   From operational mode, enter these **show** commands.

lsdesignadmin1@host:ls-product-design> **show security firewall-authentication history**
lsdesignadmin1@host:ls-product-design> **show security firewall-authentication history
  identifier** *id*
lsdesignadmin1@host:ls-product-design> **show security firewall-authentication users**
lsdesignadmin1@host:ls-product-design> **show security firewall-authentication users
  identifier** *id*

**Related
Documentation**
- Example: Configuring Access Profiles (Master Administrators Only) on page 67

- Understanding Logical System Firewall Authentication on page 65

- User Logical System Configuration Overview on page 23

- *Example: Configuring Pass-Through Authentication*

IDP in Logical Systems Overview

A Junos OS Intrusion Detection and Prevention (IDP) policy enables you to selectively enforce various attack detection and prevention techniques on network traffic passing through a logical system.

This topic includes the following sections:

IDP Policies

The master administrator configures IDP policies at the root level. Configuring an IDP policy for logical systems is similar to configuring an IDP policy on a device that is not configured for logical systems. This can include the configuration of custom attack objects.

> **NOTE:** User logical system administrators cannot create or modify IDP policies for their user logical systems. Only the master administrator can create IDP policies and bind them to user logical systems through a logical systems security profile.

> **NOTE:** The user logical system administrator can create security zones in the user logical system and assign interfaces to each security zone. Zones that are specific to user logical systems cannot be referenced in IDP policies configured by the master administrator. The master administrator can reference zones in the master logical system in an IDP policy configured for the master logical system.

The master administrator then specifies an IDP policy in the security profile that is bound to a logical system. To enable IDP in a logical system, the master administrator or user logical system administrator configures a security policy that defines the traffic to be inspected and specifies the **permit application-services idp** action.

Although the master administrator can configure multiple IDP policies, a logical system can have only one active IDP policy at a time. For user logical systems, the master administrator can either bind the same IDP policy to multiple user logical systems or bind a unique IDP policy to each user logical system. To specify the active IDP policy for the master logical system, the master administrator can *either* reference the IDP policy in the security profile that is bound to the master logical system or use the **active-policy** configuration statement at the [**edit security idp**] hierarchy level.

> ℹ️ NOTE: A commit error is generated if an IDP policy is both configured in the security profile that is bound to the master logical system and specified with the active-policy configuration statement. Use only one method to specify the active IDP policy for the master logical system.

IDP Installation and Licensing for Logical Systems

A single IDP security package is installed for all logical systems on the device. The download and install options can only be executed at the root level. The same version of the IDP attack database is shared by all logical systems.

An idp-sig license must be installed at the root level. Once IDP is enabled at the root level, it can be used with any logical system on the device.

Related Documentation
- Understanding IDP Features in Logical Systems on page 76
- Example: Configuring an IDP Policy for a User Logical System on page 119
- Example: Configuring an IDP Policy for the Master Logical System on page 79
- User Logical System Configuration Overview on page 23
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51
- *IDP Policies Overview*

Understanding IDP Features in Logical Systems

This topic includes the following sections:

- Rulebases on page 116
- Protocol Decoders on page 117
- SSL Inspection on page 117
- Inline Tap Mode on page 117
- Multi-Detectors on page 117
- Logging and Monitoring on page 118

Rulebases

A single IDP policy can contain only one instance of any type of rulebase. The following IDP rulebases are supported for logical systems:

- The Intrusion prevention system (IPS) rulebase uses attack objects to detect known and unknown attacks. It detects attacks based on stateful signature and protocol anomalies.
- The application-level distributed denial-of-service (DDoS) rulebase defines parameters to protect servers such as DNS or HTTP. The application-level DDoS rulebase defines the source match condition for traffic that should be monitored and takes an action,

such as drop the connection, drop the packet, or no action. It can also perform actions against future connections that use the same IP address.

> *i* NOTE: Status monitoring for IPS and application-level DDoS is global to the device and not on a per logical system basis.

Protocol Decoders

The Junos IDP module ships with a set of preconfigured protocol decoders. These protocol decoders have default settings for various protocol-specific contextual checks that they perform. The IDP protocol decoder configuration is global and applies to all logical systems. Only the master administrator at the root level can modify the settings at the [**edit security idp sensor-configuration**] hierarchy level.

SSL Inspection

IDP SSL inspection uses the Secure Sockets Layer (SSL) protocol suite to enable inspection of HTTP traffic encrypted in SSL.

SSL inspection configuration is global and applies to all logical systems on a device. SSL inspection can only be configured by the master administrator at the root level with the **ssl-inspection** configuration statement at the [**edit security idp sensor-configuration**] hierarchy level.

Inline Tap Mode

The inline tap mode feature provides passive, inline detection of Application Layer threats for traffic matching security policies that have the IDP application service enabled. When a device is in inline tap mode, packets pass through firewall inspection and are also copied to the independent IDP module. This allows the packets to get to the next service module without waiting for IDP processing results.

Inline tap mode is enabled or disabled for all logical systems at the root level by the master administrator. To enable inline tap mode, use the **inline-tap** configuration statement at the [**edit security forwarding-process application-services maximize-idp-sessions**] hierarchy level. Delete the inline tap mode configuration to switch the device back to regular mode.

> *i* NOTE: The device must be restarted when switching to inline tap mode or back to regular mode.

Multi-Detectors

When a new IDP security package is received, it contains attack definitions and a detector. After a new policy is loaded, it is also associated with a detector. If the policy being loaded has an associated detector that matches the detector already in use by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded

along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection.

The version of the detector is common to all logical systems.

Logging and Monitoring

Status monitoring options are available to the master administrator only. All status monitoring options under the **show security idp** and **clear security idp** CLI operational commands present global information, but not on a per logical system basis.

> *i* NOTE: SNMP monitoring for IDP is not supported on logical systems.

IDP generates event logs when an event matches an IDP policy rule in which logging is enabled.

The logical systems identification is added to the following types of IDP traffic processing logs:

- Attack logs. The following example shows an attack log for the ls-product-design logical system:

```
Oct 12 17:33:32 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In
ls-product-design at 1286930013, SIG Attack log <4.0.0.1/34327->5.0.0.1/21>
 for TCP protocol and service SERVICE_IDP application NONE by rule 1 of
rulebase IPS in policy Recommended. attack: repeat=0, action=IGNORE,
threat-severity=MEDIUM, name=FTP:USER:ROOT, NAT <0.0.0.0:0->0.0.0.0:0>,
time-elapsed=0, inbytes=0, outbytes=0, inpackets=0, outpackets=0,
intf:ls-product-design-untrust:ge-0/0/0.0->ls-product-design-trust:ge-0/0/1.0,
 packet-log-id: 65535 and misc-message -
```

- IP action logs. The following example shows an IP action log for the ls-product-design logical system:

```
Oct 13 16:56:04 8.0.0.254 RT_IDP: IDP_ATTACK_LOG_EVENT_LS: IDP: In
ls-product-design at 1287014163, TRAFFIC Attack log
<25.0.0.1/34802->15.0.0.1/21> for TCP protocol and service SERVICE_NONE
application NONE by rule 1 of rulebase IPS in policy Recommended. attack:
repeat=0, action=TRAFFIC_IPACTION_NOTIFY, threat-severity=INFO, name=_, NAT
 <0.0.0.0:0->0.0.0.0:0>, time-elapsed=0, inbytes=0, outbytes=0, inpackets=0,
 outpackets=0,
intf:ls-product-design-trust:ge-0/0/1.0->ls-product-design-untrust:plt0.3,
 packet-log-id: 0 and misc-message -
```

- Application DDoS logs. The following example shows an application DDoS log for the ls-product-design logical system:

```
Oct 11 16:29:57 8.0.0.254 RT_IDP: IDP_APPDDOS_APP_ATTACK_EVENT_LS: DDOS
Attack in ls-product-design at 1286839797 on my-http,
<ls-product-design-untrust:ge-0/0/0.0:4.0.0.1:33738->ls-product-design-trust:ge-0/0/1.0:5.0.0.1:80>
 for TCP protocol and service HTTP by rule 1 of rulebase DDOS in policy
Recommended. attack: repeats 0 action DROP threat-severity INFO,
connection-hit-rate 0,  context-name http-url-parsed, hit-rate 6,
value-hit-rate 6 time-scope PEER time-count 2  time-period 10 secs, context
 value:  ascii: /abc.html hex: 2f 61 62 63 2e 68 74 6d 6c
```

Related Documentation
- *Understanding IDP Policy Rule Bases*
- *Understanding IDP Protocol Decoders*
- *IDP SSL Overview*
- *Understanding IDP Inline Tap Mode*
- *Understanding Multiple IDP Detector Support*
- *Understanding IDP Logging*

Example: Configuring an IDP Policy for a User Logical System

The master administrator can *either* download predefined IDP policies to the device or configure custom IDP policies at the root level using custom or predefined attack objects. The master administrator is responsible for assigning an IDP policy to a user logical system. This example shows how to assign a predefined IDP policy to a user logical system.

- Requirements on page 119
- Overview on page 119
- Configuration on page 120
- Verification on page 121

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

- Read *IDP Policies Overview*.

- Assign the ls-design-profile security policy to the ls-product-design user logical system. See "Example: Configuring Logical Systems Security Profiles (Master Administrators Only)" on page 56.

- Download predefined IDP policy templates to the device. See *Downloading and Using Predefined IDP Policy Templates (CLI Procedure)*.

> **NOTE:** Activating a predefined IDP policy with the active-policy configuration statement at the [edit security idp] hierarchy level only applies to the master logical system. For a user logical system, the master administrator specifies the active IDP policy in the security profile that is bound to the user logical system.

Overview

The predefined IDP policy named Recommended contains attack objects recommended by Juniper Networks. All rules in the policy have their actions set to take the recommended

action for each attack object. You add the Recommended IDP policy to the ls-design-profile, which is bound to the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

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

    set system security-profile ls-design-profile idp-policy Recommended

Step-by-Step Procedure

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

To add a predefined IDP policy to a security profile for a user logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

        [edit]
        admin@host> **configure**
        admin@host#

2. Add the IDP policy to the security profile.

        [edit system security-profile]
        admin@host# **set ls-design-profile idp-policy Recommended**

Results

From configuration mode, confirm your configuration by entering the **show security idp** and **show system security-profile ls-design-profile** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

        [edit]
        admin@host# **show security idp**
          idp-policy Recommended {
            ...
          }
        [edit]
        admin@host# **show system security-profile ls-design-profile**
          policy {
            ...
          }
          idp-policy Recommended;
        logical-system ls-product-design;

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

Verification

Verifying the Configuration

Purpose   Verify the IDP policy assigned to the logical system.

Action   From operational mode, enter the **show security idp logical-system policy-association** command. Ensure that the IDP policy in the security profile that is bound to the logical system is correct.

```
admin@host> show security idp logical-system policy-association
  Logical system        IDP policy
  ls-product-design      Recommended
```

Related   • Example: Enabling IDP in a User Logical System Security Policy on page 121
Documentation
• IDP in Logical Systems Overview on page 74

• User Logical System Configuration Overview on page 23

Example: Enabling IDP in a User Logical System Security Policy

This example shows how to enable IDP in a security policy in a user logical system.

• Requirements on page 121
• Overview on page 121
• Configuration on page 122
• Verification on page 123

Requirements

Before you begin:

• Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.

• Use the **show system security-profiles idp-policy** command to see the security policy resources allocated to the logical system.

• Configure an IDP security policy for the user logical system as the master administrator. See "Example: Configuring an IDP Policy for a User Logical System" on page 119.

Overview

In this example, you configure the ls-product-design user logical system as shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

You enable IDP in a security policy that matches any traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone. Enabling IDP in a security policy directs matching traffic to be checked against the IDP rulebases.

> (i) NOTE: This example uses the IDP policy configured and assigned to the ls-product-design user logical system by the master administrator in "Example: Configuring an IDP Policy for a User Logical System" on page 119.

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

```
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp match source-address any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp match destination-address any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp match application any
set security policies from-zone ls-product-design-untrust to-zone ls-product-design-trust
    policy enable-idp then permit application-services idp
```

Step-by-Step Procedure

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

To configure a security policy to enable IDP in a user logical system:

1. Log in to the logical system as the user logical system administrator and enter configuration mode.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design>configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security policy that matches traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

   ```
   [edit security policies from-zone ls-product-design-untrust to-zone
       ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set policy enable-idp match source-address
       any
   lsdesignadmin1@host:ls-product-design# set policy enable-idp match
       destination-address any
   lsdesignadmin1@host:ls-product-design# set policy enable-idp match application
       any
   ```

3. Configure the security policy to enable IDP for matching traffic.

   ```
   [edit security policies from-zone ls-product-design-untrust to-zone
       ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set policy enable-idp then permit
       application-services idp
   ```

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

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security policies
  from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy enable-idp {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit {
          application-services {
            idp;
          }
        }
      }
    }
    ...
  }
```

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

Verification

Verifying Attack Matches

Purpose   Verify that attacks are being matched in network traffic.

Action   From operational mode, enter the **show security idp attack table** command.

```
admin@host> show security idp attack table
  IDP attack statistics:
    Attack name                              #Hits
    FTP:USER:ROOT                            1
```

Related Documentation
- Example: Configuring an IDP Policy for a User Logical System on page 119
- IDP in Logical Systems Overview on page 74
- User Logical System Configuration Overview on page 23

Understanding Logical System Application Identification Services

Predefined and custom application signatures identify an application by matching patterns in the first few packets of a session. Identifying applications provides the following benefits:

- Allows Intrusion Detection and Prevention (IDP) to apply appropriate attack objects to applications running on nonstandard ports.

- Improves performance by narrowing the scope of attack signatures for applications without decoders.

- Enables you to create detailed reports using AppTrack on applications passing through the device.

With logical systems, predefined and custom application signatures are global resources that are shared by all logical systems. The master administrator is responsible for downloading and installing predefined Juniper Networks application signatures and creating custom application and nested application signatures to identify applications that are not part of the predefined database.

Application identification is enabled by default.

The application system cache (ASC) saves the mapping between an application type and the corresponding destination IP address, destination port, protocol type, and service. Each user logical system has its own ASC. A user logical system administrator can display the ASC entries for their logical system with the **show services application-identification application-system-cache** command. A user logical system administrator can use the **clear services application-identification application-system-cache** command to clear the ASC entries for their logical system.

The master administrator can display or clear ASC entries for any logical system. The master administrator can also display or clear global counters with the **show services application-identification counter** and **clear services application-identification counter** commands.

Related Documentation
- *Understanding the Junos OS Application Identification Database*

- *Example: Scheduling the Application Signature Package Updates*

- *Example: Configuring Junos OS Application Identification Custom Application Signatures*

- *Understanding IDP Application Identification*

- *Understanding the Application System Cache*

- *Verifying Application System Cache Statistics*

Example: Configuring Application Firewall Services for a User Logical System

This example describes how to configure application firewall services on a user logical system by a user logical system administrator. User logical system administrators can

manage and monitor their own system application firewall rule sets and rules and manage the dynamic applications allowed or blocked on their respective logical systems.

After configuring application firewall rule sets and rules, user logical system administrators add the application firewall rule set information to the security policy on their individual logical systems.

For information about configuring an application firewall within a security policy, see *Application Firewall Overview*.

- Requirements on page 125
- Overview on page 125
- Configuration on page 126
- Verification on page 127

Requirements

Before you begin:

- Verify that the security zones are configured for the user logical system.

- Verify that the master administrator has allocated application firewall resources (appfw-rule-set and appfw-rule) in the security profile bound to the user logical system.

  For more information, see "Understanding Logical System Security Profiles (Master Administrators Only)" on page 51.

- Log in to the logical system as the user logical system administrator.

  For information about user logical system administrator role functions, see "Understanding User Logical Systems and the User Logical System Administrator Role" on page 25.

Overview

In this example you configure application firewall services on the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example creates the following application firewall configuration:

- Rule set, ls-product-design-rs1, with rules r1 and r2. When r1 is matched, telnet traffic is allowed through the firewall. When r2 is matched, web traffic is allowed through the firewall.

- Rule set, ls-product-design-rs2, with rule r1. When r1 is matched, Facebook traffic is blocked by the firewall.

All rule sets require a default rule, which specifies whether to permit or deny traffic that is not specified in any rules of a rule set. The default-rule action (permit or deny) must be the opposite from the action that is specified for the other rule(s) in the rule set.

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

```
set security application-firewall rule-sets ls-product-design-rs1 rule r1 match
    dynamic-application junos:telnet
set security application-firewall rule-sets ls-product-design-rs1 rule r1 then permit
set security application-firewall rule-sets ls-product-design-rs1 rule r2 match
    dynamic-application-group junos:web
set security application-firewall rule-sets ls-product-design-rs1 rule r2 then permit
set security application-firewall rule-sets ls-product-design-rs1 default-rule deny
set security application-firewall rule-sets ls-product-design-rs2 rule r1 match
    dynamic-application junos:facebook
set security application-firewall rule-sets ls-product-design-rs2 rule r1 then deny
set security application-firewall rule-sets ls-product-design-rs2 default-rule permit
```

Step-by-Step Procedure

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

To configure application firewall for a user logical system:

1. Log in to the user logical system as the user logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure an application firewall rule set for this logical system.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets
       ls-product-design-rs1
   ```

3. Configure a rule for this rule set and specify which dynamic applications and dynamic application groups the rule should match.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets
       ls-product-design-rs1 rule r1 match dynamic-application telnet then permit
   ```

4. Configure the default rule for this rule set and specify the action to take when the identified dynamic application is not specified in any rules of the rule set.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design# set security application-firewall rule-sets
       ls-product-design-rs1 default-rule deny
   ```

5. Repeat these steps to configure another rule set, ls-product-design-rs2, if desired.

Results

From configuration mode, confirm your configuration by entering the **show security application-firewall rule-set all** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show security application-firewall rule-set all
...
  application-firewall {
    rule-sets ls-product-design-rs1 {
      rule r1 {
        match {
          dynamic-application [junos:telnet];
        }
        then {
          permit;
        }
      }
      default-rule {
        deny;
      }
    }
    rule-sets ls-product-design-rs1 {
      rule r2 {
        match {
          dynamic-application-group [junos:web];
        }
        then {
          permit;
        }
      }
      rule-sets ls-product-design-rs2 {
        rule r1 {
          match {
            dynamic-application [junos:FACEBOOK];
          }
          then {
            deny;
          }
        }
        default-rule {
          permit;
        }
      }
```

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

Verification

Confirm that the configuration is working properly.

Verifying Application Firewall Configuration

Purpose   View the application firewall configuration on the user logical system.

Action From operational mode, enter the **show security application-firewall rule-set all** command.

```
lsdesignadmin1@host:ls-product-design> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:Telnet
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Applications: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

Related Documentation
- User Logical System Configuration Overview on page 23
- Understanding Logical System Application Firewall Services on page 85

Understanding Logical System Application Tracking Services

AppTrack is an application tracking tool that provides statistics for analyzing bandwidth usage of your network. When enabled, AppTrack collects byte, packet, and duration statistics for application flows in the specified zone. By default, when each session closes, AppTrack generates a message that provides the byte and packet counts and duration of the session, and sends it to the host device. The Security Threat Response Manager (STRM) retrieves the data and provides flow-based application visibility.

AppTrack can be enabled and configured within any logical system. Configuring AppTrack in a logical system is the same as configuring AppTrack on a device that is not configured for logical systems. An AppTrack configuration only applies to the logical system in which it is configured. The name of the logical system is added to AppTrack logs. The master administrator can configure AppTrack for any logical system while a user logical system administrator can only configure AppTrack for the logical system that they are logged in to.

> **NOTE:** The system log configuration is global on the device and must be configured by the master administrator. The user logical system administrator cannot configure system logging for a logical system.

Counters keep track of the number of log messages sent and logs that have failed. AppTrack counters are global to the device. The master administrator as well as user logical system administrators can view AppTrack counters with the **show security application-tracking counters** command.

Related
Documentation

- *Understanding AppTrack*

- *Example: Configuring AppTrack*

- Example: Configuring AppTrack for a User Logical System on page 129

Example: Configuring AppTrack for a User Logical System

This example shows how to configure the AppTrack tracking tool so you can analyze the bandwidth usage of your network.

- Requirements on page 129
- Overview on page 129
- Configuration on page 129
- Verification on page 131

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.

- (Master administrator) Configure system logging in the master logical system. See Network Monitoring and Troubleshooting Guide for Security Devices.

Overview

This example shows how to enable application tracking for the security zone ls-product-design-trust in the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

The first message is generated at session start and update messages are sent every 5 minutes after that or until the session ends. A final message is sent at session end.

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

    set security zones security-zone ls-product-design-trust application-tracking
    set security application-tracking first-update

Step-by-Step Procedure

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

To configure AppTrack for a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

        lsdesignadmin1@host:ls-product-design> **configure**
        lsdesignadmin1@host:ls-product-design#

2. Enable AppTrack for the security zone.

        [edit security]
        lsdesignadmin1@host:ls-product-design# **set zones security-zone**
          ls-product-design-trust application-tracking

3. Generate update messages at session start and at 5-minute intervals.

        [edit security]
        lsdesignadmin1@host:ls-product-design# **set application-tracking first-update**

Results

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

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

        [edit]
        lsdesignadmin1@host:ls-product-design# **show security**
        ...
          application-tracking {
            first-update;
          }
        ...
          zones {
            security-zone ls-product-design-trust {
              ...
              application-tracking;
            }
          }

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

Verification

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

Verifying AppTrack Operation

Purpose   View the AppTrack counters periodically to monitor tracking.

Action   From operational mode, enter the **show application-tracking counters** command.

Verifying Security Flow Session Statistics

Purpose   Compare byte and packet counts in logged messages with the session statistics from the **show security flow session** command output.

Action   From operational mode, enter the **show security flow session** command.

Verifying Application System Cache Statistics

Purpose   Compare cache statistics such as IP address, port, protocol, and service for an application from the **show services application-identification application-system-cache** command output.

Action   From operational mode, enter the **show services application-identification application-system-cache** command.

Verifying the Status of Application Identification Counter Values

Purpose   Compare session statistics for application identification counter values from the **show services application-identification counter** command output.

Action   From operational mode, enter the **show services application-identification counter** command.

Related Documentation

Understanding Route-Based VPN Tunnels in Logical Systems

A VPN connection can secure traffic that passes between a logical system and a remote site across a WAN. With route-based VPNs, you configure one or more security policies

in a logical system to regulate the traffic flowing through a single IP Security (IPsec) tunnel. For each IPsec tunnel, there is one set of IKE and IPsec security associations (SAs) that must be configured at the root level by the master administrator.

> **NOTE:** Only route-based VPNs are supported for logical systems. Policy-based VPNs are not supported.

In addition to configuring IKE and IPsec SAs for each VPN, the master administrator must also assign a secure tunnel (st0) interface to a user logical system. An st0 interface can only be assigned to a single user logical system. However, multiple user logical systems can each be assigned their own st0 interface.

> **NOTE:** The st0 unit 0 interface should not be assigned to a logical system, as an SA cannot be set up for this interface.

The user logical system administrator can configure the IP address and other attributes of the st0 interface assigned to the user logical system. The user logical system administrator cannot delete an st0 interface assigned to their user logical system.

For route-based VPNs, a security policy refers to a destination address and not a specific VPN tunnel. For cleartext traffic in a user logical system to be sent to the VPN tunnel for encapsulation, the user logical system administrator must make the following configurations:

- Security policy that permits traffic to a specified destination.

- Static route to the destination with the st0 interface as the next hop.

When Junos OS looks up routes in the user logical system to find the interface to use to send traffic to the destination address, it finds a static route through the st0 interface. Traffic is routed to the VPN tunnel as long as the security policy action is permit.

The master logical system and a user logical system can share a route-based VPN tunnel. An st0 interface assigned to a user logical system can also be used by the master logical system. For the master logical system, the master administrator configures a security policy that permits traffic to the remote destination and a static route to the remote destination with the st0 interface as the next hop.

VPN monitoring is configured by the master administrator in the master logical system. For the VPN monitor source interface, the master administrator must specify the st0 interface; a physical interface for a user logical system cannot be specified.

Related Documentation

- *Understanding Route-Based IPsec VPNs*

- User Logical System Configuration Overview on page 23

- Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only) on page 92

- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 133

Example: Configuring a Route-Based VPN Tunnel in a User Logical System

This example shows how to configure a route-based VPN tunnel in a user logical system.

- Requirements on page 133
- Overview on page 133
- Configuration on page 133
- Verification on page 135

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Ensure that an st0 interface is assigned to the user logical system and IKE and IPsec SAs are configured at the root level by the master administrator. See "Example: Configuring IKE and IPsec SAs for a VPN Tunnel (Master Administrators Only)" on page 92.

Overview

In this example, you configure the ls-product-design user logical system as shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

You configure the route-based VPN parameters described in Table 15 on page 133.

Table 15: User Logical System Route-Based VPN Configuration

Feature	Name	Configuration Parameters
Tunnel interface	st0 unit 1	• IPv4 protocol family (inet) • IP address 10.11.11.150/24
Static route		• Destination 192.168.168.0/24 • Next hop st0.1
Security policy	through-vpn	Permit the following traffic: • From zone: ls-product-design-trust • To zone: ls-product-design-untrust • Source address: any • Destination address: 192.168.168.0/24 • Application: any

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

```
set interfaces st0 unit 1 family inet address 10.11.11.150/24
set routing-options static route 192.168.168.0/24 next-hop st0.1
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn match source-address any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn match destination-address 192.168.168.0/24
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn match application any
set security policies from-zone ls-product-design-trust to-zone ls-product-design-untrust
    policy through-vpn then permit
```

Step-by-Step Procedure

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

To configure a route-based VPN tunnel in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   [edit]
   lsdesignadmin1@host:ls-product-design>configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure the VPN tunnel interface.

   ```
   [edit interfaces]
   lsdesignadmin1@host:ls-product-design# set st0 unit 1 family inet address
       10.11.11.150/24
   ```

3. Create a static route to the remote destination.

   ```
   [edit routing-options]
   lsdesignadmin1@host:ls-product-design# set static route 192.168.168.0/24 next-hop
       st0.1
   ```

4. Configure a security policy to permit traffic to the remote destination.

   ```
   [edit security policies from-zone ls-product-design-trust to-zone
       ls-product-design-untrust]
   lsdesignadmin1@host:ls-product-design# set policy through-vpn match
       source-address any
   lsdesignadmin1@host:ls-product-design# set policy through-vpn match
       destination-address 192.168.168.0/24
   lsdesignadmin1@host:ls-product-design# set policy through-vpn match application
       any
   lsdesignadmin1@host:ls-product-design# set policy through-vpn then permit
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces st0**, **show routing-options**, and **show security policies** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
```

```
lsdesignadmin1@host:ls-product-design# show interfaces st0
  unit 1 {
    family inet {
      address 10.11.11.150/24;
    }
  }
  lsdesignadmin1@host:ls-product-design# show routing-options
    static {
      route 192.168.168.0/24 next-hop st0.1;
    }
  [edit]
  lsdesignadmin1@host:ls-product-design# show security policies
    from-zone ls-product-design-trust to-zone ls-product-design-untrust {
      policy through-vpn {
        match {
          source-address any;
          destination-address 192.168.168.0/24;
          application any;
        }
        then {
          permit;
        }
      }
      ...
    }
```

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

Verification

Confirm that the configuration is working properly.

> **(i)** NOTE: Before starting the verification process, you need to send traffic from a host in the user logical system to a host in the 192.168.168.0/24 network. For example, initiate a ping from a host in the 12.1.1.0/24 subnet in the ls-product-design user logical system to the host 192.168.168.10.

- Verifying the IKE Phase 1 Status on page 135
- Verifying the IPsec Phase 2 Status on page 136

Verifying the IKE Phase 1 Status

Purpose Verify the IKE Phase 1 status.

Action From operational mode, enter the **show security ike security-associations** command. After obtaining an index number from the command, use the **show security ike security-associations index** *index_number* **detail** command.

For sample outputs and meanings, see the "Verification" section of *Example: Configuring a Route-Based VPN*.

Verifying the IPsec Phase 2 Status

Purpose   Verify the IPsec Phase 2 status.

Action   From operational mode, enter the **show security ipsec security-associations** command. After obtaining an index number from the command, use the **show security ipsec security-associations index** *index_number* **detail** command.

For sample outputs and meanings, see the "Verification" section of *Example: Configuring a Route-Based VPN*.

Related Documentation
- *Example: Configuring a Route-Based VPN*.
- Understanding Route-Based VPN Tunnels in Logical Systems on page 91
- User Logical System Configuration Overview on page 23

PART 4

Configuring Routing and Interfaces Features

Configuring Master Logical System Routing and Interfaces

Understanding Logical System Interfaces and Routing Instances

Logical interfaces on the device are allocated among the user logical systems by the master administrator. The user logical system administrator configures the attributes of the interfaces, including IP addresses, and assigns them to routing instances and zones.

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

Interfaces and routing instances can be configured in the master logical system and in user logical systems. Configuring an interface or routing instance in a logical system is the same as configuring an interface or routing instance on a device that is not configured for logical systems. Any routing instance created within a logical system is only applicable to that logical system.

The default routing instance, master, refers to the main inet.0 routing table in the logical system. The master routing instance is reserved and cannot be specified as a routing instance. Routes are installed in the master routing instance by default, unless a routing instance is specified. Configure global routing options and protocols for the master routing instance by including statements at the [**edit protocols**] and [**edit routing-options**] hierarchy levels in the logical system.

You can configure only virtual router routing instance type in a user logical system. Only one virtual private LAN service (VPLS) routing instance type can be configured on the device and it must be in the interconnect logical system.

The user logical system administrator can configure and view all attributes for an interface or routing instance in a user logical system. All attributes of an interface or routing instance in a user logical system are also visible to the master administrator.

Multicast is a "one source, many destinations" method of traffic distribution, which means the destinations needing to receive the information from a particular source receive the traffic stream. The master and user logical system administrators can configure a logical system to support multicast applications. The same multicast configurations to configure a device as a node in a multicast network can be used in a logical system.

Related Documentation
- Example: Configuring Interfaces and Routing Instances for a User Logical System on page 158
- User Logical System Configuration Overview on page 23
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)

This topic covers configuration of interfaces, static routes, and routing instances for the master and interconnect logical systems. It also covers configuration of logical tunnel interfaces for user logical systems.

- Requirements on page 140
- Overview on page 141
- Configuration on page 142
- Verification on page 148

Requirements

The example uses an SRX5600 device running Junos operating system (Junos OS) with logical systems.

Before you begin:

- Read "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 to understand how and where this procedure fits in the overall master administrator configuration process.
- Read "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 8

Overview

This scenario shows how to configure interfaces for the logical systems on the device, including an interconnect logical system.

- For the interconnect logical system, the example configures logical tunnel interfaces lt-0/0/0.0, lt-0/0/0.2, lt-0/0/0.4, and lt-0/0/0.6. The example configures a routing instance called vr-ic and assigns the interfaces to it.

  Because the interconnect logical system acts as a virtual switch, it is configured as a virtual private LAN service (VPLS) routing instance type. The interconnect logical system's lt-0/0/0 interfaces are configured with ethernet-vpls as the encapsulation type. The corresponding peer lt-0/0/0 interfaces in the master and user logical systems are configured with Ethernet as the encapsulation type.

  - lt-0/0/0.0 connects to lt-0/0/0.1 on the root logical system.

  - lt-0/0/0.2 connects to lt-0/0/0.3 on the ls-product-design logical system.

  - lt-0/0/0.4 connects to lt-0/0/0.5 on the ls-marketing-dept logical system.

  - lt-0/0/0.6 connects to lt-0/0/0.7 on the ls-accounting-dept logical system.

- For the master logical system, called root-logical-system, the example configures ge-0/0/4.0 and assigns it to the vr1-root routing instance. The example configures lt-0/0/0.1 to connect to lt-0/0/0.0 on the interconnect logical system and assigns it to the vr1-root routing instance. The example configures static routes to allow for communication with other logical systems and assigns them to the vr1-root routing instance.

- For the ls-product-design logical system, the example configures lt-0/0/0.3 to connect to lt-0/0/0.2 on the interconnect logical system.

- For the ls-marketing-dept logical system, the example configures lt-0/0/0.5 to connect to lt-0/0/0.4 on the interconnect logical system.

- For the ls-accounting-dept logical system, the example configures lt-0/0/0.7 to connect to lt-0/0/0.6 on the interconnect logical system.

shows the topology for this deployment including virtual routers and their interfaces for all logical systems.

Figure 4: Configuring Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers

Configuration

This topic explains how to configure interfaces for logical systems.

- Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System on page 142
- Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System on page 144
- Configuring Logical Tunnel Interfaces for the User Logical Systems on page 146

Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System

CLI Quick Configuration

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

    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 0 encapsulation
      ethernet-vpls
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 0 peer-unit 1
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 2 encapsulation
      ethernet-vpls
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 2 peer-unit 3
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 4 encapsulation
      ethernet-vpls
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 4 peer-unit 5
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 6 encapsulation
      ethernet-vpls
    set logical-systems interconnect-logical-system interfaces lt-0/0/0 unit 6 peer-unit 7
    set logical-systems interconnect-logical-system routing-instances vr-ic instance-type
      vpls
    set logical-systems interconnect-logical-system routing-instances vr-ic interface
      lt-0/0/0.0
    set logical-systems interconnect-logical-system routing-instances vr-ic interface
      lt-0/0/0.2
    set logical-systems interconnect-logical-system routing-instances vr-ic interface
      lt-0/0/0.4
    set logical-systems interconnect-logical-system routing-instances vr-ic interface
      lt-0/0/0.6

Step-by-Step Procedure

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

To configure the interconnect system lt-0/0/0 interfaces and routing instances:

1. Configure the lt-0/0/0 interfaces.

   [edit logical-systems]
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 0 encapsulation
      ethernet-vpls
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 0 peer-unit 1
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 2 encapsulation
      ethernet-vpls
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 2 peer-unit 3
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 4 encapsulation
      ethernet-vpls
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 4 peer-unit 5
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 6 encapsulation
      ethernet-vpls
   user@host# set interconnect-logical-system interfaces lt-0/0/0 unit 6 peer-unit 7

2. Configure the routing instance for the interconnect logical system and add its lt-0/0/0 interfaces to it.

   [edit logical-systems]
   user@host# set interconnect-logical-system routing-instances vr-ic instance-type
      vpls
   user@host# set interconnect-logical-system routing-instances vr-ic interface
      lt-0/0/0.0
   user@host# set interconnect-logical-system routing-instances vr-ic interface
      lt-0/0/0.2

```
user@host# set interconnect-logical-system routing-instances vr-ic interface
    lt-0/0/0.4
user@host# set interconnect-logical-system routing-instances vr-ic interface
    lt-0/0/0.6
```

Results   From configuration mode, confirm your configuration by entering the **show logical-systems interconnect-logical-system** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

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

```
user@host# show logical-systems interconnect-logical-system
  interfaces {
    lt-0/0/0 {
      unit 0 {
        encapsulation ethernet-vpls;
        peer-unit 1;
      }
      unit 2 {
        encapsulation ethernet-vpls;
        peer-unit 3;
      }
      unit 4 {
        encapsulation ethernet-vpls;
        peer-unit 5;
      }
      unit 6 {
        encapsulation ethernet-vpls;
        peer-unit 7;
      }
    }
  }
  routing-instances {
    vr-ic {
      instance-type vpls;
      interface lt-0/0/0.0;
      interface lt-0/0/0.2;
      interface lt-0/0/0.4;
      interface lt-0/0/0.6;
    }
  }
```

Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System

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

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 vlan-id 600
set interfaces ge-0/0/4 unit 0 family inet address 50.1.1.1/24
set interfaces ge-0/0/5 vlan-tagging
```

```
set interfaces ge-0/0/6 vlan-tagging
set interfaces ge-0/0/7 vlan-tagging
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 10.0.1.1/24
set routing-instances vr1-root instance-type virtual-router
set routing-instances vr1-root interface ge-0/0/4.0
set routing-instances vr1-root interface lt-0/0/0.1
set routing-instances vr1-root routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
set routing-instances vr1-root routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set routing-instances vr1-root routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
```

Step-by-Step Procedure

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

To configure the master logical system interfaces:

1. Configure the master (root) logical and lt-0/0/0.1 interfaces.

    ```
    [edit interfaces]
    user@host# set ge-0/0/4 vlan-tagging
    user@host# set ge-0/0/4 unit 0 vlan-id 600
    user@host# set ge-0/0/4 unit 0 family inet address 50.1.1.1/24
    user@host# set lt-0/0/0 unit 1 encapsulation ethernet
    user@host# set lt-0/0/0 unit 1 peer-unit 0
    user@host# set lt-0/0/0 unit 1 family inet address 10.0.1.1/24
    ```

2. Configure the interfaces for other logical systems to support VLAN tagging.

    ```
    [edit interfaces]
    user@host# set ge-0/0/5 vlan-tagging
    user@host# set ge-0/0/6 vlan-tagging
    user@host# set ge-0/0/7 vlan-tagging
    ```

3. Configure a routing instance for the master logical system, assign its interfaces to it, and configure static routes for it.

    ```
    [edit routing-instances]
    user@host# set vr1-root instance-type virtual-router
    user@host# set vr1-root interface ge-0/0/4.0
    user@host# set vr1-root interface lt-0/0/0.1
    user@host# set vr1-root routing-options static route 12.1.1.0/24 next-hop 10.0.1.2
    user@host# set vr1-root routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
    user@host# set vr1-root routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
    ```

Results

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

```
[edit]
user@host# show interfaces
  ge-0/0/4 {
    vlan-tagging;
    unit 0 {
      vlan-id 600;
```

```
            family inet {
               address 50.1.1.1/24;
            }
         }
      }
      ge-0/0/5 {
         vlan-tagging;
      }
      ge-0/0/6 {
         vlan-tagging;
      }
      ge-0/0/7 {
         vlan-tagging;
      }
      lt-0/0/0 {
         unit 1 {
            encapsulation ethernet;
            peer-unit 0;
            family inet {
               address 10.0.1.1/24;
            }
         }
      }
   }

[edit]
user@host# show routing-instances
   vr1-root {
      instance-type virtual-router;
      interface ge-0/0/4.0;
      interface lt-0/0/0.1;
      routing-options {
         static {
            route 14.1.1.0/24 next-hop 10.0.1.4;
            route 12.1.1.0/24 next-hop 10.0.1.2;
            route 13.1.1.0/24 next-hop 10.0.1.3;
         }
      }
   }
```

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

Configuring Logical Tunnel Interfaces for the User Logical Systems

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

```
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems ls-product-design interfaces lt-0/0/0 unit 3 family inet address
   10.0.1.2/24
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems ls-marketing-dept interfaces lt-0/0/0 unit 5 family inet address
   10.0.1.3/24
```

```
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 encapsulation ethernet
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 peer-unit 6
set logical-systems ls-accounting-dept interfaces lt-0/0/0 unit 7 family inet address
    10.0.1.4/24
```

Step-by-Step
Procedure

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

1. Configure the lt-0/0/0 interface for the first user logical system:

   ```
   [edit logical-systems]
   user@host# set ls-product-design interfaces lt-0/0/0 unit 3 encapsulation ethernet
   user@host# set ls-product-design interfaces lt-0/0/0 unit 3 peer-unit 2
   user@host# set ls-product-design interfaces lt-0/0/0 unit 3 family inet address
       10.0.1.2/24
   ```

2. Configure the lt-0/0/0 interface for the second user logical system.

   ```
   [edit logical-systems]
   user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 encapsulation ethernet
   user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 peer-unit 4
   user@host# set ls-marketing-dept interfaces lt-0/0/0 unit 5 family inet address
       10.0.1.3/24 face
   ```

3. Configure the lt-0/0/0 interface for the third user logical system.

   ```
   [edit logical-systems]
   user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 encapsulation ethernet
   user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 peer-unit 6
   user@host# set ls-accounting-dept interfaces lt-0/0/0 unit 7 family inet address
       10.0.1.4/24
   ```

Results

From configuration mode, confirm your configuration by entering the **show logical-systems ls-product-design interfaces lt-0/0/0**, **show logical-systems ls-marketing-dept interfaces lt-0/0/0** , and **show logical-systems ls-accounting-dept interfaces lt-0/0/0** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host#  show logical-systems ls-product-design interfaces lt-0/0/0

lt-0/0/0 {
  unit 3 {
    encapsulation ethernet;
    peer-unit 2;
    family inet {
      address 10.0.1.2/24;
    }
  }
}
user@host# show logical-systems ls-marketing-dept interfaces lt-0/0/0
  lt-0/0/0 {
    unit 5 {
      encapsulation ethernet;
      peer-unit 4;
      family inet {
        address 10.0.1.3/24;
```

```
                }
              }
            }
          }
        }
        user@host# show logical-systems ls-accounting-dept interfaces lt-0/0/0
        lt-0/0/0 {
          unit 7 {
            encapsulation ethernet;
            peer-unit 6;
            family inet {
              address 10.0.1.4/24;
            }
          }
        }
```

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

Verification

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

- Verifying That the Static Routes Configured for the Master Administrator Are Correct on page 148

Verifying That the Static Routes Configured for the Master Administrator Are Correct

Purpose Verify if you can send data from the master logical system to the other logical systems.

Action From operational mode, use the **ping** command.

Related Documentation
- Understanding the Master Logical System and the Master Administrator Role on page 19
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 8

Example: Configuring OSPF Routing Protocol for the Master Logical System

This example shows how to configure OSPF for the master logical system.

- Requirements on page 149
- Overview on page 149
- Configuration on page 149
- Verification on page 151

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Example: Configuring a Root Password for the Device (Master Administrators Only)" on page 39.

- Configure logical interfaces ge-0/0/4.0 and lt-0/0/0.1 for the master logical system and assign them to the vr1-root routing instance. See "Example: Configuring Interfaces, Routing Instances, and Static Routes for the Master and Interconnect Logical Systems and Logical Tunnel Interfaces for the User Logical Systems (Master Administrators Only)" on page 140.

Overview

In this example, you configure OSPF for the master logical system, called root-logical-system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example enables OSPF routing on the ge-0/0/4.0 and lt-0/0/0.1 interfaces in the master logical system. You configure the following routing policies to export routes from the Junos OS routing table into OSPF in the vr1-root routing instance:

- ospf-redist-direct—Routes learned from directly connected interfaces.

- ospf-redist-static—Static routes.

- ospf-to-ospf—Routes learned from OSPF.

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

```
set policy-options policy-statement ospf-redist-direct from protocol direct
set policy-options policy-statement ospf-redist-direct then accept
set policy-options policy-statement ospf-redist-static from protocol static
set policy-options policy-statement ospf-redist-static then accept
set policy-options policy-statement ospf-to-ospf from protocol ospf
set policy-options policy-statement ospf-to-ospf then accept
set routing-instances vr1-root protocols ospf export ospf-redist-direct
set routing-instances vr1-root protocols ospf export ospf-redist-static
set routing-instances vr1-root protocols ospf export ospf-to-ospf
set routing-instances vr1-root protocols ospf area 0.0.0.1 interface ge-0/0/4.0
set routing-instances vr1-root protocols ospf area 0.0.0.1 interface lt-0/0/0.1
```

Step-by-Step Procedure

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

To configure OSPF for the master logical system:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   admin@host> configure
   admin@host#
   ```

2. Create routing policies that accept routes.

   ```
   [edit policy-options]
   admin@host# set policy-statement ospf-redist-direct from protocol direct
   admin@host# set policy-statement ospf-redist-direct then accept
   admin@host# set policy-statement ospf-redist-static from protocol static
   admin@host# set policy-statement ospf-redist-static then accept
   admin@host# set policy-statement ospf-to-ospf from protocol ospf
   admin@host# set policy-statement ospf-to-ospf then accept
   ```

3. Apply the routing policies to routes exported from the Junos OS routing table into OSPF.

   ```
   [edit routing-instances]
   admin@host# set vr1-root protocols ospf export ospf-redist-direct
   admin@host# set vr1-root protocols ospf export ospf-redist-static
   admin@host# set vr1-root protocols ospf export ospf-to-ospf
   ```

4. Enable OSPF on the logical interfaces.

   ```
   [edit routing-instances]
   admin@host# set vr1-root protocols ospf area 0.0.0.1 interface ge-0/0/4.0
   admin@host# set vr1-root protocols ospf area 0.0.0.1 interface lt-0/0/0.1
   ```

Results

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

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
  admin@host# show policy-options
  policy-statement ospf-redist-direct {
    from protocol direct;
    then accept;
  }
  policy-statement ospf-redist-static {
    from protocol static;
    then accept;
  }
  policy-statement ospf-to-ospf {
    from protocol ospf;
    then accept;
```

```
        }
      [edit]
   admin@host# show routing-instances
     vr1-root {
       ...
       protocols {
         ospf {
           export [ ospf-redist-direct ospf-to-ospf ospf-redist-static ];
           area 0.0.0.1 {
             interface lt-0/0/0.1;
             interface ge-0/0/4.0;
           }
         }
       }
     }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying OSPF Interfaces on page 151
- Verifying OSPF Neighbors on page 151
- Verifying OSPF Routes on page 151

Verifying OSPF Interfaces

Purpose   Verify OSPF-enabled interfaces.

Action   From the CLI, enter the **show ospf interface instance vr1-root** command.

```
admin@host>  show ospf interface instance vr1-root
Interface          State    Area            DR ID          BDR ID          Nbrs
lt-0/0/0.1         DR       0.0.0.0         10.0.1.1       0.0.0.0            0
ge-0/0/4.0         DR       0.0.0.1         10.0.1.1       0.0.0.0            0
```

Verifying OSPF Neighbors

Purpose   Verify OSPF neighbors.

Action   From the CLI, enter the **show ospf neighbor instance vr1-root** command.

```
admin@host>  show ospf neighbor instance vr1-root
Address    Interface  State  ID      Pri Dead
10.0.1.2   plt0.3     Full   0.0.0.0   128   39
```

Verifying OSPF Routes

Purpose   Verify OSPF routes.

Action   From the CLI, enter the **show ospf route instance vr1-root** command.

```
admin@host>  show ospf route instance vr1-root
```

```
Topology default Route Table:

Prefix            Path   Route     NH      Metric NextHop      Nexthop
                  Type   Type      Type           Interface    Address/LSP
10.0.1.0/24       Intra  Network   IP          1  lt-0/0/0.1
12.12.1.0/24      Intra  Network   IP          1  ge-0/0/4.0
```

Related Documentation

- Understanding Logical System Interfaces and Routing Instances on page 139

- Example: Configuring OSPF Routing Protocol for a User Logical System on page 160

- Verifying an OSPF Configuration

Configuring User Logical System Routing, Interfaces, and NAT Features

Understanding Logical System Network Address Translation

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses.

Any combination of static, destination, or source NAT can be configured in the root or user logical systems. Configuring NAT in a logical system is the same as configuring NAT in a root system. The master administrator can configure and monitor NAT in the master logical system as well as any user logical system.

For each user logical system, the master administrator can configure the maximum and reserved numbers for the following NAT resources:

- Source NAT pools and destination NAT pools
- IP addresses in source NAT pools with and without port address translation
- Rules for source, destination, and static NAT
- Persistent NAT bindings
- IP addresses that support port overloading

From a user logical system, the user logical system administrator can use the operational command **show system security-profile** with a NAT option to view the number of NAT resources allocated to the user logical system.

> ℹ️ **NOTE:** The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of NAT resources applied to the master logical system. The number of resources configured in the master logical system count toward the maximum number of NAT resources available on the device.

From a user logical system, the user logical system administrator can use the **show security nat** command to view the information about NAT for the user logical system. From the master logical system, the master administrator can use the same command to view information for the master logical system, a specific user logical system, or all logical systems.

Related Documentation

- Example: Configuring Network Address Translation for a User Logical System on page 154
- User Logical System Configuration Overview on page 23
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51
- *Introduction to NAT*

Example: Configuring Network Address Translation for a User Logical System

This example shows how to configure static NAT for a user logical system.

- Requirements on page 154
- Overview on page 154
- Configuration on page 155
- Verification on page 156

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator. See "User Logical System Configuration Overview" on page 23.
- Use the **show system security-profile nat-static-rule** command to see the static NAT resources allocated to the logical system.
- Configure security policies. See "Example: Configuring Security Policies in a User Logical System" on page 106.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

Devices in the ls-product-design-untrust zone access a specific host in the ls-product-design-trust zone by way of the address 12.1.1.200/32. For packets that enter the ls-product-design logical system from the ls-product-design-untrust zone with the destination IP address 12.1.1.200/32, the destination IP address is translated to the 12.1.1.100/32. This example configures the static NAT described in .

Table 16: User Logical System Static NAT Configuration

Feature	Name	Configuration Parameters
Static NAT rule set	rs1	• Rule r1 to match packets from the ls-product-design-untrust zone with destination address 12.1.1.200/32. • Destination IP address in matching packets is translated to 12.1.1.100/32.
Proxy ARP		Address 12.1.1.200 on interface lt-0/0/0.3.

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

```
set security nat static rule-set rs1 from zone ls-product-design-untrust
set security nat static rule-set rs1 rule r1 match destination-address 12.1.1.200/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 12.1.1.100/32
set security nat proxy-arp interface lt-0/0/0.3 address 12.1.1.200/32
```

Step-by-Step Procedure

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

To configure NAT in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a static NAT rule set.

   ```
   [edit security nat static]
   lsdesignadmin1@host:ls-product-design# set rule-set rs1 from zone
       ls-product-design-untrust
   ```

3. Configure a rule that matches packets and translates the destination address in the packets.

   ```
   [edit security nat static]
   lsdesignadmin1@host:ls-product-design# set rule-set rs1 rule r1 match
       destination-address 12.1.1.200/32
   lsdesignadmin1@host:ls-product-design# set rule-set rs1 rule r1 then static-nat prefix
       12.1.1.100/32
   ```

4. Configure proxy ARP.

```
[edit security nat]
lsdesignadmin1@host:ls-product-design# set proxy-arp interface lt-0/0/0.3 address
    12.1.1.200/32
```

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

```
lsdesignadmin1@host:ls-product-design# show security nat
static {
  rule-set rs1 {
    from zone ls-product-design-untrust;
    rule r1 {
      match {
        destination-address 12.1.1.200/32;
      }
      then {
        static-nat prefix 12.1.1.100/32;
      }
    }
  }
}
  proxy-arp {
    interface lt-0/0/0.3 {
      address {
        12.1.1.200/32;
      }
    }
  }
}
```

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

Verification

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

Verifying Static NAT Configuration

Purpose   Verify that there is traffic matching the static NAT rule set.

Action   From operational mode, enter the **show security nat static rule** command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose   Verify that NAT is being applied to the specified traffic.

Action   From operational mode, enter the **show security flow session** command.

- User Logical System Configuration Overview on page 23

- Understanding Logical System Network Address Translation on page 153

- *Static NAT Configuration Overview*

Understanding Logical System Interfaces and Routing Instances

Logical interfaces on the device are allocated among the user logical systems by the master administrator. The user logical system administrator configures the attributes of the interfaces, including IP addresses, and assigns them to routing instances and zones.

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. There can be multiple routing tables for a single routing instance—for example, unicast IPv4, unicast IPv6, and multicast IPv4 routing tables can exist in a single routing instance. Routing protocol parameters and options control the information in the routing tables.

Interfaces and routing instances can be configured in the master logical system and in user logical systems. Configuring an interface or routing instance in a logical system is the same as configuring an interface or routing instance on a device that is not configured for logical systems. Any routing instance created within a logical system is only applicable to that logical system.

The default routing instance, master, refers to the main inet.0 routing table in the logical system. The master routing instance is reserved and cannot be specified as a routing instance. Routes are installed in the master routing instance by default, unless a routing instance is specified. Configure global routing options and protocols for the master routing instance by including statements at the [**edit protocols**] and [**edit routing-options**] hierarchy levels in the logical system.

You can configure only virtual router routing instance type in a user logical system. Only one virtual private LAN service (VPLS) routing instance type can be configured on the device and it must be in the interconnect logical system.

The user logical system administrator can configure and view all attributes for an interface or routing instance in a user logical system. All attributes of an interface or routing instance in a user logical system are also visible to the master administrator.

Multicast is a "one source, many destinations" method of traffic distribution, which means the destinations needing to receive the information from a particular source receive the traffic stream. The master and user logical system administrators can configure a logical system to support multicast applications. The same multicast configurations to configure a device as a node in a multicast network can be used in a logical system.

- Example: Configuring Interfaces and Routing Instances for a User Logical System on page 158

- User Logical System Configuration Overview on page 23

- Understanding User Logical Systems and the User Logical System Administrator Role on page 25

Example: Configuring Interfaces and Routing Instances for a User Logical System

This example shows how to configure interfaces and routing instances for a user logical system.

- Requirements on page 158
- Overview on page 158
- Configuration on page 158

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Determine which logical interfaces and, optionally, which logical tunnel interfaces are allocated to your user logical system by the master administrator. The master administrator configures the logical tunnel interfaces. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

Overview

This example configures the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example configures the interfaces and routing instances described in Table 17 on page 158.

Table 17: User Logical System Interface and Routing Instance Configuration

Feature	Name	Configuration Parameters
Interface	ge-0/0/5.1	- IP address 12.1.1.1/24 - VLAN ID 700
Routing instance	pd-vr1	- Instance type: virtual router - Includes interfaces ge-0/0/5.1 and lt-0/0/0.3 - Static routes: - 13.1.1.0/24 next-hop 10.0.1.3 - 14.1.1.0/24 next-hop 10.0.1.4 - 12.12.1.0/24 next-hop 10.0.1.1

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

```
set interfaces ge-0/0/5 unit 1 family inet address 12.1.1.1/24
set interfaces ge-0/0/5 unit 1 vlan-id 700
set routing-instances pd-vr1 instance-type virtual-router
set routing-instances pd-vr1 interface ge-0/0/5.1
set routing-instances pd-vr1 interface lt-0/0/0.3
set routing-instances pd-vr1 routing-options static route 13.1.1.0/24 next-hop 10.0.1.3
set routing-instances pd-vr1 routing-options static route 14.1.1.0/24 next-hop 10.0.1.4
set routing-instances pd-vr1 routing-options static route 12.12.1.0/24 next-hop 10.0.1.1
```

Step-by-Step Procedure

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

To configure an interface and a routing instance in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure the logical interface for a user logical system.

   ```
   [edit interfaces]
   lsdesignadmin1@host:ls-product-design# set ge-0/0/5 unit 1 family inet address
       12.1.1.1/24
   lsdesignadmin1@host:ls-product-design# set ge-0/0/5 unit 1 vlan-id 700
   ```

3. Configure the routing instance and assign interfaces.

   ```
   [edit routing-instances]
   lsdesignadmin1@host:ls-product-design# set pd-vr1 instance-type virtual-router
   lsdesignadmin1@host:ls-product-design# set pd-vr1 interface ge-0/0/5.1
   lsdesignadmin1@host:ls-product-design# set pd-vr1 interface lt-0/0/0.3
   ```

4. Configure static routes.

   ```
   [edit routing-instances]
   lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route
       13.1.1.0/24 next-hop 10.0.1.3
   lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route
       14.1.1.0/24 next-hop 10.0.1.4
   lsdesignadmin1@host:ls-product-design# set pd-vr1 routing-options static route
       12.12.1.0/24 next-hop 10.0.1.1
   ```

Results

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

> **NOTE:** The master administrator configures the lt-0/0/0.3 interface. Thus, the lt-0/0/0.3 configuration appears in the **show interfaces** output even though you did not configure this item.

```
lsdesignadmin1@host:ls-product-design# show interfaces
```

```
                    ge-0/0/5 {
                       unit 1 {
                          vlan-id 700;
                          family inet {
                             address 12.1.1.1/24;
                          }
                       }
                    }
                       lt-0/0/0 {
                          unit 3 {
                             encapsulation ethernet;
                             peer-unit 2;
                             family inet {
                                address 10.0.1.2/24;
                             }
                          }
                       }
                    lsdesignadmin1@host:ls-product-design# show routing-instances
                    pd-vr1 {
                       instance-type virtual-router;
                       interface ge-0/0/5.1;
                       interface lt-0/0/0.3;
                       routing-options {
                          static {
                             route 13.1.1.0/24 next-hop 10.0.1.3;
                             route 14.1.1.0/24 next-hop 10.0.1.4;
                             route 12.12.1.0/24 next-hop 10.0.1.1;
                          }
                       }
                    }
```

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

Related Documentation

- User Logical System Configuration Overview on page 23

- Understanding Logical System Interfaces and Routing Instances on page 139

Example: Configuring OSPF Routing Protocol for a User Logical System

This example shows how to configure OSPF for a user logical system.

- Requirements on page 160
- Overview on page 161
- Configuration on page 161
- Verification on page 163

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Configure logical interface ge-0/0/5.1. Assign ge-0/0/5.1 and lt-0/0/0.3 to the pd-vr1 routing instance. See "Example: Configuring Interfaces and Routing Instances for a User Logical System" on page 158.

Overview

In this example, you configure OSPF for the ls-product-design user logical system, shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.

This example enables OSPF routing on the ge-0/0/5.1 and lt-0/0/0.3 interfaces in the ls-product-design user logical system. You configure the following routing policies to export routes from the Junos OS routing table into OSPF in the pd-vr1 routing instance:

- ospf-redist-direct—Routes learned from directly connected interfaces.

- ospf-redist-static—Static routes.

- ospf-to-ospf—Routes learned from OSPF.

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

```
set policy-options policy-statement ospf-redist-direct from protocol direct
set policy-options policy-statement ospf-redist-direct then accept
set policy-options policy-statement ospf-redist-static from protocol static
set policy-options policy-statement ospf-redist-static then accept
set policy-options policy-statement ospf-to-ospf from protocol ospf
set policy-options policy-statement ospf-to-ospf then accept
set routing-instances pd-vr1 protocols ospf export ospf-redist-direct
set routing-instances pd-vr1 protocols ospf export ospf-redist-static
set routing-instances pd-vr1 protocols ospf export ospf-to-ospf
set routing-instances pd-vr1 protocols ospf area 0.0.0.1 interface ge-0/0/5.1
set routing-instances pd-vr1 protocols ospf area 0.0.0.1 interface lt-0/0/0.3
```

Step-by-Step Procedure

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

To configure OSPF for the user logical system:

1. Log in to the user logical system as the user logical system administrator and enter configuration mode.

   lsdesignadmin1@host:ls-product-design> **configure**
   lsdesignadmin1@host:ls-product-design#

2. Create routing policies that accept routes.

   [edit policy-options]

```
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-direct
     from protocol direct
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-direct
     then accept
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-static
     from protocol static
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-redist-static
     then accept
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-to-ospf from
     protocol ospf
lsdesignadmin1@host:ls-product-design# set policy-statement ospf-to-ospf then
     accept
```

3. Apply the routing policies to routes exported from the Junos OS routing table into OSPF.

    ```
    [edit routing-instances]
    lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export
         ospf-redist-direct
    lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export
         ospf-redist-static
    lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf export
         ospf-to-ospf
    ```

4. Enable OSPF on the logical interfaces.

    ```
    [edit routing-instances]
    lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf area 0.0.0.1
         interface ge-0/0/5.1
    lsdesignadmin1@host:ls-product-design# set pd-vr1 protocols ospf area 0.0.0.1
         interface lt-0/0/0.3
    ```

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

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
lsdesignadmin1@host:ls-product-design# show policy-options
  policy-statement ospf-redist-direct {
    from protocol direct;
    then accept;
  }
  policy-statement ospf-redist-static {
    from protocol static;
    then accept;
  }
  policy-statement ospf-to-ospf {
    from protocol ospf;
    then accept;
  }
[edit]
lsdesignadmin1@host:ls-product-design# show routing-instances
```

```
pd-vr1 {
   ...
   protocols {
      ospf {
         export [ ospf-redist-direct ospf-to-ospf ospf-redist-static ];
         area 0.0.0.1 {
            interface lt-0/0/0.3;
            interface ge-0/0/5.1;
         }
      }
   }
}
```

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

Verification

Confirm that the configuration is working properly.

-
-
-

Verifying OSPF Interfaces

Purpose Verify OSPF-enabled interfaces.

Action From the CLI, enter the **show ospf interface instance pd-vr1** command.

```
lsdesignadmin1@host:ls-product-design> show ospf interface instance pd-vr1
Interface          State  Area          DR ID         BDR ID          Nbrs
lt-0/0/0.3         DR     0.0.0.0       10.0.1.2      0.0.0.0            0
ge-0/0/5.1         DR     0.0.0.1       10.0.1.2      0.0.0.0            0
```

Verifying OSPF Neighbors

Purpose Verify OSPF neighbors.

Action From the CLI, enter the **show ospf neighbor instance pd-vr1** command.

```
lsdesignadmin1@host:ls-product-design> show ospf neighbor instance pd-vr1
Address    Interface  State   ID       Pri Dead
10.0.1.1   plt0.1     Full    0.0.0.0   128   39
```

Verifying OSPF Routes

Purpose Verify OSPF routes.

Action From the CLI, enter the **show ospf route instance pd-vr1** command.

```
lsdesignadmin1@host:ls-product-design> show ospf route instance pd-vr1
Topology default Route Table:

Prefix             Path  Route     NH      Metric NextHop     Nexthop
                   Type  Type      Type           Interface   Address/LSP
```

```
10.0.1.0/24        Intra Network    IP          1 lt-0/0/0.3
12.12.1.0/24       Intra Network    IP          1 ge-0/0/5.1
```

Related Documentation

- Understanding Logical System Interfaces and Routing Instances on page 139

- Example: Configuring OSPF Routing Protocol for the Master Logical System on page 148

- Verifying an OSPF Configuration

Configuring Logical Systems in Chassis Cluster

Configuring Logical Systems When Device is in Chassis Cluster Mode

Understanding Logical Systems in the Context of Chassis Cluster

The behavior of a chassis cluster whose nodes consist of SRX Series devices running logical systems is the same as that of a cluster whose SRX Series nodes in the cluster are not running logical systems. No difference exists between events that cause a node to fail over. In particular, if a link associated with a single logical system fails, then the device fails over to another node in the cluster.

The master administrator configures the chassis cluster (including both primary and secondary nodes) before he or she creates and configures the logical systems. Each node in the cluster has the same configuration, as is the case for nodes in a cluster not running logical systems. All logical system configurations are synchronized and replicated between both nodes in the cluster.

When you use SRX Series devices running logical systems within a chassis cluster, you must purchase and install the same number of licenses for each node in the chassis cluster. Logical systems licenses pertain to a single chassis, or node, within a chassis cluster and not to the cluster collectively.

Related
Documentation
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 168
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 201
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 8
- Understanding Logical Systems for SRX Series Services Gateways on page 3
- *Chassis Cluster Overview*

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only)

This example shows how to configure logical systems in a basic active/passive chassis cluster.

> ℹ️ NOTE: The master administrator configures the chassis cluster and creates logical systems (including an optional interconnect logical system), administrators, and security profiles. Either the master administrator or the user logical system administrator configures a user logical system. The configuration is synchronized between nodes in the cluster.

- Requirements on page 168
- Overview on page 169
- Configuration on page 172
- Verification on page 195

Requirements

Before you begin:

- Obtain two high-end SRX Series Services Gateways with identical hardware configurations. See *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*. This chassis cluster deployment scenario includes the configuration of the SRX Series device for connections to an MX240 edge router and an EX8208 Ethernet Switch.

- Physically connect the two devices (back-to-back for the fabric and control ports) and ensure that they are the same models. You can configure both the fabric and control ports on the SRX5000 line. See *Connecting SRX Series Devices to Create a Chassis Cluster*.

- Set the chassis cluster ID and node ID on each device and reboot the devices to enable clustering. See *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*.

> ℹ️ NOTE: For this example, chassis cluster and logical system configuration is performed on the primary (node 0) device at the root level by the master administrator. Log in to the device as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

NOTE: When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of logical system licenses for each node in the chassis cluster. Logical system licenses pertain to a single chassis or node within a chassis cluster and not to the cluster collectively. See "Understanding Licenses for Logical Systems on SRX Series Devices" on page 7.

Overview

In this example, the basic active/passive chassis cluster consists of two devices:

- One device actively provides logical systems, along with maintaining control of the chassis cluster.

- The other device passively maintains its state for cluster failover capabilities should the active device become inactive.

NOTE: Logical systems in an active/active chassis cluster are configured in a similar manner as for logical systems in an active/passive chassis cluster. For active/active chassis clusters, there can be multiple redundancy groups that can be primary on different nodes.

The master administrator configures the following logical systems on the primary device (node 0):

- Master logical system—The master administrator configures a security profile to provision portions of the system's security resources to the master logical system and configures the resources of the master logical system.

- User logical systems LSYS1 and LSYS2 and their administrators—The master administrator also configures security profiles to provision portions of the system's security resources to user logical systems. The user logical system administrator can then configure interfaces, routing, and security resources allocated to his or her logical system.

- Interconnect logical system LSYS0 that connects logical systems on the device—The master administrator configures logical tunnel interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for the establishment of tunnels.

NOTE: This example does not describe configuring features such as NAT, IDP, or VPNs for a logical system. See "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 and "User Logical System Configuration Overview" on page 23 for more information about features that can be configured for logical systems.

If you are performing proxy ARP in a chassis cluster configuration, you must apply the proxy ARP configuration to the reth interfaces rather than the member interfaces because the reth interfaces contain the logical configurations. See *Configuring Proxy ARP (CLI Procedure)*.

Topology

shows the topology used in this example.

Figure 5: Logical Systems in a Chassis Cluster

Configuration

Chassis Cluster Configuration (Master Administrator)

CLI Quick
Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, 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.

On {primary:node0}

```
set chassis cluster control-ports fpc 0 port 0
set chassis cluster control-ports fpc 6 port 0
set interfaces fab0 fabric-options member-interfaces ge-1/1/0
set interfaces fab1 fabric-options member-interfaces ge-7/1/0
set groups node0 system host-name SRX5800-1
set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
set groups node1 system host-name SRX5800-2
set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
set apply-groups "${node}"
set chassis cluster reth-count 5
set chassis cluster redundancy-group 0 node 0 priority 200
set chassis cluster redundancy-group 0 node 1 priority 100
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 100
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/0/2 gigether-options redundant-parent reth2
set interfaces ge-1/0/3 gigether-options redundant-parent reth3
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/2 gigether-options redundant-parent reth2
set interfaces ge-7/0/3 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet address 95.99.99.1/8
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
```

Step-by-Step Procedure

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

To configure a chassis cluster:

> **NOTE:** Perform the following steps on the primary device (node 0). They are automatically copied over to the secondary device (node 1) when you execute a **commit** command.

---

1. Configure control ports for the clusters.

   ```
   [edit chass cluster]
   user@host# set control-ports fpc 0 port 0
   user@host# set control-ports fpc 6 port 0
   ```

2. Configure the fabric (data) ports of the cluster that are used to pass RTOs in active/passive mode.

   ```
   [edit interfaces]
   user@host# set fab0 fabric-options member-interfaces ge-1/1/0
   user@host# set fab1 fabric-options member-interfaces ge-7/1/0
   ```

3. Assign some elements of the configuration to a specific member. Configure out-of-band management on the fxp0 interface of the SRX Services Gateway using separate IP addresses for the individual control planes of the cluster.

   ```
   [edit]
   user@host# set groups node0 system host-name SRX5800-1
   user@host# set groups node0 interfaces fxp0 unit 0 family inet address
       10.157.90.24/9
   user@host# set groups node0 system backup-router 10.157.64.1 destination
       0.0.0.0/0
   user@host# set groups node1 system host-name SRX5800-2
   user@host# set groups node1 interfaces fxp0 unit 0 family inet address
       10.157.90.23/19
   user@host# set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
   user@host# set apply-groups "${node}"
   ```

4. Configure redundancy groups for chassis clustering.

   ```
   [edit chassis cluster]
   user@host# set reth-count 5
   user@host# set redundancy-group 0 node 0 priority 200
   user@host# set redundancy-group 0 node 1 priority 100
   user@host# set redundancy-group 1 node 0 priority 200
   user@host# set redundancy-group 1 node 1 priority 100
   ```

5. Configure the data interfaces on the platform so that in the event of a data plane failover, the other chassis cluster member can take over the connection seamlessly.

   ```
   [edit interfaces]
   user@host# set ge-1/0/0 gigether-options redundant-parent reth0
   user@host# set ge-1/0/1 gigether-options redundant-parent reth1
   user@host# set ge-1/0/2 gigether-options redundant-parent reth2
   ```

```
user@host# set ge-1/0/3 gigether-options redundant-parent reth3
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/2 gigether-options redundant-parent reth2
user@host# set ge-7/0/3 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet address 95.99.99.1/8
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

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

```
user@host> show configuration
version ;
groups {
    node0 {
        system {
            host-name SRX58001;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.24/9;
                    }
                }
            }
        }
    }
    node1 {
        system {
            host-name SRX58002;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.23/19;
                    }
                }
            }
        }
    }
}
apply-groups "${node}";
chassis {
    cluster {
        control-link-recovery;
        reth-count 5;
        control-ports {
            fpc 0 port 0;
            fpc 6 port 0;
```

```
            }
            redundancy-group 0 {
                node 0 priority 200;
                node 1 priority 100;
            }
            redundancy-group 1 {
                node 0 priority 200;
                node 1 priority 100;
            }
        }
    }
    interfaces {
        ge-1/0/0 {
            gigether-options {
                redundant-parent reth0;
            }
        }
        ge-1/0/1 {
            gigether-options {
                redundant-parent reth1;
            }
        }
        ge-1/0/2 {
            gigether-options {
                redundant-parent reth2;
            }
        }
        ge-1/0/3 {
            gigether-options {
                redundant-parent reth3;
            }
        }
        ge-7/0/0 {
            gigether-options {
                redundant-parent reth0;
            }
        }
        ge-7/0/1 {
            gigether-options {
                redundant-parent reth1;
            }
        }
        ge-7/0/2 {
            gigether-options {
                redundant-parent reth2;
            }
        }
        ge-7/0/3 {
            gigether-options {
                redundant-parent reth3;
            }
        }
        fab0 {
            fabric-options {
                member-interfaces {
                    ge-1/1/0;
                }
            }
        }
        fab1 {
            fabric-options {
```

```
            member-interfaces {
                ge-7/1/0;
            }
        }
    }
    reth0 {
        redundant-ether-options {
            redundancy-group 1;
        }
        unit 0 {
            family inet {
                address 95.99.99.1/8;
            }
        }
    }
    reth1 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth2 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
    reth3 {
        redundant-ether-options {
            redundancy-group 1;
        }
    }
}
```

Logical System Configuration (Master Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, 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.

---

i NOTE: You are prompted to enter and then reenter plain-text passwords.

On {primary:node0}

```
set logical-systems LSYS1
set logical-systems LSYS2
set logical-systems LSYS0
set system login class lsys1 logical-system LSYS1
set system login class lsys1 permissions all
set system login user lsys1admin full-name lsys1-admin
set system login user lsys1admin class lsys1
set user lsys1admin authentication plain-text-password
set system login class lsys2 logical-system LSYS2
set system login class lsys2 permissions all
set system login user lsys2admin full-name lsys2-admin
```

```
set system login user lsys2admin class lsys2
set system login user lsys2admin authentication plain-text-password
set system security-profile SP-root policy maximum 200
set system security-profile SP-root policy reserved 100
set system security-profile SP-root zone maximum 200
set system security-profile SP-root zone reserved 100
set system security-profile SP-root flow-session maximum 200
set system security-profile SP-root flow-session reserved 100
set system security-profile SP-root root-logical-system
set system security-profile SP0 logical-system LSYS0
set system security-profile SP1 policy maximum 100
set system security-profile SP1 policy reserved 50
set system security-profile SP1 zone maximum 100
set system security-profile SP1 zone reserved 50
set system security-profile SP1 flow-session maximum 100
set system security-profile SP1 flow-session reserved 50
set system security-profile SP1 logical-system LSYS1
set system security-profile SP2 policy maximum 100
set system security-profile SP2 policy reserved 50
set system security-profile SP2 zone maximum 100
set system security-profile SP2 zone reserved 50
set system security-profile SP2 flow-session maximum 100
set system security-profile SP2 flow-session reserved 50
set system security-profile SP2 logical-system LSYS2
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 2.1.1.1/24
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface reth0.0
set routing-instances vr0 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr0 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr0 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
set security zones security-zone root-trust host-inbound-traffic system-services all
set security zones security-zone root-trust host-inbound-traffic protocols all
set security zones security-zone root-trust interfaces reth0.0
set security zones security-zone root-untrust host-inbound-traffic system-services all
set security zones security-zone root-untrust host-inbound-traffic protocols all
set security zones security-zone root-untrust interfaces lt-0/0/0.1
set security policies from-zone root-trust to-zone root-untrust policy
   root-Trust_to_root-Untrust match source-address any
set security policies from-zone root-trust to-zone root-untrust policy
   root-Trust_to_root-Untrust match destination-address any
set security policies from-zone root-trust to-zone root-untrust policy
   root-Trust_to_root-Untrust match application any
set security policies from-zone root-trust to-zone root-untrust policy
   root-Trust_to_root-Untrust then permit
set security policies from-zone root-untrust to-zone root-trust policy
   root-Untrust_to_root-Trust match source-address any
set security policies from-zone root-untrust to-zone root-trust policy
   root-Untrust_to_root-Trust match destination-address any
set security policies from-zone root-untrust to-zone root-trust policy
   root-Untrust_to_root-Trust match application any
set security policies from-zone root-untrust to-zone root-trust policy
   root-Untrust_to_root-Trust then permit
```

```
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust match source-address any
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust match destination-address any
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust match application any
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust then permit
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    match source-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    match destination-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    match application any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    then permit
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet address 2.1.1.3/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet address 2.1.1.5/24
```

Step-by-Step Procedure

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

To create logical systems and user logical system administrators and configure the master and interconnect logical systems:

1. Create the interconnect and user logical systems.

   ```
   [edit logical-systems]
   user@host# set LSYS0
   user@host# set LSYS1
   user@host# set LSYS2
   ```

2. Configure user logical system administrators.

   a. Configure the user logical system administrator for LSYS1.

      ```
      [edit system login]
      user@host# set class lsys1 logical-system LSYS1
      user@host# set class lsys1 permissions all
      user@host# set user lsys1admin full-name lsys1-admin
      user@host# set user lsys1admin class lsys1
      user@host# set user lsys1admin authentication plain-text-password
      ```

b. Configure the user logical system administrator for LSYS2.

[edit system login]
user@host# set class lsys2 logical-system LSYS2
user@host# set class lsys2 permissions all
user@host# set user lsys2admin full-name lsys2-admin
user@host# set user lsys2admin class lsys2
user@host# set user lsys2admin authentication plain-text-password

3. Configure security profiles and assign them to logical systems.

a. Configure a security profile and assign it to the root logical system.

[edit system security-profile]
user@host# set SP-root policy maximum 200
user@host# set SP-root policy reserved 100
user@host# set SP-root zone maximum 200
user@host# set SP-root zone reserved 100
user@host# set SP-root flow-session maximum 200
user@host# set SP-root flow-session reserved 100
user@host# set SP-root root-logical-system

b. Assign a dummy security profile containing no resources to the interconnect logical system LSYS0.

[edit system security-profile]
user@host# set SP0 logical-system LSYS0

c. Configure a security profile and assign it to LSYS1.

[edit system security-profile]
user@host# set SP1 policy maximum 100
user@host# set SP1 policy reserved 50
user@host# set SP1 zone maximum 100
user@host# set SP1 zone reserved 50
user@host# set SP1 flow-session maximum 100
user@host# set SP1 flow-session reserved 50
user@host# set SP1 logical-system LSYS1

d. Configure a security profile and assign it to LSYS2.

[edit system security-profile]
user@host# set SP2 policy maximum 100
user@host# set SP2 policy reserved 50
user@host# set SP2 zone maximum 100
user@host# set SP2 zone reserved 50
user@host# set SP2 flow-session maximum 100
user@host# set SP2 flow-session reserved 50
user@host# set SP2 logical-system LSYS2

4. Configure the master logical system.

a. Configure logical tunnel interfaces.

[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation ethernet
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet address 2.1.1.1/24

b. Configure a routing instance.

```
[edit routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface reth0.0
user@host# set vr0 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
user@host# set vr0 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
user@host# set vr0 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
```

c. Configure zones.

```
[edit security zones]
user@host# set security-zone root-trust host-inbound-traffic system-services
    all
user@host# set security-zone root-trust host-inbound-traffic protocols all
user@host# set security-zone root-trust interfaces reth0.0
user@host# set security-zone root-untrust host-inbound-traffic system-services
    all
user@host# set security-zone root-untrust host-inbound-traffic protocols all
user@host# set security-zone root-untrust interfaces lt-0/0/0.1
```

d. Configure security policies.

```
[edit security policies from-zone root-trust to-zone root-untrust]
user@host# set policy root-Trust_to_root-Untrust match source-address any
user@host# set policy root-Trust_to_root-Untrust match destination-address
    any
user@host# set policy root-Trust_to_root-Untrust match application any
user@host# set policy root-Trust_to_root-Untrust then permit
```

```
[edit security policies from-zone root-untrust to-zone root-trust]
user@host# set policy root-Untrust_to_root-Trust match source-address any
user@host# set policy root-Untrust_to_root-Trust match destination-address
    any
user@host# set policy root-Untrust_to_root-Trust match application any
user@host# set policy root-Untrust_to_root-Trust then permit
```

```
[edit security policies from-zone root-untrust to-zone root-untrust]
user@host# set policy root-Untrust_to_root-Untrust match source-address any
user@host# set policy root-Untrust_to_root-Untrust match destination-address
    any
user@host# set policy root-Untrust_to_root-Untrust match application any
user@host# set policy root-Untrust_to_root-Untrust then permit
```

```
[edit security policies from-zone root-trust to-zone root-trust]
user@host# set policy root-Trust_to_root-Trust match source-address any
user@host# set policy root-Trust_to_root-Trust match destination-address any
user@host# set policy root-Trust_to_root-Trust match application any
user@host# set policy root-Trust_to_root-Trust then permit
```

5. Configure the interconnect logical system.

a. Configure logical tunnel interfaces.

```
[edit logical-systems LSYS0 interfaces]
user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
user@host# set lt-0/0/0 unit 0 peer-unit 1
user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
```

```
                                   user@host# set lt-0/0/0 unit 2 peer-unit 3
                                   user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
                                   user@host# set lt-0/0/0 unit 4 peer-unit 5
```

   b. Configure the VPLS routing instance.

```
                                   [edit logical-systems LSYS0 routing-instances]
                                   user@host# set vr instance-type vpls
                                   user@host# set vr interface lt-0/0/0.0
                                   user@host# set vr interface lt-0/0/0.2
                                   user@host# set vr interface lt-0/0/0.4
```

6. Configure logical tunnel interfaces for the user logical systems.

   a. Configure logical tunnel interfaces for LSYS1.

```
                                   [edit logical-systems LSYS1 interfaces ]
                                   user@host# set lt-0/0/0 unit 3 encapsulation ethernet
                                   user@host# set lt-0/0/0 unit 3 peer-unit 2
                                   user@host# set lt-0/0/0 unit 3 family inet address 2.1.1.3/24
```

   b. Configure logical tunnel interfaces for LSYS2.

```
                                   [edit logical-systems LSYS2 interfaces ]
                                   user@host# set lt-0/0/0 unit 5 encapsulation ethernet
                                   user@host# set lt-0/0/0 unit 5 peer-unit 4
                                   user@host# set lt-0/0/0 unit 5 family inet address 2.1.1.5/24
```

Results   From configuration mode, confirm the configuration for LSYS0 by entering the **show
           logical-systems LSYS0** command. If the output does not display the intended
           configuration, repeat the configuration instructions in this example to correct it.

```
                         [edit]
                         user@host# show logical-systems LSYS0
                         interfaces {
                           lt-0/0/0 {
                             unit 0 {
                               encapsulation ethernet-vpls;
                               peer-unit 1;
                             }
                             unit 2 {
                               encapsulation ethernet-vpls;
                               peer-unit 3;
                             }
                             unit 4 {
                               encapsulation ethernet-vpls;
                               peer-unit 5;
                             }
                           }
                         }
                           routing-instances {
                             vr {
                               instance-type vpls;
                               interface lt-0/0/0.0;
                               interface lt-0/0/0.2;
                               interface lt-0/0/0.4;
```

```
      }
    }
```

From configuration mode, confirm the configuration for the master logical system by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show interfaces
lt-0/0/0 {
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet {
      address 2.1.1.1/24;
    }
  }
}
  ge-1/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-1/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-1/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-1/0/3 {
    gigether-options {
      redundant-parent reth3;
    }
  }
  ge-7/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-7/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-7/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-7/0/3 {
    gigether-options {
```

```
                redundant-parent reth3;
            }
        }
        fab0 {
            fabric-options {
                member-interfaces {
                    ge-1/1/0;
                }
            }
        }
        fab1 {
            fabric-options {
                member-interfaces {
                    ge-7/1/0;
                }
            }
        }
        reth0 {
            redundant-ether-options {
                redundancy-group 1;
            }
            unit 0 {
                family inet {
                    address 95.99.99.1/8;
                }
            }
        }
        reth1 {
            redundant-ether-options {
                redundancy-group 1;
            }
        }
        reth2 {
            redundant-ether-options {
                redundancy-group 1;
            }
        }
        reth3 {
            redundant-ether-options {
                redundancy-group 1;
            }
        }
[edit]
user@host# show routing-instances
vr0 {
    instance-type virtual-router;
    interface lt-0/0/0.1;
    interface reth0.0;
    routing-options {
        static {
            route 85.0.0.0/8 next-hop 2.1.1.3;
            route 75.0.0.0/8 next-hop 2.1.1.3;
            route 65.0.0.0/8 next-hop 2.1.1.5;
        }
    }
}
```

```
[edit]
user@host# show security
policies {
  from-zone root-trust to-zone root-untrust {
    policy root-Trust_to_root-Untrust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone root-untrust to-zone root-trust {
    policy root-Untrust_to_root-Trust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone root-untrust to-zone root-untrust {
    policy root-Untrust_to_root-Untrust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone root-trust to-zone root-trust {
    policy root-Trust_to_root-Trust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
}
  zones {
    security-zone root-trust {
      host-inbound-traffic {
        system-services {
```

```
                        all;
                    }
                    protocols {
                        all;
                    }
                }
                interfaces {
                    reth0.0;
                }
            }
            security-zone root-untrust {
                host-inbound-traffic {
                    system-services {
                        all;
                    }
                    protocols {
                        all;
                    }
                }
                interfaces {
                    lt-0/0/0.1;
                }
            }
        }
    }
```

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

User Logical System Configuration (User Logical System Administrator)

CLI Quick
Configuration

To quickly configure user logical systems, 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.

Enter the following commands while logged in as the user logical system administrator for LSYS1:

```
set interfaces reth1 unit 0 family inet address 85.88.88.1/8
set interfaces reth2 unit 0 family inet address 75.77.77.1/8
set routing-instances vr11 instance-type virtual-router
set routing-instances vr11 interface lt-0/0/0.3
set routing-instances vr11 interface reth1.0
set routing-instances vr11 routing-options static route 65.0.0.0/8 next-hop 2.1.1.5
set routing-instances vr11 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
set routing-instances vr12 instance-type virtual-router
set routing-instances vr12 interface reth2.0
set routing-instances vr12 routing-options interface-routes rib-group inet vr11vr12v4
set routing-instances vr12 routing-options static route 85.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 95.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 65.0.0.0/8 next-table vr11.inet.0
set routing-instances vr12 routing-options static route 2.1.1.0/24 next-table vr11.inet.0
set routing-options rib-groups vr11vr12v4 import-rib vr11.inet.0
set routing-options rib-groups vr11vr12v4 import-rib vr12.inet.0
set security zones security-zone lsys1-trust host-inbound-traffic system-services all
set security zones security-zone lsys1-trust host-inbound-traffic protocols all
```

```
set security zones security-zone lsys1-trust interfaces reth1.0
set security zones security-zone lsys1-trust interfaces lt-0/0/0.3
set security zones security-zone lsys1-untrust host-inbound-traffic system-services all
set security zones security-zone lsys1-untrust host-inbound-traffic protocols all
set security zones security-zone lsys1-untrust interfaces reth2.0
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust match application any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust then permit
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    match application any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    then permit
```

Enter the following commands while logged in as the user logical system administrator
for LSYS2:

```
set interfaces reth3 unit 0 family inet address 65.66.66.1/8
set routing-instances vr2 instance-type virtual-router
set routing-instances vr2 interface lt-0/0/0.5
set routing-instances vr2 interface reth3.0
set routing-instances vr2 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr2 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3
set routing-instances vr2 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1
set security zones security-zone lsys2-trust host-inbound-traffic system-services all
set security zones security-zone lsys2-trust host-inbound-traffic protocols all
set security zones security-zone lsys2-trust interfaces reth3.0
set security zones security-zone lsys2-untrust host-inbound-traffic system-services all
set security zones security-zone lsys2-untrust host-inbound-traffic protocols all
set security zones security-zone lsys2-untrust interfaces lt-0/0/0.5
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust match source-address any
```

```
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust match application any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust then permit
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust match application any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust then permit
```

Step-by-Step Procedure

NOTE: The user logical system administrator performs the following configuration while logged in to his or her user logical system. The master administrator can also configure a user logical system at the [edit logical-systems *logical-system*] hierarchy level.

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

To configure the LSYS1 user logical system:

1. Configure interfaces.

    ```
    [edit interfaces]
    lsys1-admin@host:LSYS1# set reth1 unit 0 family inet address 85.88.88.1/8
    lsys1-admin@host:LSYS1# set reth2 unit 0 family inet address 75.77.77.1/8
    ```

2. Configure routing.

    ```
    [edit routing-instances]
    lsys1-admin@host:LSYS1# set vr11 instance-type virtual-router
    lsys1-admin@host:LSYS1# set vr11 interface lt-0/0/0.3
    ```

lsys1-admin@host:LSYS1# set vr11 interface reth1.0
lsys1-admin@host:LSYS1# set vr11 routing-options static route 65.0.0.0/8 next-hop
    2.1.1.5
lsys1-admin@host:LSYS1# set vr11 routing-options static route 95.0.0.0/8 next-hop
    2.1.1.1
lsys1-admin@host:LSYS1# set vr12 instance-type virtual-router
lsys1-admin@host:LSYS1# set vr12 interface reth2.0
lsys1-admin@host:LSYS1# set vr12 routing-options interface-routes rib-group inet
    vr11vr12v4
lsys1-admin@host:LSYS1# set vr12 routing-options static route 85.0.0.0/8 next-table
    vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 95.0.0.0/8 next-table
    vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 65.0.0.0/8 next-table
    vr11.inet.0
lsys1-admin@host:LSYS1# set vr12 routing-options static route 2.1.1.0/24 next-table
    vr11.inet.0

[edit routing-options]
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v4 import-rib vr11.inet.0
lsys1-admin@host:LSYS1# set rib-groups vr11vr12v4 import-rib vr12.inet.0

3. Configure zones and security policies.

[edit security zones]
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
    system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
    protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces reth1.0
lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces lt-0/0/0.3
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
    system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
    protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust interfaces reth2.0

[edit security policies from-zone lsys1-trust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match source-address
    any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match
    destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match application
    any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match source-address
    any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match
    destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match application
    any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
    source-address any

lsys1-admin@host:LSYS1# **set policy lsys1untrust-to-lsys1untrust match destination-address any**
lsys1-admin@host:LSYS1# **set policy lsys1untrust-to-lsys1untrust match application any**
lsys1-admin@host:LSYS1# **set policy lsys1untrust-to-lsys1untrust then permit**

[edit security policies from-zone lsys1-trust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# **set policy lsys1trust-to-lsys1trust match source-address any**
lsys1-admin@host:LSYS1# **set policy lsys1trust-to-lsys1trust match destination-address any**
lsys1-admin@host:LSYS1# **set policy lsys1trust-to-lsys1trust match application any**
lsys1-admin@host:LSYS1# **set policy lsys1trust-to-lsys1trust then permit**

Step-by-Step Procedure

To configure the LSYS2 user logical system:

1. Configure interfaces.

   [edit interfaces]
   lsys2-admin@host:LSYS2# **set reth3 unit 0 family inet address 65.66.66.1/8**

2. Configure routing.

   [edit routing-instances]
   lsys2-admin@host:LSYS2# **set vr2 instance-type virtual-router**
   lsys2-admin@host:LSYS2# **set vr2 interface lt-0/0/0.5**
   lsys2-admin@host:LSYS2# **set vr2 interface reth3.0**
   lsys2-admin@host:LSYS2# **set vr2 routing-options static route 75.0.0.0/8 next-hop 2.1.1.3**
   lsys2-admin@host:LSYS2# **set vr2 routing-options static route 85.0.0.0/8 next-hop 2.1.1.3**
   lsys2-admin@host:LSYS2# **set vr2 routing-options static route 95.0.0.0/8 next-hop 2.1.1.1**

3. Configure zones and security policies.

   [edit security zones]
   lsys2-admin@host:LSYS2# **set security-zone lsys2-trust host-inbound-traffic system-services all**
   lsys2-admin@host:LSYS2# **set security-zone lsys2-trust host-inbound-traffic protocols all**
   lsys2-admin@host:LSYS2# **set security-zone lsys2-trust interfaces reth3.0**
   lsys2-admin@host:LSYS2# **set security zones security-zone lsys2-untrust host-inbound-traffic system-services all**
   lsys2-admin@host:LSYS2# **set security-zone lsys2-untrust host-inbound-traffic protocols all**
   lsys2-admin@host:LSYS2# **set security-zone lsys2-untrust interfaces lt-0/0/0.5**

   [edit security policies from-zone lsys2-trust to-zone lsys2-untrust]
   lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust match source-address any**
   lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust match destination-address any**
   lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust match application any**
   lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust then permit**

   [edit security policies from-zone from-zone lsys2-untrust to-zone lsys2-trust]

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust match source-address any**

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust match destination-address any**

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust match application any**

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust then permit**

[edit security policies from-zone lsys2-untrust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust match source-address any**

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust match destination-address any**

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust match application any**

lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust then permit**

[edit security policies from-zone lsys2-trust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust match source-address any**

lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust match destination-address any**

lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust match application any**

lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust then permit**

Results   From configuration mode, confirm the configuration for LSYS1 by entering the **show interfaces**, **show routing-instances**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys1-admin@host:LSYS1# show interfaces
interfaces {
  lt-0/0/0 {
    unit 3 {
      encapsulation ethernet;
      peer-unit 2;
      family inet {
        address 2.1.1.3/24;
      }
    }
  }
  reth1 {
    unit 0 {
      family inet {
        address 85.88.88.1/8;
      }
    }
  }
  reth2 {
    unit 0 {
      family inet {
        address 75.77.77.1/8;
      }
    }
```

```
        }
      }
      [edit]
      lsys1-admin@host:LSYS1# show routing-instances
      routing-instances {
        vr11 {
          instance-type virtual-router;
          interface lt-0/0/0.3;
          interface reth1.0;
          routing-options {
            static {
              route 65.0.0.0/8 next-hop 2.1.1.5;
              route 95.0.0.0/8 next-hop 2.1.1.1;
            }
          }
        }
        vr12 {
          instance-type virtual-router;
          interface reth2.0;
          routing-options {
            interface-routes {
              rib-group inet vr11vr12v4;
            }
            static {
              route 85.0.0.0/8 next-table vr11.inet.0;
              route 95.0.0.0/8 next-table vr11.inet.0;
              route 65.0.0.0/8 next-table vr11.inet.0;
              route 2.1.1.0/24 next-table vr11.inet.0;
            }
          }
        }
      }
      [edit]
      lsys1-admin@host:LSYS1# show routing-options
      rib-groups {
        vr11vr12v4 {
          import-rib [ vr11.inet.0 vr12.inet.0 ];
        }
      }
      [edit]
      lsys1-admin@host:LSYS1# show security
      security {
        policies {
          from-zone lsys1-trust to-zone lsys1-untrust {
            policy lsys1trust-to-lsys1untrust {
              match {
                source-address any;
                destination-address any;
                application any;
              }
              then {
                permit;
              }
            }
          }
          from-zone lsys1-untrust to-zone lsys1-trust {
```

```
                    policy lsys1untrust-to-lsys1trust {
                      match {
                        source-address any;
                        destination-address any;
                        application any;
                      }
                      then {
                        permit;
                      }
                    }
                  }
                  from-zone lsys1-untrust to-zone lsys1-untrust {
                    policy lsys1untrust-to-lsys1untrust {
                      match {
                        source-address any;
                        destination-address any;
                        application any;
                      }
                      then {
                        permit;
                      }
                    }
                  }
                  from-zone lsys1-trust to-zone lsys1-trust {
                    policy lsys1trust-to-lsys1trust {
                      match {
                        source-address any;
                        destination-address any;
                        application any;
                      }
                      then {
                        permit;
                      }
                    }
                  }
                }
                zones {
                  security-zone lsys1-trust {
                    host-inbound-traffic {
                      system-services {
                        all;
                      }
                      protocols {
                        all;
                      }
                    }
                    interfaces {
                      reth1.0;
                      lt-0/0/0.3;
                    }
                  }
                  security-zone lsys1-untrust {
                    host-inbound-traffic {
                      system-services {
                        all;
                      }
                  }
```

```
        protocols {
          all;
        }
      }
      interfaces {
        reth2.0;
      }
    }
  }
}
```

From configuration mode, confirm the configuration for LSYS2 by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
lsys2-admin@host:LSYS2# show interfaces
[edit]
  interfaces {
  lt-0/0/0 {
    unit 5 {
      encapsulation ethernet;
      peer-unit 4;
      family inet {
        address 2.1.1.5/24;
      }
    }
  }
  reth3 {
    unit 0 {
      family inet {
        address 65.66.66.1/8;
      }
    }
  }
}
[edit]
lsys2-admin@host:LSYS2# show routing-instances
routing-instances {
  vr2 {
    instance-type virtual-router;
    interface lt-0/0/0.5;
    interface reth3.0;
    routing-options {
      static {
        route 75.0.0.0/8 next-hop 2.1.1.3;
        route 85.0.0.0/8 next-hop 2.1.1.3;
        route 95.0.0.0/8 next-hop 2.1.1.1;
      }
    }
  }
}
[edit]
lsys2-admin@host:LSYS2# show security
security {
  policies {
```

```
                    from-zone lsys2-trust to-zone lsys2-untrust {
                        policy lsys2trust-to-lsys2untrust {
                            match {
                                source-address any;
                                destination-address any;
                                application any;
                            }
                            then {
                                permit;
                            }
                        }
                    }
                    from-zone lsys2-untrust to-zone lsys2-trust {
                        policy lsys2untrust-to-lsys2trust {
                            match {
                                source-address any;
                                destination-address any;
                                application any;
                            }
                            then {
                                permit;
                            }
                        }
                    }
                    from-zone lsys2-untrust to-zone lsys2-untrust {
                        policy lsys2untrust-to-lsys2untrust {
                            match {
                                source-address any;
                                destination-address any;
                                application any;
                            }
                            then {
                                permit;
                            }
                        }
                    }
                    from-zone lsys2-trust to-zone lsys2-trust {
                        policy lsys2trust-to-lsys2trust {
                            match {
                                source-address any;
                                destination-address any;
                                application any;
                            }
                            then {
                                permit;
                            }
                        }
                    }
                }
                zones {
                    security-zone lsys2-trust {
                        host-inbound-traffic {
                            system-services {
                                all;
                            }
                            protocols {
```

```
                all;
              }
            }
            interfaces {
              reth3.0;
            }
          }
          security-zone lsys2-untrust {
            host-inbound-traffic {
              system-services {
                all;
              }
              protocols {
                all;
              }
            }
            interfaces {
              lt-0/0/0.5;
            }
          }
        }
      }
    }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying Chassis Cluster Status on page 195
- Troubleshooting Chassis Cluster with Logs on page 196
- Verifying Logical System Licenses on page 196
- Verifying Logical System License Usage on page 196
- Verifying Intra-Logical System Traffic on a Logical System on page 197
- Verifying Intra-Logical System Traffic Within All Logical Systems on page 197
- Verifying Traffic Between User Logical Systems on page 198

Verifying Chassis Cluster Status

Purpose Verify the chassis cluster status, failover status, and redundancy group information.

Action From operational mode, enter the **show chassis cluster status** command.

```
{primary:node0}
show chassis cluster status
Cluster ID: 1
Node                    Priority          Status     Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                 200             primary       no       no
    node1                 100             secondary     no       no

Redundancy group: 1 , Failover count: 1
```

```
node0                          200        primary        no        no
node1                          100        secondary      no        no
```

Troubleshooting Chassis Cluster with Logs

Purpose   Identify any chassis cluster issues by looking at the logs on both nodes.

Action   From operational mode, enter these **show log** commands.

```
user@host> show log jsrpd
user@host> show log chassisd
user@host> show log messages
user@host> show log dcd
user@host> show traceoptions
```

Verifying Logical System Licenses

Purpose   Verify information about logical system licenses.

Action   From operational mode, enter the **show system license status logical-system all** command.

```
{primary:node0}
user@host> show system license status logical-system all
node0:
--------------------------------------------------------------------------
Logical system license status:

logical system name                    license status
root-logical-system                    enabled
LSYS0                                   enabled
LSYS1                                   enabled
LSYS2                                   enabled
```

Verifying Logical System License Usage

Purpose   Verify information about logical system license usage.

> *i*   NOTE: The actual number of licenses used is only displayed on the primary node.

Action   From operational mode, enter the **show system license** command.

```
{primary:node0}
user@host> show system license
License usage:
                               Licenses     Licenses     Licenses     Expiry
  Feature name                    used     installed      needed
  logical-system                     4            25            0     permanent

Licenses installed:
  License identifier: JUNOS305013
  License version: 2
  Valid for device: JN110B54BAGB
  Features:
```

```
      logical-system-25 - Logical System Capacity
        permanent
```

Verifying Intra-Logical System Traffic on a Logical System

Purpose   Verify information about currently active security sessions within a logical system.

Action   From operational mode, enter the **show security flow session logical-system LSYS1** command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000114, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1782, Valid
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 33, Bytes: 1881

  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 28, Bytes: 2329
Total sessions: 1

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:

Session ID: 90000001, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14388, Valid
  In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

Verifying Intra-Logical System Traffic Within All Logical Systems

Purpose   Verify information about currently active security sessions on all logical systems.

Action   From operational mode, enter the **show security flow session logical-system all** command.

```
{primary:node0}
user@host> show security flow session logical-system all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
```

```
                       Total sessions: 0

                       Flow Sessions on FPC2 PIC0:
                       Total sessions: 0

                       Flow Sessions on FPC2 PIC1:

                       Session ID: 90000114, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
                       Timeout: 1776, Valid
                       Logical system: LSYS1
                         In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 33, Bytes: 1881

                         Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 28, Bytes: 2329
                       Total sessions: 1

                       node1:
                       --------------------------------------------------------------------------

                       Flow Sessions on FPC0 PIC1:
                       Total sessions: 0

                       Flow Sessions on FPC2 PIC0:
                       Total sessions: 0

                       Flow Sessions on FPC2 PIC1:

                       Session ID: 90000001, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
                       Timeout: 14382, Valid
                       Logical system: LSYS1
                         In: 85.88.88.2/34538 --> 75.77.77.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
                         Out: 75.77.77.2/23 --> 85.88.88.2/34538;tcp, If: reth2.0, Pkts: 0, Bytes: 0
                       Total sessions: 1
```

Verifying Traffic Between User Logical Systems

Purpose   Verify information about currently active security sessions between logical systems.

Action   From operational mode, enter the **show security flow session logical-system**
logical-system-name command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000094, Policy name: root-Untrust_to_root-Trust/5, State: Active,
Timeout: 1768, Valid
  In: 75.77.77.2/34590 --> 95.99.99.2/23;tcp, If: lt-0/0/0.1, Pkts: 23, Bytes:
1351
  Out: 95.99.99.2/23 --> 75.77.77.2/34590;tcp, If: reth0.0, Pkts: 22, Bytes: 1880
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

```
node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000002, Policy name: root-Untrust_to_root-Trust/5, State: Backup,
Timeout: 14384, Valid
  In: 75.77.77.2/34590 --> 95.99.99.2/23;tcp, If: lt-0/0/0.1, Pkts: 0, Bytes: 0
  Out: 95.99.99.2/23 --> 75.77.77.2/34590;tcp, If: reth0.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system LSYS2

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000089, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1790, Valid
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 40, Bytes:
2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 32, Bytes: 2114
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000002, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14398, Valid
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system all

node0:
--------------------------------------------------------------------------
```

```
Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000088, Policy name: lsys1trust-to-lsys1trust/11, State: Active,
Timeout: 1782, Valid
Logical system: LSYS1
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: reth1.0, Pkts: 40, Bytes: 2252

  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: lt-0/0/0.3, Pkts: 32, Bytes:
2114

Session ID: 80000089, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1782, Valid
Logical system: LSYS2
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 40, Bytes:
2252
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 32, Bytes: 2114
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000001, Policy name: lsys1trust-to-lsys1trust/11, State: Backup,
Timeout: 14382, Valid
Logical system: LSYS1
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0

Session ID: 80000002, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14390, Valid
Logical system: LSYS2
  In: 85.88.88.2/34539 --> 65.66.66.2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 65.66.66.2/23 --> 85.88.88.2/34539;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Related Documentation

- Understanding Logical Systems in the Context of Chassis Cluster on page 167

- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 201

- *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*

- *Chassis Cluster Overview*

Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only)

This example shows how to configure logical systems in a basic active/passive chassis cluster with IPv6 addresses.

> **NOTE:** The master administrator configures the chassis cluster and creates logical systems (including an optional interconnect logical system), administrators, and security profiles. Either the master administrator or the user logical system administrator configures a user logical system. The configuration is synchronized between nodes in the cluster.

- Requirements on page 201
- Overview on page 202
- Configuration on page 205
- Verification on page 229

Requirements

Before you begin:

- Obtain two high-end SRX Series Services Gateways with identical hardware configurations. See *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*. This chassis cluster deployment scenario includes the configuration of the SRX Series device for connections to an MX240 edge router and an EX8208 Ethernet Switch.

- Physically connect the two devices (back-to-back for the fabric and control ports) and ensure that they are the same models. You can configure both the fabric and control ports on the SRX5000 line.

- Set the chassis cluster ID and node ID on each device and reboot the devices to enable clustering. See *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*.

> **NOTE:** For this example, chassis cluster and logical system configuration is performed on the primary (node 0) device at the root level by the master administrator. Log in to the device as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.

NOTE: When you use SRX Series devices running logical systems in a chassis cluster, you must purchase and install the same number of logical system licenses for each node in the chassis cluster. Logical system licenses pertain to a single chassis or node within a chassis cluster and not to the cluster collectively. See "Understanding Licenses for Logical Systems on SRX Series Devices" on page 7.

Overview

In this example, the basic active/passive chassis cluster consists of two devices:

- One device actively provides logical systems, along with maintaining control of the chassis cluster.

- The other device passively maintains its state for cluster failover capabilities should the active device become inactive.

NOTE: Logical systems in an active/active chassis cluster are configured in a similar manner as for logical systems in an active/passive chassis cluster. For active/active chassis clusters, there can be multiple redundancy groups that can be primary on different nodes.

The master administrator configures the following logical systems on the primary device (node 0):

- Master logical system—The master administrator configures a security profile to provision portions of the system's security resources to the master logical system and configures the resources of the master logical system.

- User logical systems LSYS1 and LSYS2 and their administrators—The master administrator also configures security profiles to provision portions of the system's security resources to user logical systems. The user logical system administrator can then configure interfaces, routing, and security resources allocated to his or her logical system.

- Interconnect logical system LSYS0 that connects logical systems on the device—The master administrator configures logical tunnel interfaces between the interconnect logical system and each logical system. These peer interfaces effectively allow for the establishment of tunnels.

> **NOTE:** This example does not describe configuring features such as NAT, IDP, or VPNs for a logical system. See "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 and "User Logical System Configuration Overview" on page 23 for more information about features that can be configured for logical systems.
>
> If you are performing proxy ARP in a chassis cluster configuration, you must apply the proxy ARP configuration to the reth interfaces rather than the member interfaces because the reth interfaces contain the logical configurations. See *Configuring Proxy ARP (CLI Procedure)*.

Topology

Figure 6 on page 204 shows the topology used in this example.

Figure 6: Logical Systems in a Chassis Cluster (IPv6)

Configuration

Chassis Cluster Configuration with IPv6 Addresses (Master Administrator)

CLI Quick Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, 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.

On {primary:node0}

```
set chassis cluster control-ports fpc 0 port 0
set chassis cluster control-ports fpc 6 port 0
set interfaces fab0 fabric-options member-interfaces ge-1/1/0
set interfaces fab1 fabric-options member-interfaces ge-7/1/0
set groups node0 system host-name SRX5800-1
set groups node0 interfaces fxp0 unit 0 family inet address 10.157.90.24/9
set groups node0 system backup-router 10.157.64.1 destination 0.0.0.0/0
set groups node1 system host-name SRX5800-2
set groups node1 interfaces fxp0 unit 0 family inet address 10.157.90.23/19
set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0
set apply-groups "${node}"
set chassis cluster reth-count 5
set chassis cluster redundancy-group 0 node 0 priority 200
set chassis cluster redundancy-group 0 node 1 priority 100
set chassis cluster redundancy-group 1 node 0 priority 200
set chassis cluster redundancy-group 1 node 1 priority 100
set interfaces ge-1/0/0 gigether-options redundant-parent reth0
set interfaces ge-1/0/1 gigether-options redundant-parent reth1
set interfaces ge-1/0/2 gigether-options redundant-parent reth2
set interfaces ge-1/0/3 gigether-options redundant-parent reth3
set interfaces ge-7/0/0 gigether-options redundant-parent reth0
set interfaces ge-7/0/1 gigether-options redundant-parent reth1
set interfaces ge-7/0/2 gigether-options redundant-parent reth2
set interfaces ge-7/0/3 gigether-options redundant-parent reth3
set interfaces reth0 redundant-ether-options redundancy-group 1
set interfaces reth0 unit 0 family inet6 address 9995::1/64
set interfaces reth1 redundant-ether-options redundancy-group 1
set interfaces reth2 redundant-ether-options redundancy-group 1
set interfaces reth3 redundant-ether-options redundancy-group 1
```

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

To configure a chassis cluster:

> NOTE: Perform the following steps on the primary device (node 0). They are automatically copied over to the secondary device (node 1) when you execute a **commit** command.

1. Configure control ports for the clusters.

   [edit chassis cluster]
   user@host# **set control-ports fpc 0 port 0**
   user@host# **set control-ports fpc 6 port 0**

2. Configure the fabric (data) ports of the cluster that are used to pass RTOs in active/passive mode.

   [edit interfaces]
   user@host# **set fab0 fabric-options member-interfaces ge-1/1/0**
   user@host# **set fab1 fabric-options member-interfaces ge-7/1/0**

3. Assign some elements of the configuration to a specific member. Configure out-of-band management on the fxp0 interface of the SRX Services Gateway using separate IP addresses for the individual control planes of the cluster.

   [edit]
   user@host# **set groups node0 system host-name SRX5800-1**
   user@host# **set groups node0 interfaces fxp0 unit 0 family inet address**
       10.157.90.24/9
   user@host# **set groups node0 system backup-router 10.157.64.1 destination**
       0.0.0.0/0
   user@host# **set groups node1 system host-name SRX5800-2**
   user@host# **set groups node1 interfaces fxp0 unit 0 family inet address**
       10.157.90.23/19
   user@host# **set groups node1 system backup-router 10.157.64.1 destination 0.0.0.0/0**
   user@host# **set apply-groups "${node}"**

4. Configure redundancy groups for chassis clustering.

   [edit chassis cluster]
   user@host# **set reth-count 5**
   user@host# **set redundancy-group 0 node 0 priority 200**
   user@host# **set redundancy-group 0 node 1 priority 100**
   user@host# **set redundancy-group 1 node 0 priority 200**
   user@host# **set redundancy-group 1 node 1 priority 100**

5. Configure the data interfaces on the platform so that in the event of a data plane failover, the other chassis cluster member can take over the connection seamlessly.

   [edit interfaces]
   user@host# **set ge-1/0/0 gigether-options redundant-parent reth0**
   user@host# **set ge-1/0/1 gigether-options redundant-parent reth1**
   user@host# **set ge-1/0/2 gigether-options redundant-parent reth2**

```
user@host# set ge-1/0/3 gigether-options redundant-parent reth3
user@host# set ge-7/0/0 gigether-options redundant-parent reth0
user@host# set ge-7/0/1 gigether-options redundant-parent reth1
user@host# set ge-7/0/2 gigether-options redundant-parent reth2
user@host# set ge-7/0/3 gigether-options redundant-parent reth3
user@host# set reth0 redundant-ether-options redundancy-group 1
user@host# set reth0 unit 0 family inet6 address 9995::1/64
user@host# set reth1 redundant-ether-options redundancy-group 1
user@host# set reth2 redundant-ether-options redundancy-group 1
user@host# set reth3 redundant-ether-options redundancy-group 1
```

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

```
user@host> show configuration
version ;
groups {
    node0 {
        system {
            host-name SRX58001;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.24/9;
                    }
                }
            }
        }
    }
    node1 {
        system {
            host-name SRX58002;
            backup-router 10.157.64.1 destination 0.0.0.0/0;

        }
        interfaces {
            fxp0 {
                unit 0 {
                    family inet {
                        address 10.157.90.23/19;
                    }
                }
            }
        }
    }
}
apply-groups "${node}";
chassis {
    cluster {
        control-link-recovery;
        reth-count 5;
        control-ports {
            fpc 0 port 0;
            fpc 6 port 0;
```

```
                    }
                    redundancy-group 0 {
                        node 0 priority 200;
                        node 1 priority 100;
                    }
                    redundancy-group 1 {
                        node 0 priority 200;
                        node 1 priority 100;
                    }
                }
            }
            interfaces {
                ge-1/0/0 {
                    gigether-options {
                        redundant-parent reth0;
                    }
                }
                ge-1/0/1 {
                    gigether-options {
                        redundant-parent reth1;
                    }
                }
                ge-1/0/2 {
                    gigether-options {
                        redundant-parent reth2;
                    }
                }
                ge-1/0/3 {
                    gigether-options {
                        redundant-parent reth3;
                    }
                }
                ge-7/0/0 {
                    gigether-options {
                        redundant-parent reth0;
                    }
                }
                ge-7/0/1 {
                    gigether-options {
                        redundant-parent reth1;
                    }
                }
                ge-7/0/2 {
                    gigether-options {
                        redundant-parent reth2;
                    }
                }
                ge-7/0/3 {
                    gigether-options {
                        redundant-parent reth3;
                    }
                }
                fab0 {
                    fabric-options {
                        member-interfaces {
                            ge-1/1/0;
                        }
                    }
                }
                fab1 {
                    fabric-options {
```

```
                    member-interfaces {
                        ge-7/1/0;
                    }
                }
            }
            reth0 {
                redundant-ether-options {
                    redundancy-group 1;
                }
                unit 0 {
                    family inet6 {
                        address 9995::1/64;
                    }
                }
            }
            reth1 {
                redundant-ether-options {
                    redundancy-group 1;
                }
            }
            reth2 {
                redundant-ether-options {
                    redundancy-group 1;
                }
            }
            reth3 {
                redundant-ether-options {
                    redundancy-group 1;
                }
            }
        }
}
```

Logical System Configuration with IPv6 Addresses (Master Administrator)

CLI Quick
Configuration

To quickly create logical systems and user logical system administrators and configure the master and interconnect logical systems, 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.

> *i* NOTE: You are prompted to enter and then reenter plain-text passwords.

On {primary:node0}

```
set logical-systems LSYS1
set logical-systems LSYS2
set logical-systems LSYS0
set system login class lsys1 logical-system LSYS1
set system login class lsys1 permissions all
set system login user lsys1admin full-name lsys1-admin
set system login user lsys1admin class lsys1
set user lsys1admin authentication plain-text-password
set system login class lsys2 logical-system LSYS2
set system login class lsys2 permissions all
set system login user lsys2admin full-name lsys2-admin
```

```
set system login user lsys2admin class lsys2
set system login user lsys2admin authentication plain-text-password
set system security-profile SP-root policy maximum 200
set system security-profile SP-root policy reserved 100
set system security-profile SP-root zone maximum 200
set system security-profile SP-root zone reserved 100
set system security-profile SP-root flow-session maximum 200
set system security-profile SP-root flow-session reserved 100
set system security-profile SP-root root-logical-system
set system security-profile SP0 logical-system LSYS0
set system security-profile SP1 policy maximum 100
set system security-profile SP1 policy reserved 50
set system security-profile SP1 zone maximum 100
set system security-profile SP1 zone reserved 50
set system security-profile SP1 flow-session maximum 100
set system security-profile SP1 flow-session reserved 50
set system security-profile SP1 logical-system LSYS1
set system security-profile SP2 policy maximum 100
set system security-profile SP2 policy reserved 50
set system security-profile SP2 zone maximum 100
set system security-profile SP2 zone reserved 50
set system security-profile SP2 flow-session maximum 100
set system security-profile SP2 flow-session reserved 50
set system security-profile SP2 logical-system LSYS2
set interfaces lt-0/0/0 unit 1 encapsulation ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet6 address 2111::1/64
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface reth0.0
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 8885::/64 next-hop
    2111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 7775::/64 next-hop
    2111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 6665::/64 next-hop
    2111::5
set security zones security-zone root-trust host-inbound-traffic system-services all
set security zones security-zone root-trust host-inbound-traffic protocols all
set security zones security-zone root-trust interfaces reth0.0
set security zones security-zone root-untrust host-inbound-traffic system-services all
set security zones security-zone root-untrust host-inbound-traffic protocols all
set security zones security-zone root-untrust interfaces lt-0/0/0.1
set security policies from-zone root-trust to-zone root-untrust policy
    root-Trust_to_root-Untrust match source-address any
set security policies from-zone root-trust to-zone root-untrust policy
    root-Trust_to_root-Untrust match destination-address any
set security policies from-zone root-trust to-zone root-untrust policy
    root-Trust_to_root-Untrust match application any
set security policies from-zone root-trust to-zone root-untrust policy
    root-Trust_to_root-Untrust then permit
set security policies from-zone root-untrust to-zone root-trust policy
    root-Untrust_to_root-Trust match source-address any
set security policies from-zone root-untrust to-zone root-trust policy
    root-Untrust_to_root-Trust match destination-address any
set security policies from-zone root-untrust to-zone root-trust policy
    root-Untrust_to_root-Trust match application any
```

set security policies from-zone root-untrust to-zone root-trust policy
    root-Untrust_to_root-Trust then permit
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust match source-address any
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust match destination-address any
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust match application any
set security policies from-zone root-untrust to-zone root-untrust policy
    root-Untrust_to_root-Untrust then permit
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    match source-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    match destination-address any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    match application any
set security policies from-zone root-trust to-zone root-trust policy root-Trust_to_root-Trust
    then permit
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet6 address 2111::3/64
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet6 address 2111::5/64

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

To create logical systems and user logical system administrators and configure the master and interconnect logical systems:

1. Create the interconnect and user logical systems.

    [edit logical-systems]
    user@host# **set LSYS0**
    user@host# **set LSYS1**
    user@host# **set LSYS2**

2. Configure user logical system administrators.

    a. Configure the user logical system administrator for LSYS1.

        [edit system login]
        user@host# **set class lsys1 logical-system LSYS1**
        user@host# **set class lsys1 permissions all**
        user@host# **set user lsys1admin full-name lsys1-admin**

user@host# set user lsys1admin class lsys1
user@host# set user lsys1admin authentication plain-text-password

b. Configure the user logical system administrator for LSYS2.

[edit system login]
user@host# set class lsys2 logical-system LSYS2
user@host# set class lsys2 permissions all
user@host# set user lsys2admin full-name lsys2-admin
user@host# set user lsys2admin class lsys2
user@host# set user lsys2admin authentication plain-text-password

3. Configure security profiles and assign them to logical systems.

a. Configure a security profile and assign it to the root logical system.

[edit system security-profile]
user@host# set SP-root policy maximum 200
user@host# set SP-root policy reserved 100
user@host# set SP-root zone maximum 200
user@host# set SP-root zone reserved 100
user@host# set SP-root flow-session maximum 200
user@host# set SP-root flow-session reserved 100
user@host# set SP-root root-logical-system

b. Assign a dummy security profile containing no resources to the interconnect
    logical system LSYS0.

[edit system security-profile]
user@host# set SP0 logical-system LSYS0

c. Configure a security profile and assign it to LSYS1.

[edit system security-profile]
user@host# set SP1 policy maximum 100
user@host# set SP1 policy reserved 50
user@host# set SP1 zone maximum 100
user@host# set SP1 zone reserved 50
user@host# set SP1 flow-session maximum 100
user@host# set SP1 flow-session reserved 50
user@host# set SP1 logical-system LSYS1

d. Configure a security profile and assign it to LSYS2.

[edit system security-profile]
user@host# set SP2 policy maximum 100
user@host# set SP2 policy reserved 50
user@host# set SP2 zone maximum 100
user@host# set SP2 zone reserved 50
user@host# set SP2 flow-session maximum 100
user@host# set SP2 flow-session reserved 50
user@host# set SP2 logical-system LSYS2

4. Configure the master logical system.

a. Configure logical tunnel interfaces.

[edit interfaces]
user@host# set lt-0/0/0 unit 1 encapsulation ethernet

```
user@host# set lt-0/0/0 unit 1 peer-unit 0
user@host# set lt-0/0/0 unit 1 family inet6 address 2111::1/64
```

b. Configure a routing instance.

```
[edit routing-instances]
user@host# set vr0 instance-type virtual-router
user@host# set vr0 interface lt-0/0/0.1
user@host# set vr0 interface reth0.0
user@host# set vr0 routing-options rib vr0.inet6.0 static route 8885::/64
    next-hop 2111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 7775::/64 next-hop
    2111::3
user@host# set vr0 routing-options rib vr0.inet6.0 static route 6665::/64
    next-hop 2111::5
```

c. Configure zones.

```
[edit security zones]
user@host# set security-zone root-trust host-inbound-traffic system-services
    all
user@host# set security-zone root-trust host-inbound-traffic protocols all
user@host# set security-zone root-trust interfaces reth0.0
user@host# set security-zone root-untrust host-inbound-traffic system-services
    all
user@host# set security-zone root-untrust host-inbound-traffic protocols all
user@host# set security-zone root-untrust interfaces lt-0/0/0.1
```

d. Configure security policies.

```
[edit security policies from-zone root-trust to-zone root-untrust]
user@host# set policy root-Trust_to_root-Untrust match source-address any
user@host# set policy root-Trust_to_root-Untrust match destination-address
    any
user@host# set policy root-Trust_to_root-Untrust match application any
user@host# set policy root-Trust_to_root-Untrust then permit

[edit security policies from-zone root-untrust to-zone root-trust]
user@host# set policy root-Untrust_to_root-Trust match source-address any
user@host# set policy root-Untrust_to_root-Trust match destination-address
    any
user@host# set policy root-Untrust_to_root-Trust match application any
user@host# set policy root-Untrust_to_root-Trust then permit

[edit security policies from-zone root-untrust to-zone root-untrust]
user@host# set policy root-Untrust_to_root-Untrust match source-address any
user@host# set policy root-Untrust_to_root-Untrust match destination-address
    any
user@host# set policy root-Untrust_to_root-Untrust match application any
user@host# set policy root-Untrust_to_root-Untrust then permit

[edit security policies from-zone root-trust to-zone root-trust]
user@host# set policy root-Trust_to_root-Trust match source-address any
user@host# set policy root-Trust_to_root-Trust match destination-address any
user@host# set policy root-Trust_to_root-Trust match application any
```

> user@host# **set policy root-Trust_to_root-Trust then permit**

5. Configure the interconnect logical system.

   a. Configure logical tunnel interfaces.

      [edit logical-systems LSYS0 interfaces]
      user@host# **set lt-0/0/0 unit 0 encapsulation ethernet-vpls**
      user@host# **set lt-0/0/0 unit 0 peer-unit 1**
      user@host# **set lt-0/0/0 unit 2 encapsulation ethernet-vpls**
      user@host# **set lt-0/0/0 unit 2 peer-unit 3**
      user@host# **set lt-0/0/0 unit 4 encapsulation ethernet-vpls**
      user@host# **set lt-0/0/0 unit 4 peer-unit 5**

   b. Configure the VPLS routing instance.

      [edit logical-systems LSYS0 routing-instances]
      user@host# **set vr instance-type vpls**
      user@host# **set vr interface lt-0/0/0.0**
      user@host# **set vr interface lt-0/0/0.2**
      user@host# **set vr interface lt-0/0/0.4**

6. Configure logical tunnel interfaces for the user logical systems.

   a. Configure logical tunnel interfaces for LSYS1.

      [edit logical-systems LSYS1 interfaces]
      user@host# **set lt-0/0/0 unit 3 encapsulation ethernet**
      user@host# **set lt-0/0/0 unit 3 peer-unit 2**
      user@host# **set lt-0/0/0 unit 3 family inet6 address 2111::3/64**

   b. Configure logical tunnel interfaces for LSYS2.

      [edit logical-systems LSYS2 interfaces]
      user@host# **set lt-0/0/0 unit 5 encapsulation ethernet**
      user@host# **set lt-0/0/0 unit 5 peer-unit 4**
      user@host# **set lt-0/0/0 unit 5 family inet6 address 2111::5/64**

Results   From configuration mode, confirm the configuration for LSYS0 by entering the **show logical-systems LSYS0** command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show logical-systems LSYS0
interfaces {
  lt-0/0/0 {
    unit 0 {
      encapsulation ethernet-vpls;
      peer-unit 1;
    }
    unit 2 {
      encapsulation ethernet-vpls;
      peer-unit 3;
    }
    unit 4 {
      encapsulation ethernet-vpls;
      peer-unit 5;
```

```
      }
    }
  }
  routing-instances {
    vr {
      instance-type vpls;
      interface lt-0/0/0.0;
      interface lt-0/0/0.2;
      interface lt-0/0/0.4;
    }
  }
```

From configuration mode, confirm the configuration for the master logical system by
entering the **show interfaces**, **show routing-instances**, and **show security** commands. If
the output does not display the intended configuration, repeat the configuration
instructions in this example to correct it.

```
[edit]
user@host# show interfaces
lt-0/0/0 {
  unit 1 {
    encapsulation ethernet;
    peer-unit 0;
    family inet6 {
      address 2111::1/64;
    }
  }
}
  ge-1/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-1/0/1 {
    gigether-options {
      redundant-parent reth1;
    }
  }
  ge-1/0/2 {
    gigether-options {
      redundant-parent reth2;
    }
  }
  ge-1/0/3 {
    gigether-options {
      redundant-parent reth3;
    }
  }
  ge-7/0/0 {
    gigether-options {
      redundant-parent reth0;
    }
  }
  ge-7/0/1 {
    gigether-options {
      redundant-parent reth1;
```

```
        }
      }
      ge-7/0/2 {
        gigether-options {
          redundant-parent reth2;
        }
      }
      ge-7/0/3 {
        gigether-options {
          redundant-parent reth3;
        }
      }
      fab0 {
        fabric-options {
          member-interfaces {
            ge-1/1/0;
          }
        }
      }
      fab1 {
        fabric-options {
          member-interfaces {
            ge-7/1/0;
          }
        }
      }
      reth0 {
        redundant-ether-options {
          redundancy-group 1;
        }
        unit 0 {
          family inet6 {
            address 9995::1/64;
          }
        }
      }
      reth1 {
        redundant-ether-options {
          redundancy-group 1;
        }
      }
      reth2 {
        redundant-ether-options {
          redundancy-group 1;
        }
      }
      reth3 {
        redundant-ether-options {
          redundancy-group 1;
        }
      }
[edit]
user@host# show routing-instances
vr0 {
  instance-type virtual-router;
  interface lt-0/0/0.1;
```

```
interface reth0.0;
routing-options {
  rib vr0.inet6.0 {
    static {
      route 8885::/64 next-hop 2111::3;
      route 7775::/64 next-hop 2111::3;
      route 6665::/64 next-hop 2111::5;
    }
  }
}
}
[edit]
user@host# show security
policies {
  from-zone root-trust to-zone root-untrust {
    policy root-Trust_to_root-Untrust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone root-untrust to-zone root-trust {
    policy root-Untrust_to_root-Trust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone root-untrust to-zone root-untrust {
    policy root-Untrust_to_root-Untrust {
      match {
        source-address any;
        destination-address any;
        application any;
      }
      then {
        permit;
      }
    }
  }
  from-zone root-trust to-zone root-trust {
    policy root-Trust_to_root-Trust {
      match {
        source-address any;
        destination-address any;
        application any;
```

```
                    }
                  then {
                    permit;
                  }
                }
              }
            }
          }
        zones {
          security-zone root-trust {
            host-inbound-traffic {
              system-services {
                all;
              }
              protocols {
                all;
              }
            }
            interfaces {
              reth0.0;
            }
          }
          security-zone root-untrust {
            host-inbound-traffic {
              system-services {
                all;
              }
              protocols {
                all;
              }
            }
            interfaces {
              lt-0/0/0.1;
            }
          }
        }
      }
```

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

User Logical System Configuration with IPv6 (User Logical System Administrator)

CLI Quick Configuration

To quickly configure user logical systems, 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.

Enter the following commands while logged in as the user logical system administrator for LSYS1:

```
set interfaces reth1 unit 0 family inet6 address 8885::1/64
set interfaces reth2 unit 0 family inet6 address 7775::1/64
set routing-instances vr11 instance-type virtual-router
set routing-instances vr11 interface lt-0/0/0.3
set routing-instances vr11 interface reth1.0
set routing-instances vr11 routing-options rib vr11.inet6.0 static route 6665::/64 next-hop
    2111::5
```

set routing-instances vr11 routing-options rib vr11.inet6.0 static route 9995::/64 next-hop
    2111::1
set routing-instances vr12 instance-type virtual-router
set routing-instances vr12 interface reth2.0
set routing-instances vr12 routing-options interface-routes rib-group inet6 vr11vr12v6
set routing-instances vr12 rrouting-options rib vr12.inet6.0 static route 8885::/64
    next-table vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 9995::/64 next-table
    vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 6665::/64 next-table
    vr11.inet6.0
set routing-instances vr12 routing-options rib vr12.inet6.0 static route 2111::/64 next-table
    vr11.inet6.0
set routing-options rib-groups vr11vr12v6 import-rib vr11.inet6.0
set routing-options rib-groups vr11vr12v6 import-rib vr12.inet6.0
set security zones security-zone lsys1-trust host-inbound-traffic system-services all
set security zones security-zone lsys1-trust host-inbound-traffic protocols all
set security zones security-zone lsys1-trust interfaces reth1.0
set security zones security-zone lsys1-trust interfaces lt-0/0/0.3
set security zones security-zone lsys1-untrust host-inbound-traffic system-services all
set security zones security-zone lsys1-untrust host-inbound-traffic protocols all
set security zones security-zone lsys1-untrust interfaces reth2.0
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust match application any
set security policies from-zone lsys1-trust to-zone lsys1-untrust policy
    lsys1trust-to-lsys1untrust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-trust policy
    lsys1untrust-to-lsys1trust then permit
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust match source-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust match destination-address any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust match application any
set security policies from-zone lsys1-untrust to-zone lsys1-untrust policy
    lsys1untrust-to-lsys1untrust then permit
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    match source-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    match destination-address any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    match application any
set security policies from-zone lsys1-trust to-zone lsys1-trust policy lsys1trust-to-lsys1trust
    then permit

Enter the following commands while logged in as the user logical system administrator for LSYS2:

```
set interfaces reth3 unit 0 family inet6 address 6665::1/64
set routing-instances vr2 instance-type virtual-router
set routing-instances vr2 interface lt-0/0/0.5
set routing-instances vr2 interface reth3.0
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 7775::/64 next-hop
    2111::3
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 8885::/64 next-hop
    2111::3
set routing-instances vr2 routing-options rib vr2.inet6.0 static route 9995::/64 next-hop
    2111::1
set security zones security-zone lsys2-trust host-inbound-traffic system-services all
set security zones security-zone lsys2-trust host-inbound-traffic protocols all
set security zones security-zone lsys2-trust interfaces reth3.0
set security zones security-zone lsys2-untrust host-inbound-traffic system-services all
set security zones security-zone lsys2-untrust host-inbound-traffic protocols all
set security zones security-zone lsys2-untrust interfaces lt-0/0/0.5
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust match application any
set security policies from-zone lsys2-trust to-zone lsys2-untrust policy
    lsys2trust-to-lsys2untrust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-trust policy
    lsys2untrust-to-lsys2trust then permit
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust match source-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust match destination-address any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust match application any
set security policies from-zone lsys2-untrust to-zone lsys2-untrust policy
    lsys2untrust-to-lsys2untrust then permit
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust match source-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust match destination-address any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust match application any
set security policies from-zone lsys2-trust to-zone lsys2-trust policy
    lsys2trust-to-lsys2trust then permit
```

Step-by-Step Procedure

NOTE: The user logical system administrator performs the following configuration while logged in to his or her user logical system. The master administrator can also configure a user logical system at the [edit logical-systems *logical-system*] hierarchy level.

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

To configure the LSYS1 user logical system:

1. Configure interfaces.

   ```
   [edit interfaces]
   lsys1-admin@host:LSYS1# set reth1 unit 0 family inet6 address 8885::1/64
   lsys1-admin@host:LSYS1# set reth2 unit 0 family inet6 address 7775::1/64
   ```

2. Configure routing.

   ```
   [edit routing-instances]
   lsys1-admin@host:LSYS1# set vr11 instance-type virtual-router
   lsys1-admin@host:LSYS1# set vr11 interface lt-0/0/0.3
   lsys1-admin@host:LSYS1# set vr11 interface reth1.0
   lsys1-admin@host:LSYS1# set vr11 routing-options rib vr11.inet6.0 static route
       6665::/64 next-hop 2111::5
   lsys1-admin@host:LSYS1# set vr11 routing-options rib vr11.inet6.0 static route
       9995::/64 next-hop 2111::1
   lsys1-admin@host:LSYS1# set vr12 instance-type virtual-router
   lsys1-admin@host:LSYS1# set vr12 interface reth2.0
   lsys1-admin@host:LSYS1# set vr12 routing-options interface-routes rib-group inet6
       vr11vr12v6
   lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
       8885::/64 next-table vr11.inet6.0
   lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
       9995::/64 next-table vr11.inet6.0
   lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
       6665::/64 next-table vr11.inet6.0
   lsys1-admin@host:LSYS1# set vr12 routing-options rib vr12.inet6.0 static route
       2111::/64 next-table vr11.inet6.0

   [edit routing-options]
   lsys1-admin@host:LSYS1# set rib-groups vr11vr12v6 import-rib vr11.inet6.0
   lsys1-admin@host:LSYS1# set rib-groups vr11vr12v6 import-rib vr12.inet6.0
   ```

3. Configure zones and security policies.

   ```
   [edit security zones]
   lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
       system-services all
   lsys1-admin@host:LSYS1# set security-zone lsys1-trust host-inbound-traffic
       protocols all
   lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces reth1.0
   lsys1-admin@host:LSYS1# set security-zone lsys1-trust interfaces lt-0/0/0.3
   ```

lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
    system-services all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust host-inbound-traffic
    protocols all
lsys1-admin@host:LSYS1# set security-zone lsys1-untrust interfaces reth2.0

[edit security policies from-zone lsys1-trust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match source-address
    any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match
    destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust match application
    any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match source-address
    any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match
    destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust match application
    any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1trust then permit

[edit security policies from-zone lsys1-untrust to-zone lsys1-untrust]
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
    source-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match
    destination-address any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust match application
    any
lsys1-admin@host:LSYS1# set policy lsys1untrust-to-lsys1untrust then permit

[edit security policies from-zone lsys1-trust to-zone lsys1-trust]
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match source-address
    any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match
    destination-address any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust match application any
lsys1-admin@host:LSYS1# set policy lsys1trust-to-lsys1trust then permit

Step-by-Step Procedure To configure the LSYS2 user logical system:

1. Configure interfaces.

    [edit interfaces]
    lsys2-admin@host:LSYS2# set reth3 unit 0 family inet6 address 6665::1/64

2. Configure routing.

    [edit routing-instances]
    lsys2-admin@host:LSYS2# set vr2 instance-type virtual-router
    lsys2-admin@host:LSYS2# set vr2 interface lt-0/0/0.5
    lsys2-admin@host:LSYS2# set vr2 interface reth3.0
    lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route
        7775::/64 next-hop 2111::3
    lsys2-admin@host:LSYS2# set vr2 routing-options rib vr2.inet6.0 static route
        8885::/64 next-hop 2111::3

lsys2-admin@host:LSYS2# **set vr2 routing-options rib vr2.inet6.0 static route 9995::/64 next-hop 2111::1**

3. Configure zones and security policies.

[edit security zones]
lsys2-admin@host:LSYS2# **set security-zone lsys2-trust host-inbound-traffic system-services all**
lsys2-admin@host:LSYS2# **set security-zone lsys2-trust host-inbound-traffic protocols all**
lsys2-admin@host:LSYS2# **set security-zone lsys2-trust interfaces reth3.0**
lsys2-admin@host:LSYS2# **set security zones security-zone lsys2-untrust host-inbound-traffic system-services all**
lsys2-admin@host:LSYS2# **set security-zone lsys2-untrust host-inbound-traffic protocols all**
lsys2-admin@host:LSYS2# **set security-zone lsys2-untrust interfaces lt-0/0/0.5**

[edit security policies from-zone lsys2-trust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust match source-address any**
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust match destination-address any**
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust match application any**
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2untrust then permit**

[edit security policies from-zone from-zone lsys2-untrust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust match source-address any**
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust match destination-address any**
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust match application any**
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2trust then permit**

[edit security policies from-zone lsys2-untrust to-zone lsys2-untrust]
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust match source-address any**
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust match destination-address any**
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust match application any**
lsys2-admin@host:LSYS2# **set policy lsys2untrust-to-lsys2untrust then permit**

[edit security policies from-zone lsys2-trust to-zone lsys2-trust]
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust match source-address any**
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust match destination-address any**
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust match application any**
lsys2-admin@host:LSYS2# **set policy lsys2trust-to-lsys2trust then permit**

Results   From configuration mode, confirm the configuration for LSYS1 by entering the **show interfaces**, **show routing-instances**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys1-admin@host:LSYS1# show interfaces
interfaces {
  lt-0/0/0 {
    unit 3 {
      encapsulation ethernet;
      peer-unit 2;
      family inet6 {
        address 2111::3/64;
      }
    }
  }
  reth1 {
    unit 0 {
      family inet6 {
        address 8885::1/64;
      }
    }
  }
  reth2 {
    unit 0 {
      family inet6 {
        address 7775::1/64;
      }
    }
  }
}
[edit]
lsys1-admin@host:LSYS1# show routing-instances
routing-instances {
  vr11 {
    instance-type virtual-router;
    interface lt-0/0/0.3;
    interface reth1.0;
    routing-options {
      rib vr11.inet6.0 {
        static {
          route 6665::/64 next-hop 2111::5;
          route 9995::/64 next-hop 2111::1;
        }
      }
    }
  }
  vr12 {
    instance-type virtual-router;
    interface reth2.0;
    routing-options {
      interface-routes {
        rib-group inet6 vr11vr12v6;
      }
      rib vr12.inet6.0 {
        static {
          route 8885::/64 next-table vr11.inet6.0;
          route 9995::/64 next-table vr11.inet6.0;
          route 6665::/64 next-table vr11.inet6.0;
          route 2111::/64 next-table vr11.inet6.0;
```

```
                }
              }
            }
          }
        }
[edit]
lsys1-admin@host:LSYS1# show routing-options
rib-groups {
  vr11vr12v6 {
    import-rib [ vr11.inet6.0 vr12.inet6.0 ];
  }
}
[edit]
lsys1-admin@host:LSYS1# show security
security {
  policies {
    from-zone lsys1-trust to-zone lsys1-untrust {
      policy lsys1trust-to-lsys1untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys1-untrust to-zone lsys1-trust {
      policy lsys1untrust-to-lsys1trust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys1-untrust to-zone lsys1-untrust {
      policy lsys1untrust-to-lsys1untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys1-trust to-zone lsys1-trust {
      policy lsys1trust-to-lsys1trust {
        match {
          source-address any;
```

```
                        destination-address any;
                        application any;
                      }
                      then {
                        permit;
                      }
                    }
                  }
                }
                zones {
                  security-zone lsys1-trust {
                    host-inbound-traffic {
                      system-services {
                        all;
                      }
                      protocols {
                        all;
                      }
                    }
                    interfaces {
                      reth1.0;
                      lt-0/0/0.3;
                    }
                  }
                  security-zone lsys1-untrust {
                    host-inbound-traffic {
                      system-services {
                        all;
                      }
                      protocols {
                        all;
                      }
                    }
                    interfaces {
                      reth2.0;
                    }
                  }
                }
              }
```

From configuration mode, confirm the configuration for LSYS2 by entering the **show interfaces**, **show routing-instances**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
lsys2-admin@host:LSYS2# show interfaces
interfaces {
  lt-0/0/0 {
    unit 5 {
      encapsulation ethernet;
      peer-unit 4;
      family inet6 {
        address 2111::5/64;
      }
    }
```

```
        }
        reth3 {
          unit 0 {
            family inet6 {
              address 6665::1/64;
            }
          }
        }
      }
    }
[edit]
lsys2-admin@host:LSYS2# show routing-instances
routing-instances {
  vr2 {
    instance-type virtual-router;
    interface lt-0/0/0.5;
    interface reth3.0;
    routing-options {
      rib vr2.inet6.0 {
        static {
          route 7775::/64 next-hop 2111::3;
          route 8885::/64 next-hop 2111::3;
          route 9995::/64 next-hop 2111::1;
        }
      }
    }
  }
}
[edit]
lsys2-admin@host:LSYS2# show security
security {
  policies {
    from-zone lsys2-trust to-zone lsys2-untrust {
      policy lsys2trust-to-lsys2untrust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys2-untrust to-zone lsys2-trust {
      policy lsys2untrust-to-lsys2trust {
        match {
          source-address any;
          destination-address any;
          application any;
        }
        then {
          permit;
        }
      }
    }
    from-zone lsys2-untrust to-zone lsys2-untrust {
```

```
                              policy lsys2untrust-to-lsys2untrust {
                                match {
                                  source-address any;
                                  destination-address any;
                                  application any;
                                }
                                then {
                                  permit;
                                }
                              }
                            }
                            from-zone lsys2-trust to-zone lsys2-trust {
                              policy lsys2trust-to-lsys2trust {
                                match {
                                  source-address any;
                                  destination-address any;
                                  application any;
                                }
                                then {
                                  permit;
                                }
                              }
                            }
                          }
                          zones {
                            security-zone lsys2-trust {
                              host-inbound-traffic {
                                system-services {
                                  all;
                                }
                                protocols {
                                  all;
                                }
                              }
                              interfaces {
                                reth3.0;
                              }
                            }
                            security-zone lsys2-untrust {
                              host-inbound-traffic {
                                system-services {
                                  all;
                                }
                                protocols {
                                  all;
                                }
                              }
                              interfaces {
                                lt-0/0/0.5;
                              }
                            }
                          }
                        }
```

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

Verification

Confirm that the configuration is working properly.

Verifying Chassis Cluster Status (IPv6)

Purpose   Verify the chassis cluster status, failover status, and redundancy group information.

Action   From operational mode, enter the **show chassis cluster status** command.

```
{primary:node0}
show chassis cluster status
Cluster ID: 1
Node                    Priority          Status     Preempt  Manual failover

Redundancy group: 0 , Failover count: 1
    node0                   200           primary        no       no
    node1                   100           secondary      no       no

Redundancy group: 1 , Failover count: 1
    node0                   200           primary        no       no
    node1                   100           secondary      no       no
```

Troubleshooting Chassis Cluster with Logs (IPv6)

Purpose   Use these logs to identify any chassis cluster issues. You should run these logs on both nodes.

Action   From operational mode, enter these **show log** commands.

```
user@host> show log jsrpd
user@host> show log chassisd
user@host> show log messages
user@host> show log dcd
user@host> show traceoptions
```

Verifying Logical System Licenses (IPv6)

Purpose   Verify information about logical system licenses.

Action   From operational mode, enter the **show system license status logical-system all** command.

```
{primary:node0}
user@host> show system license status logical-system all
```

```
node0:
------------------------------------------------------------------------
Logical system license status:

logical system name                license status
root-logical-system                enabled
LSYS0                              enabled
LSYS1                              enabled
LSYS2                              enabled
```

Verifying Logical System License Usage (IPv6)

Purpose   Verify information about logical system license usage.

> **NOTE:** The actual number of licenses used is only displayed on the primary node.

Action   From operational mode, enter the **show system license** command.

```
{primary:node0}
user@host> show system license
License usage:
                              Licenses    Licenses    Licenses    Expiry
  Feature name                   used   installed      needed
  logical-system                    4          25           0   permanent

Licenses installed:
  License identifier: JUNOS305013
  License version: 2
  Valid for device: JN110B54BAGB
  Features:
    logical-system-25 - Logical System Capacity
      permanent
```

Verifying Intra-Logical System Traffic on a Logical System (IPv6)

Purpose   Verify information about currently active security sessions within a logical system.

Action   From operational mode, enter the **show security flow session logical-system LSYS1** command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1
node0:
------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000115, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1784, Valid
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 22, Bytes: 1745
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 19, Bytes: 2108
Total sessions: 1

Flow Sessions on FPC2 PIC0:
```

```
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0


node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000006, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14392, Valid
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Verifying Intra-Logical System Traffic Within All Logical Systems (IPv6)

Purpose   Verify information about currently active security sessions on all logical systems.

Action   From operational mode, enter the **show security flow session logical-system all** command.

```
{primary:node0}
user@host> show security flow session logical-system all
node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000115, Policy name: lsys1trust-to-lsys1untrust/8, State: Active,
Timeout: 1776, Valid
Logical system: LSYS1
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 22, Bytes: 1745
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 19, Bytes: 2108
Total sessions: 1

Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:

Session ID: 10000006, Policy name: lsys1trust-to-lsys1untrust/8, State: Backup,
Timeout: 14384, Valid
Logical system: LSYS1
  In: 8885::2/34564 --> 7775::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 7775::2/23 --> 8885::2/34564;tcp, If: reth2.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

```
Flow Sessions on FPC2 PIC0:
Total sessions: 0

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Verifying Traffic Between User Logical Systems (IPv6)

Purpose   Verify information about currently active security sessions between logical systems.

Action   From operational mode, enter the **show security flow session logical-system** *logical-system-name* command.

```
{primary:node0}
user@host> show security flow session logical-system LSYS1

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000118, Policy name: lsys1trust-to-lsys1trust/11, State: Active,
Timeout: 1792, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 65, Bytes: 6701
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000010, Policy name: lsys1trust-to-lsys1trust/11, State: Backup,
Timeout: 14388, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system LSYS2

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
```

```
Session ID: 80000119, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1788, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 65, Bytes: 6701
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000011, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14380, Valid
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 1

Flow Sessions on FPC2 PIC1:
Total sessions: 0

{primary:node0}
user@host> show security flow session logical-system all

node0:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:

Session ID: 80000118, Policy name: lsys1trust-to-lsys1trust/11, State: Active,
Timeout: 1784, Valid
Logical system: LSYS1
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 65, Bytes: 6701

Session ID: 80000119, Policy name: lsys2untrust-to-lsys2trust/13, State: Active,
 Timeout: 1784, Valid
Logical system: LSYS2
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 91, Bytes: 6802
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 65, Bytes: 6701
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0

node1:
--------------------------------------------------------------------------

Flow Sessions on FPC0 PIC1:
Total sessions: 0

Flow Sessions on FPC2 PIC0:
```

```
Session ID: 80000010, Policy name: lsys1trust-to-lsys1trust/11, State: Backup,
Timeout: 14378, Valid
Logical system: LSYS1
  In: 8885::2/34565 --> 6665::2/23;tcp, If: reth1.0, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: lt-0/0/0.3, Pkts: 0, Bytes: 0

Session ID: 80000011, Policy name: lsys2untrust-to-lsys2trust/13, State: Backup,
 Timeout: 14376, Valid
Logical system: LSYS2
  In: 8885::2/34565 --> 6665::2/23;tcp, If: lt-0/0/0.5, Pkts: 0, Bytes: 0
  Out: 6665::2/23 --> 8885::2/34565;tcp, If: reth3.0, Pkts: 0, Bytes: 0
Total sessions: 2

Flow Sessions on FPC2 PIC1:
Total sessions: 0
```

Related Documentation

- Understanding Logical Systems in the Context of Chassis Cluster on page 167

- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 168

- *Example: Configuring an Active/Passive Chassis Cluster On a High-End SRX Series Services Gateway*

- *Chassis Cluster Overview*

PART 6

Configuring IPv6 for Logical Systems

Configuring IPv6 Addresses for Logical Systems

IPv6 Addresses in Logical Systems Overview

IP version 6 (IPv6) increases the size of an IP address from the 32 bits that compose an IPv4 address to 128 bits. Each extra bit given to an address doubles the size of its address space. IPv6 has a much larger address space than the soon-to-be exhausted IPv4 address space.

IPv6 addresses can be configured in logical systems for the following features:

- Interfaces
- Firewall authentication
- Flows
- Routing (BGP only)
- Zones and security policies
- Screen options
- Network Address Translation (except for interface NAT)
- Administrative operations such as Telnet, SSH, HTTPS, and other utilities
- Chassis clusters

> ⓘ NOTE: An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

Related Documentation
- *Understanding IPv6 Address Space, Addressing, Address Format, and Address Types*
- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 239
- Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (IPv6) (Master Administrators Only) on page 201
- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 238

Understanding IPv6 Dual-Stack Lite in Logical Systems

IPv6 dual-stack lite (DS-Lite) allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content. A DS-Lite softwire initiator at the customer edge encapsulates IPv4 packets into IPv6 packets while a softwire concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4 NAT translations.

A specific softwire concentrator and the set of softwire initiators that connect with that softwire concentrator can belong to only one logical system. The master administrator configures the maximum and reserved numbers of softwire initiators that can be connected to a softwire concentrator in a logical system using the **dslite-softwire-initiator** configuration statement at the [**edit system security-profile resources**] hierarchy level. The default maximum value is the system maximum; the default reserved value is 0.

> ⓘ NOTE: The master administrator can configure a security profile for the master logical system that specifies the maximum and reserved numbers of softwire initiators that can connect to a softwire concentrator configured for the master logical system. The number of softwire initiators configured in the master logical system count toward the maximum number of softwire initiators available on the device.

The user logical system administrator can configure softwire concentrators for their user logical system and the master administrator can configure softwire concentrators for the master logical system at the [**edit security softwires**] hierarchy level. The master administrator can also configure softwire concentrators for a user logical system at the [**edit logical-systems** *logical-system* **security softwires**] hierarchy level.

NOTE: The softwire concentrator IPv6 address can match an IPv6 address configured on either a physical interface or a loopback interface.

Related Documentation

- Example: Configuring IPv6 Dual-Stack Lite for a User Logical System on page 254
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51
- *Understanding IPv6 Dual-Stack Lite*

Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only)

This topic covers configuration of IPv6 interfaces, static routes, and routing instances for the master and interconnect logical systems. It also covers configuration of IPv6 logical tunnel interfaces for user logical systems.

- Requirements on page 239
- Overview on page 239
- Configuration on page 241
- Verification on page 247

Requirements

Before you begin:

- See "SRX Series Logical System Master Administrator Configuration Tasks Overview" on page 20 to understand how and where this procedure fits in the overall master administrator configuration process.
- See "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40.
- See "Understanding the Interconnect Logical System and Logical Tunnel Interfaces" on page 8.

Overview

This scenario shows how to configure interfaces for the logical systems on the device, including an interconnect logical system.

- For the interconnect logical system, the example configures logical tunnel interfaces lt-0/0/0.0, lt-0/0/0.2, and lt-0/0/0.4. The example configures a routing instance called vr and assigns the interfaces to it.

  Because the interconnect logical system acts as a virtual switch, it is configured as a VPLS routing instance type. The interconnect logical system's lt-0/0/0 interfaces are configured with ethernet-vpls as the encapsulation type. The corresponding peer

lt-0/0/0 interfaces in the master and user logical systems are configured with Ethernet as the encapsulation type.

- lt-0/0/0.0 connects to lt-0/0/0.1 on the root logical system.

- lt-0/0/0.2 connects to lt-0/0/0.3 on the LSYS1 logical system.

- lt-0/0/0.4 connects to lt-0/0/0.5 on the LSYS2 logical system.

- For the master logical system, called root-logical-system, the example configures ge-5/0/0 and assigns it to the vr0 routing instance. The example configures lt-0/0/0.1 to connect to lt-0/0/0.0 on the interconnect logical system and assigns it to the vr0 routing instance. The example configures static routes to allow for communication with other logical systems and assigns them to the vr0 routing instance.

- For the LSYS1 logical system, the example configures lt-0/0/0.3 to connect to lt-0/0/0.2 on the interconnect logical system.

- For the LSYS2 logical system, the example configures lt-0/0/0.5 to connect to lt-0/0/0.4 on the interconnect logical system.

shows the topology for this deployment including virtual routers and their interfaces for all IPv6 logical systems.

Figure 7: Configuring IPv6 Logical Tunnel Interfaces, Logical Interfaces, and Virtual Routers

Configuration

This topic explains how to configure interfaces for logical systems.

- Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System on page 242
- Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System on page 243
- Configuring Logical Tunnel Interfaces for the User Logical Systems on page 245

Configuring Logical Tunnel Interfaces and a Routing Instance for the Interconnect Logical System

CLI Quick Configuration

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

```
set forwarding-options family inet6 mode flow-based
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 0 peer-unit 1
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 2 peer-unit 3
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 encapsulation ethernet-vpls
set logical-systems LSYS0 interfaces lt-0/0/0 unit 4 peer-unit 5
set logical-systems LSYS0 routing-instances vr instance-type vpls
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.0
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.2
set logical-systems LSYS0 routing-instances vr interface lt-0/0/0.4
```

Step-by-Step Procedure

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

To configure the interconnect system lt-0/0/0 interfaces and routing instances:

1. Enable flow-based forwarding for IPv6 traffic.

    ```
    [edit security]
    user@host# set forwarding-options family inet6 mode flow-based
    ```

2. Configure the lt-0/0/0 interfaces.

    ```
    [edit logical-systems LSYS0 interfaces]
    user@host# set lt-0/0/0 unit 0 encapsulation ethernet-vpls
    user@host# set lt-0/0/0 unit 0 peer-unit 1
    user@host# set lt-0/0/0 unit 2 encapsulation ethernet-vpls
    user@host# set lt-0/0/0 unit 2 peer-unit 3
    user@host# set lt-0/0/0 unit 4 encapsulation ethernet-vpls
    user@host# set lt-0/0/0 unit 4 peer-unit 5
    ```

3. Configure the routing instance for the interconnect logical system and add its lt-0/0/0 interfaces to it.

    ```
    [edit logical-systems LSYS0 routing-instances]
    user@host# set vr instance-type vpls
    user@host# set vr interface lt-0/0/0.0
    user@host# set vr interface lt-0/0/0.2
    user@host# set vr interface lt-0/0/0.4
    ```

Results

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

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

```
user@host# show logical-systems LSYS0
interfaces {
  lt-0/0/0 {
    unit 0 {
      encapsulation ethernet-vpls;
      peer-unit 1;
    }
    unit 2 {
      encapsulation ethernet-vpls;
      peer-unit 3;
      }
      unit 4 {
        encapsulation ethernet-vpls;
        peer-unit 5;
      }
    }
  }
    routing-instances {
      vr {
        instance-type vpls;
        interface lt-0/0/0.0;
        interface lt-0/0/0.2;
        interface lt-0/0/0.4;
      }
    }
```

Configuring Interfaces, a Routing Instance, and Static Routes for the Master Logical System

CLI Quick
Configuration

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

```
set interfaces ge-5/0/0 vlan-tagging
set interfaces ge-5/0/0 unit 0 vlan-id 600
set interfaces lt-0/0/0 unit 1 encapsulation Ethernet
set interfaces lt-0/0/0 unit 1 peer-unit 0
set interfaces lt-0/0/0 unit 1 family inet address 1.1.1.1/24
set interfaces lt-0/0/0 unit 1 family inet6 address 1111::1/64
set interfaces ge-5/0/0 unit 0 family inet address 99.99.99.1/24
set interfaces ge-5/0/0 unit 0 family inet6 address 9999::1/64
set routing-instances vr0 instance-type virtual-router
set routing-instances vr0 interface lt-0/0/0.1
set routing-instances vr0 interface ge-5/0/0.0
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 7777::/64 next-hop
   1111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 8888::/64 next-hop
   1111::3
set routing-instances vr0 routing-options rib vr0.inet6.0 static route 6666::/64 next-hop
   1111::5
set routing-instances vr0 routing-options static route 77.77.77.0/24 next-hop 1.1.1.3
set routing-instances vr0 routing-options static route 88.88.88.0/24 next-hop 1.1.1.3
set routing-instances vr0 routing-options static route 66.66.66.0/24 next-hop 1.1.1.5
```

Step-by-Step Procedure

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

To configure the master logical system interfaces:

1. Configure the master (root) logical system and lt-0/0/0.1 interfaces.

   ```
   [edit interfaces]
   user@host# set ge-5/0/0 vlan-tagging
   user@host# set ge-5/0/0 unit 0 vlan-id 600
   user@host# set lt-0/0/0 unit 1 encapsulation Ethernet
   user@host# set lt-0/0/0 unit 1 peer-unit 0
   user@host# set lt-0/0/0 unit 1 family inet address 1.1.1.1/24
   user@host# set lt-0/0/0 unit 1 family inet6 address 1111::1/64
   user@host# set ge-5/0/0 unit 0 family inet address 99.99.99.1/24
   user@host# set ge-5/0/0 unit 0 family inet6 address 9999::1/64
   ```

2. Configure a routing instance for the master logical system, assign its interfaces to it, and configure static routes for it.

   ```
   [edit interfaces routing-instances]
   user@host# set vr0 instance-type virtual-router
   user@host# set vr0 interface lt-0/0/0.1
   user@host# set vr0 interface ge-5/0/0.0
   user@host# set vr0 routing-options rib vr0.inet6.0 static route 7777::/64 next-hop
       1111::3
   user@host# set vr0 routing-options rib vr0.inet6.0 static route 8888::/64 next-hop
       1111::3
   user@host# set vr0 routing-options rib vr0.inet6.0 static route 6666::/64 next-hop
       1111::5
   user@host# set vr0 routing-options static route 77.77.77.0/24 next-hop 1.1.1.3
   user@host# set vr0 routing-options static route 88.88.88.0/24 next-hop 1.1.1.3
   user@host# set vr0 routing-options static route 66.66.66.0/24 next-hop 1.1.1.5
   ```

Results

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

```
[edit]
user@host# show interfaces
ge-5/0/0 {
    vlan-tagging;
    unit 0 {
        vlan-id 600;
        family inet {
            address 99.99.99.1/24;
        }
        family inet 6{
            address 9999::1/64;
        }
    }
}
lt-0/0/0 {
    unit 1 {
        encapsulation ethernet;
```

```
            peer-unit 0;
            family inet {
               address 1.1.1.1/24;
            }
            family inet 6{
               address 1111::1/64;
            }
         }
      }

      [edit]
      user@host# show routing-instances
      vr0 {
         instance-type virtual-router;
         interface ge-5/0/0.0;
         interface lt-0/0/0;
         routing-options {
            rib vr0.inet6.0 {
            static {
               route 8888::/64 next-hop 1111::3;
               route 7777::/64 next-hop 1111::3;
               route 6666::/64 next-hop 1111::5;
            }
            }
            static {
               route 77.77.77.0/24 next-hop 1.1.1.3;
               route 88.88.88.0/24 next-hop 1.1.1.3;
               route 66.66.66.0/24 next-hop 1.1.1.5;
            }
         }
      }
```

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

Configuring Logical Tunnel Interfaces for the User Logical Systems

CLI Quick Configuration

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

```
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 encapsulation ethernet
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 peer-unit 2
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet address 1.1.1.3/24
set logical-systems LSYS1 interfaces lt-0/0/0 unit 3 family inet6 address 1111::3/64
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 encapsulation ethernet
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 peer-unit 4
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet address 1.1.1.5/24
set logical-systems LSYS2 interfaces lt-0/0/0 unit 5 family inet6 address 1111::5/64
```

Step-by-Step Procedure

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

1.   Configure the lt-0/0/0 interface for the first user logical system:

```
[edit logical-systems LSYS1 interfaces lt-0/0/0 unit 3]
user@host# set encapsulation ethernet
user@host# set peer-unit 2
user@host# set family inet address 1.1.1.3/24
user@host# set family inet6 address 1111::3/64
```

2. Configure the lt-0/0/0 interface for the second user logical system.

```
[edit logical-systems LSYS2 interfaces lt-0/0/0 unit 5]
user@host# set encapsulation ethernet
user@host# set peer-unit 4
user@host# set family inet address 1.1.1.5/24
user@host# set family inet6 address 1111::5/64
```

Results   From configuration mode, confirm your configuration by entering the **show logical-systems LSYS1 interfaces lt-0/0/0**, and **show logical-systems LSYS2 interfaces lt-0/0/0** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host#  show logical-systems LSYS1 interfaces lt-0/0/0

lt-0/0/0 {
    unit 3 {
        encapsulation ethernet;
        peer-unit 2;
        family inet {
            address 1.1.1.3/24;
        }
        family inet 6{
            address 1111::3/64;
        }
    }
}
```

```
user@host#  show logical-systems LSYS2 interfaces lt-0/0/0

lt-0/0/0 {
    unit 5 {
        encapsulation ethernet;
        peer-unit 4;
        family inet {
            address 1.1.1.5/24;
        }
        family inet 6{
            address 1111::5/64;
        }
    }
}
```

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

Verification

Verifying That the Static Routes Configured for the Master Administrator Are Correct

Purpose   Confirm that the configuration is working properly. Verify if you can send data from the master logical system to the other logical systems.

Action   From operational mode, use the **ping** command.

Related Documentation
- Understanding the Master Logical System and the Master Administrator Role on page 19
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25
- Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 8
- Example: Configuring IPv6 Zones for a User Logical System on page 247
- Example: Configuring IPv6 Security Policies for a User Logical System on page 250

Example: Configuring IPv6 Zones for a User Logical System

This example shows how to configure IPv6 zones for a user logical system.

- Requirements on page 247
- Overview on page 248
- Configuration on page 248

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator.

  See "User Logical System Configuration Overview" on page 23.

- Ensure that forwarding options for inet6 is flow-based. Otherwise, you must configure it and reset the device.

  Use the **show security forwarding-options** command to check the configuration.

  ....

  i NOTE: Only the user logical system administrator can configure the forwarding options.

  ....

Overview

This example configures the ls-product-design user logical system described in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40

This example creates the IPv6 zones and address books described in Table 18 on page 248.

Table 18: User Logical System Zone and Address Book Configuration

Feature	Name	Configuration Parameters
Zones	ls-product-design-trust	• Bind to interface ge-0/0/5.1. • TCP reset enabled.
	ls-product-design-untrust	• Bind to interface lt-0/0/0.3.
Address books	product-design-internal	• Address product-designers: 3002::1/96 • Attach to zone ls-product-design-trust
	product-design-external	• Address marketing: 3003::1/24 • Address accounting: 3004::1/24 • Address others: 3002::2/24 • Address set otherlsys: marketing, accounting • Attach to zone ls-product-design-untrust

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

```
set logical-system lsys1 security address-book product-design-internal address
    product-designers 3002::1/96
set logical-system lsys1 security address-book product-design-internal attach zone
    ls-product-design-trust
set logical-system lsys1 security address-book product-design-external address marketing
    3003::1/24
set logical-system lsys1 security address-book product-design-external address accounting
    3004::1/24
set logical-system lsys1 security address-book product-design-external address others
    3002::2/24
set logical-system lsys1 security address-book product-design-external address-set
    otherlsys address marketing
set logical-system lsys1 security address-book product-design-external address-set
    otherlsys address accounting
set logical-system lsys1 security address-book product-design-external attach zone
    ls-product-design-untrust
set logical-system lsys1 security zones security-zone ls-product-design-trust tcp-rst
set logical-system lsys1 security zones security-zone ls-product-design-trust interfaces
    ge-0/0/5.1
```

set logical-system lsys1 security zones security-zone ls-product-design-untrust interfaces
lt-0/0/0.3

Step-by-Step
Procedure

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

To configure IPv6 zones in a user logical system:

1. Log in to the user logical system as the logical system administrator and enter
    configuration mode.

    lsdesignadmin1@host:ls-product-design> configure
    lsdesignadmin1@host:ls-product-design#

2. Configure a security zone and assign it to an interface.

    [edit logical-system lsys1 security zones]
    lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-trust
        interfaces ge-0/0/5.1

3. Configure the TCP-Reset parameter for the zone.

    [edit logical-system lsys1 security zones security-zone ls-product-design-trust]
    lsdesignadmin1@host:ls-product-design# set tcp-rst

4. Configure a security zone and assign it to an interface.

    [edit logical-system lsys1 security zones]
    lsdesignadmin1@host:ls-product-design# set security-zone ls-product-design-untrust
        interfaces lt-0/0/0.3

5. Create global address book entries.

    [edit logical-system lsys1 security]
    lsdesignadmin1@host:ls-product-design# set address-book product-design-internal
        address product-designers 3002::1/96
    lsdesignadmin1@host:ls-product-design# set address-book product-design-external
        address marketing 3003::1/24
    lsdesignadmin1@host:ls-product-design# set address-book product-design-external
        address accounting 3004::1/24
    lsdesignadmin1@host:ls-product-design# set address-book product-design-external
        address others 3002::2/24
    lsdesignadmin1@host:ls-product-design# set address-book product-design-external
        address-set otherlsys address marketing
    lsdesignadmin1@host:ls-product-design# set address-book product-design-external
        address-set otherlsys address accounting

6. Attach address books to zones.

    [edit logical-system lsys1 security]
    lsdesignadmin1@host:ls-product-design# set address-book product-design-internal
        attach zone ls-product-design-trust
    lsdesignadmin1@host:ls-product-design# set address-book product-design-external
        attach zone ls-product-design-untrust

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

```
lsdesignadmin1@host:ls-product-design# show security zones
address-book {
    product-design-internal {
        address product-designers 3002::1/96;
        attach {
            zone ls-product-design-trust;
        }
    }
    product-design-external {
        address marketing 3003::1/24;
        address accounting 3004::1/24;
        address others 3002::2/24;
        address-set otherlsys {
            address marketing;
            address accounting;
        }
        attach {
            zone ls-product-design-untrust;
        }
    }
}
zones {
    security-zone ls-product-design-trust {
        tcp-rst;
        interfaces {
            ge-0/0/5.1;
        }
    }
    security-zone ls-product-design-untrust {
        interfaces {
            lt-0/0/0.3;
        }
    }
}
```

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

Related Documentation
- Understanding Logical System Zones on page 97
- User Logical System Configuration Overview on page 23
- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 239
- Example: Configuring IPv6 Security Policies for a User Logical System on page 250

Example: Configuring IPv6 Security Policies for a User Logical System

This example shows how to configure IPv6 security policies for a user logical system.

- Requirements on page 251
- Overview on page 251

Requirements

Before you begin:

- Log in to the user logical system as the logical system administrator.

  See "User Logical System Configuration Overview" on page 23.

- Use the **show system security-profiles policy** command to see the security policy resources allocated to the logical system.

- Configure zones and address books.

  See "Example: Configuring IPv6 Zones for a User Logical System" on page 247

Overview

This example shows how to configure the security policies described in Table 19 on page 251.

Table 19: User Logical System Security Policies Configuration

Policy Name	Configuration Parameters
permit-all-to-otherlsys	Permit the following traffic: - From zone: ls-product-design-trust - To zone: ls-product-design-untrust - Source address: product-designers - Destination address: otherlsys - Application: any
permit-all-from-otherlsys	Permit the following traffic: - From zone: ls-product-design-untrust - To zone: ls-product-design-trust - Source address: otherlsys - Destination address: product-designers - Application: any

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

```
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
    ls-product-design-untrust policy permit-all-to-otherlsys match source-address
    product-designers
```

```
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
    ls-product-design-untrust policy permit-all-to-otherlsys match destination-address
    otherlsys
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
    ls-product-design-untrust policy permit-all-to-otherlsys match application any
set logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
    ls-product-design-untrust policy permit-all-to-otherlsys then permit
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
    ls-product-design-trust policy permit-all-from-otherlsys match source-address otherlsys
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
    ls-product-design-trust policy permit-all-from-otherlsys match destination-address
    product-designers
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
    ls-product-design-trust policy permit-all-from-otherlsys match application any
set logical-systems lsys1 security policies from-zone ls-product-design-untrust to-zone
    ls-product-d esign-trust policy permit-all-from-otherlsys then permit
```

Step-by-Step Procedure

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

To configure IPv6 security policies for a user logical system:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

   ```
   lsdesignadmin1@host:ls-product-design> configure
   lsdesignadmin1@host:ls-product-design#
   ```

2. Configure a security policy that permits traffic from the ls-product-design-trust zone to the ls-product-design-untrust zone.

   ```
   [edit logical-systems lsys1 security policies from-zone ls-product-design-trust to-zone
       ls-product-design-untrust]
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
       source-address product-designers
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
       destination-address otherlsys
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys match
       application any
   lsdesignadmin1@host:ls-product-design# set policy permit-all-to-otherlsys then
       permit
   ```

3. Configure a security policy that permits traffic from the ls-product-design-untrust zone to the ls-product-design-trust zone.

   ```
   [edit logical-systems lsys1 security policies from-zone ls-product-design-untrust
       to-zone ls-product-design-trust]
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
       source-address otherlsys
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
       destination-address product-designers
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys match
       application any
   lsdesignadmin1@host:ls-product-design# set policy permit-all-from-otherlsys then
       permit
   ```

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

```
lsdesignadmin1@host:ls-product-design# show security policies
from-zone ls-product-design-trust to-zone ls-product-design-untrust {
    policy permit-all-to-otherlsys {
        match {
            source-address product-designers;
            destination-address otherlsys;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone ls-product-design-untrust to-zone ls-product-design-trust {
    policy permit-all-from-otherlsys {
        match {
            source-address otherlsys;
            destination-address product-designers;
            application any;
        }
        then {
            permit;
        }
    }
}
```

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

Verification

Verifying Policy Configuration

Purpose Verify information about policies and rules.

Action From operational mode, enter the **show security policies detail** command to display a summary of all policies configured on the logical system.

Related Documentation

- Understanding Logical System Security Policies on page 104

- User Logical System Configuration Overview on page 23

- *Troubleshooting Security Policies*

- Example: Configuring IPv6 Zones for a User Logical System on page 247

- Example: Configuring IPv6 for the Master, Interconnect, and User Logical Systems (Master Administrators Only) on page 239

Example: Configuring IPv6 Dual-Stack Lite for a User Logical System

This example shows how to configure a softwire concentrator for a user logical system.

Requirements

Before you begin:

- Log in to the user logical system as the user logical system administrator. See "User Logical System Configuration Overview" on page 23.

- Use the **show system security-profile dslite-softwire-initiator** command to see the number softwire initiators that can be connected to a softwire concentrator in the logical system.

Overview

This example shows how to configure a softwire concentrator to decapsulate IPv4-in-IPv6 packets in the ls-product-design user logical system shown in "Example: Creating User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40. The IPv6 address of the softwire concentrator is 3000::1 and the name of the softwire configuration is sc_1.

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

```
set security softwires softwire-name sc_1 softwire-concentrator 3000::1 softwire-type
    IPv4-in-IPv6
```

Step-by-Step Procedure

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

To configure an IPv6 DS-Lite softwire concentrator:

1. Log in to the user logical system as the logical system administrator and enter configuration mode.

    ```
    lsdesignadmin1@host:ls-product-design> configure
    lsdesignadmin1@host:ls-product-design#
    ```

2. Specify the address of the softwire concentrator and the softwire type.

    ```
    [edit security]
    ```

lsdesignadmin1@host:ls-product-design# **set softwires softwire-name sc_1 softwire-concentrator 3000::1 softwire-type IPv4-in-IPv6**

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

```
[edit]
lsdesignadmin1@host:ls-product-design# show security softwires
softwire-name sc_1 {
    softwire-concentrator 3000::1;
    softwire-type IPv4-in-IPv6;
}
```

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

Verification

Verifying the DS-Lite Configuration

Purpose   Verify that the softwire initiators can connect to the softwire concentrator configured in the user logical system.

Action   From operational mode, enter the **show security softwires** command.

If a softwire initiator is not connected, the operational output looks like this:

```
lsdesignadmin1@host:ls-product-design> show security softwires
Softwire Name           SC Address          Status    Number of SI connected
   sc_1                 3000::1             Active        0
```

If a softwire initiator is connected, the operational output looks like this:

```
lsdesignadmin1@host:ls-product-design> show security softwires
Softwire Name           SC Address          Status    Number of SI connected
   sc_1                 3000::1             Connected     1
```

Related Documentation

- Understanding IPv6 Dual-Stack Lite in Logical Systems on page 238
- User Logical System Configuration Overview on page 23

PART 7

Configuring System Resources Allocation

System Resources Allocation (Master Administrators Only)

Understanding CPU Allocation and Control

When device CPU utilization is low, logical systems can acquire and use CPU resources above their allocated reserve quotas as long as the system-wide utilization remains within a stable range. CPU utilization on a device should never reach 100 percent because a device running at 100 percent CPU utilization might be slow to respond to management or system events or be unable to handle traffic bursts.

CPU resources are used on a first-come first-served basis. Without controls, logical systems can compete for CPU resources and drive CPU utilization up to 100 percent. You cannot rely on the configuration of static resources, such as security policies and zones, to directly control CPU usage because a logical system with small numbers of static resources allocated could still consume a large amount of CPU. Instead, the master administrator can enable CPU resource control and configure CPU utilization parameters for logical systems.

> **NOTE:** Only the master administrator can enable CPU control and configure CPU utilization parameters. User logical system administrators can use the **show system security-profile cpu** command to view CPU utilization for their logical systems.

This topic includes the following sections:

CPU Control

The master administrator enables CPU control with the **cpu-control** configuration statement at the [**edit system security-profile resources**] hierarchy level.

> **NOTE:** The resources security profile is a special security profile that contains global settings that apply to all logical systems in the device. Other security profiles configured by the master administrator are bound to specific logical systems.

When CPU control is enabled, the master administrator can then configure the following CPU utilization parameters:

- A reserved CPU quota is the percentage of CPU utilization that is guaranteed for a logical system.
- The CPU control target is the upper limit, in percent, for system-wide CPU utilization on the device under normal operating conditions.

Reserved CPU Utilization Quota for Logical Systems

A configured reserved CPU quota guarantees that a specified percentage of CPU is always available to a logical system. During runtime, CPU utilization by each logical system is measured every two seconds. The reserved CPU quota is used to calculate the amount of CPU each logical system can use based on the runtime utilization.

The master administrator specifies the reserved CPU quota in a logical system security profile with the **cpu reserved** configuration statement at the [**edit system security-profile** *profile-name*] hierarchy level. The security profile is bound to one or more logical systems. Unlike other resources that are allocated to a logical system in a security profile, no maximum allowed quota can be configured for CPU utilization.

The Junos OS software checks to ensure that the sum of reserved CPU quotas for all logical systems on the device is less than 90 percent of the CPU control target value. If CPU control is enabled and reserved CPU quotas are not configured, the default reserved CPU quota for the master logical system is 1 percent and the default reserved CPU quota for user logical systems is 0 percent. The master administrator can configure reserved CPU quotas even if CPU control is not enabled. The master administrator can enable or disable CPU control without changing security profiles.

> **CAUTION:** The master logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota because traffic loss could occur.

CPU Control Target

CPU control target is the upper limit, in percent, for CPU utilization on the device under normal operating conditions. If CPU utilization on the device surpasses the configured target value, the Junos OS software initiates controls to bring CPU utilization between the target value and 90 percent of the target value. For example, if the CPU control target value is 80 and CPU utilization on the device surpasses 80 percent, then controls are initiated to bring CPU utilization within the range of 72 (90 percent of 80) and 80 percent.

During runtime, CPU utilization by each logical system is measured every two seconds. Dropping packets reduces the CPU usage for a logical system. If the CPU usage of a logical system exceeds its quota, CPU utilization control drops the packets received on that logical system. The packet drop rate is calculated every two seconds based on CPU utilization of all logical systems.

The master administrator configures the CPU control target with the **cpu-control-target** configuration statement at the [**edit system security-profile resources**] hierarchy level. A stable level of CPU utilization should be relatively close to 100 percent but allow for bursts in CPU utilization. The master administrator should configure the CPU control target level based on an understanding of the usage pattern of the logical system's deployment on the device.

CPU control must be enabled for the Junos OS software to control CPU usage. If the master administrator enables CPU control without specifying a CPU control target value, the default CPU control target is 80 percent.

Shared CPU Resources and CPU Quotas

The sum of the reserved CPU quotas for all logical systems on the device must be less than 90 percent of the CPU control target; the difference is called the shared CPU resource. The shared CPU resource is dynamically allocated among the logical systems that need additional CPU. This means that a logical system can use more CPU than its reserved CPU quota.

The CPU quota for a logical system is the sum of its reserved CPU quota and its portion of the shared CPU resource. If multiple logical systems need more CPU resources, they split the shared CPU resource based on the relative weights of their reserved CPU quotas. Logical systems with larger reserved CPU quotas receive larger portions of the shared CPU resource. The goal for CPU control is to keep the actual CPU utilization of a logical system at its CPU quota. If a logical system's CPU needs are greater than its CPU quota, packets are dropped for that logical system.

The following scenarios illustrate CPU control for logical systems. In each scenario, the CPU control target value is 80, which means that CPU controls will keep the maximum system-wide CPU utilization between 72 and 80 percent. The reserved CPU quotas for the logical systems are configured as follows: master and lsys1 logical systems are 10 percent each and the lsys2 logical system is 5 percent.

CPU Utilization Scenario 1

In this scenario, each of the three logical systems needs 40 percent of CPU.
Table 20 on page 262 shows the CPU quotas for each logical system. Because the CPU
needed by each logical system is greater than its CPU quota, packets are dropped for
each logical system.

Table 20: CPU Utilization Scenario 1

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
master	40%	28.8%	Yes
lsys1	40%	28.8%	Yes
lsys2	40%	14.4%	Yes

CPU Utilization Scenario 2

In this scenario, the master logical system needs 25 percent of CPU while the two user
logical systems need 40 percent. Table 21 on page 262 shows the CPU quota for the master
logical system is equal to the CPU it needs, so no packets are dropped for the master
logical system and CPU control monitors the CPU utilization of the master logical system.
Packets are dropped for lsys1 and lsys2.

Table 21: CPU Utilization Scenario 2

Logical System	Needed CPU	CPU Quotas	Packets Dropped?
master	25%	25%	No
lsys1	40%	31.3%	Yes
lsys2	40%	15.6%	Yes

CPU Utilization Scenario 3

In this scenario, the master and lsys2 logical systems need 5 percent and 3 percent of
CPU, respectively, while lsys1 needs 40 percent. Table 22 on page 262 shows system-wide
CPU utilization is 48 percent, which is less than 72 percent (90 percent of the CPU control
target), so no packets are dropped and CPU control monitors all logical systems.

Table 22: CPU Utilization Scenario 3

Logical System	Needed CPU	CPU Quota	Packets Dropped?
master	5%	5%	No
lsys1	40%	40%	No
lsys2	3%	3%	No

Monitoring CPU Utilization

CPU utilization can be monitored by either the master administrator or the user logical system administrators. The master administrator can monitor CPU utilization for the master logical system, a specified user logical system, or all logical systems. User logical system administrators can only monitor CPU utilization for their logical system.

The **show system security-profile cpu** command shows the usage and drop rate in addition to the reserved CPU quota configured for the logical system. During runtime, CPU utilization by each logical system is measured every two seconds. The usage and drop rates displayed are the values at the interval prior to when the **show** command is run. If the **detail** option is not specified, the utilization of the central point (CP) and the average utilization of all services processing units (SPUs) is shown. The **detail** option displays the CPU utilization on each SPU.

The CPU utilization log file **lsys-cpu-utilization-log** contains utilization data for all logical systems on the device. Only the master administrator can view the log file with the **show log lsys-cpu-utilization-log** command.

Related Documentation

- Example: Configuring CPU Utilization (Master Administrators Only) on page 263
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

Example: Configuring CPU Utilization (Master Administrators Only)

The master administrator can enable CPU control and configure CPU utilization parameters. This example shows how to enable CPU utilization control and configure CPU utilization quotas and a control target.

- Requirements on page 263
- Overview on page 263
- Configuration on page 264
- Verification on page 265

Requirements

Before you begin:

- Log in to the master logical system as the master administrator. See "Understanding the Master Logical System and the Master Administrator Role" on page 19.
- Bind security profiles to the master logical system and user logical systems configured on the device. See "Example: Configuring Logical Systems Security Profiles (Master Administrators Only)" on page 56.

Overview

In this example, you enable CPU control and set the CPU control target to be 85 percent. You allocate reserved CPU quotas to the logical systems shown in "Example: Creating

User Logical Systems, Their Administrators, Their Users, and an Interconnect Logical System (Master Administrators Only)" on page 40. The logical systems are bound to the security profiles shown in Table 23 on page 264 and are assigned the reserved CPU quotas in the security profiles.

Table 23: Logical Systems, Security Profiles, and Reserved CPU Quotas

Logical System	Security Profile	Reserved CPU Quotas
root-logical-system (master)	master-profile	2 percent
ls-product-design	ls-design-profile	2 percent
ls-marketing-dept, ls-accounting-dept	ls-accnt-mrkt-profile	1 percent

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

```
set system security-profile resources cpu-control
set system security-profile resources cpu-control-target 85
set system security-profile master-profile cpu reserved 2
set system security-profile ls-design-profile cpu reserved 2
set system security-profile ls-accnt-mrkt-profile cpu reserved 1
```

Step-by-Step Procedure

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

To configure CPU utilization control parameters:

1. Log in to the master logical system as the master administrator and enter configuration mode.

   ```
   [edit]
   admin@host> configure
   admin@host#
   ```

2. Enable CPU control.

   ```
   [edit system security-profile resources]
   admin@host# set cpu-control
   ```

3. Configure the CPU control target.

   ```
   [edit system security-profile resources]
   admin@host# set cpu-control-target 85
   ```

4. Configure the reserved CPU quotas in the security profiles.

   ```
   [edit system]
   admin@host# set security-profile security-profile master-profile cpu reserved 2
   admin@host# set security-profile security-profile ls-design-profile cpu reserved 2
   ```

```
admin@host# set security-profile security-profile ls-accnt-mrkt-profile cpu reserved
    1
```

Results   From configuration mode, confirm your configuration by entering the **show system security-profile** command. If the output does not display the intended configuration, repeat the \ instructions in this example to correct the configuration.

For brevity, this **show** command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
admin@host# show system security-profile
  resources {
    cpu-control;
    cpu-control-target 85;
  }
    ls-accnt-mrkt-profile {
      ...
      cpu {
        reserved 1;
      }
      logical-system [ ls-marketing-dept ls-accounting-dept ];
    }
    ls-design-profile {
      ...
      cpu {
        reserved 2;
      }
      logical-system ls-product-design;
    }
    master-profile {
      ...
      cpu {
        reserved 2;
      }
      logical-system root-logical-system;
    }
```

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

Verification

Confirm that the configuration is working properly.

- Verifying CPU Utilization on page 265

Verifying CPU Utilization

Purpose   Display the configured reserved CPU quota, the actual CPU usage, and the drop rate.

Action   From operational mode, enter the **show system security-profile cpu logical-system all** command.

```
admin@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 85.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)
drop rate(%)
root-logical-system    master-profile CP               0.10%        2.00%
       0.00%
root-logical-system    master-Profile SPU              0.25%        2.00%
       0.00%
ls-product-design      ls-design-profile CP            0.53%        2.00%
       0.00%
ls-product-design      ls-design-profile SPU           0.26%        2.00%
       0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP         0.10%        1.00%
       0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU        0.15%        1.00%
       0.00%
ls-accounting-dept     ls-acct-mrkt-profile CP         0.23%        1.00%
       0.00%
ls-accounting-dept     ls-acct-mrkt-profile SPU        0.34%        1.00%
       0.00%
```

Related Documentation

- Understanding CPU Allocation and Control on page 259

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

Example: Deleting an SRX Series Services Gateway Logical System (Master Administrators Only)

This example shows how to delete a logical system configured for an SRX Series Services Gateway device running logical systems. Only the master administrator can delete a logical system.

- Requirements on page 266

- Overview on page 266

- Configuration on page 267

- Verification on page 269

Requirements

The example uses an SRX5600 device running Junos OS with Logical Systems.

Alternatively, follow those instructions substituting your own configuration values.

Overview

This example shows how to delete a logical system, which you can do at any time. However, if you have configured the device to include the maximum number of logical systems that are supported you must first delete an existing logical system before you can add another one.

Deletion of a logical system is a simple procedure that includes these tasks:

- Remove from the logical system the security profile that is bound to it.

  Note that in this step you are not deleting the security profile—it might be used for other logical systems—but simply detaching it from the logical system that you intend to delete.

- Detach from the logical system any login classes that are associated with it.

  Removing them from the logical system does not delete the login classes.

- Delete the logical system.

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

```
delete system security-profile ls-design-profile logical-system ls-product-design
delete system login class ls-design-admin logical-system ls-product-design
delete system login class ls-design-user logical-system ls-product-design
delete logical-system ls-product-design
```

Step-by-Step Procedure

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

To delete a logical system:

1. Determine that the logical system that you want to delete exists.

   ```
   [edit]
   user@host# show logical-systems ?
   interconnect-logical-system Logical system name
   ls-accounting-dept Logical system name
   ls-marketing-dept Logical system name
   ls-product-design Logical system name
   ```

2. Delete the security profile.

   a. Verify that security profile that you intend to detach from the logical system is bound to it.

      ```
      [edit]
      user@host# show system security-profile ls-design-profile
      logical-system [ ls-product-design ];
      ```

   b. Detach the security profile from the logical system.

      ```
      [edit]
      ```

```
user@host# delete system security-profile ls-design-profile logical-system
    ls-product-design
```

4. Delete the login classes.

   a. Display the login class and login user configurations for the user logical system
      administrator.

   ```
   user@host> show configuration system login class ls-design-admin
   logical-system ls-product-design;
   permissions all;
   user@host> show configuration system login user lsdesignadmin1
   full-name lsdesignadmin1;
   uid 2006;
   class ls-design-admin;
   authentication {
      encrypted-password "$ABC123"; ## SECRET-DATA
      }
   ```

   b. Detach the login class for the administrator from the logical system.

   ```
   [edit]
   user@host# delete system login class ls-design-admin logical-system
       ls-product-design
   ```

   c. Display the login class and login user configurations for the user.

   ```
   user@host> show configuration system login class ls-design-user
   logical-system ls-product-design;
   permissions view;
   user@host> show configuration system login user lsdesignuser1
   full-name lsdesignuser1
   uid 2007;
   class ls-design-user;
   authentication {
      encrypted-password "$ABC123"; ## SECRET-DATA
      }
   ```

   d. Detach the login class for the user from the logical system.

   ```
   user@host# delete system login class ls-design-user logical-system
       ls-product-design
   ```

5. Delete the logical system.

   ```
   [edit]
   user@host# delete logical-system ls-product-design
   ```

Results From configuration mode, confirm your configuration by entering the **show logical-systems**
command. In this case, the logical system that you deleted should not be included in
displayed list of logical systems configured for the device. If the output does not display
the intended configuration, repeat the configuration instructions in this example to correct
it.

```
user@host# show logical-systems
interconnect-logical-system Logical system name
ls-accounting-dept Logical system name
```

interconnect-logical-system Logical system name
ls-marketing-dept Logical system name

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

Verification

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

- Verifying That the Correct Logical System and Its Profile and Attached Class Were Deleted on page 269

Verifying That the Correct Logical System and Its Profile and Attached Class Were Deleted

Purpose   Verify if the logical system has been deleted using the show command described previously.

Related Documentation
- Understanding User Logical Systems and the User Logical System Administrator Role on page 25
- Understanding Logical Systems for SRX Series Services Gateways on page 3

PART 8

Troubleshooting

Troubleshooting Logical Systems (Master Administrators Only)

Understanding Security Logs and Logical Systems

Security logs are system log messages that include security events. If a device is configured for logical systems, security logs generated within the context of a logical system use the name *logname*_**LS** (for example, **IDP_ATTACK_LOG_EVENT_LS**). The logical system version of a log has the same set of attributes as the log for devices that are not configured for logical systems, but it also includes logical-system-name as the first attribute.

The following security log shows the attributes for the IDP_ATTACK_LOG_EVENT log for a device that is *not* configured for logical systems:

```
IDP_ATTACK_LOG_EVENT {
    help "IDP attack log";
    description "IDP Attack log generated for attack";
    type event;
    args timestamp message-type source-address source-port destination-address
 destination-port protocol-name service-name application-name rule-name
rulebase-name policy-name repeat-count action threat-severity attack-name
nat-source-address nat-source-port nat-destination-address nat-destination-port
 elapsed-time inbound-bytes outbound-bytes inbound-packets outbound-packets
source-zone-name source-interface-name destination-zone-name
destination-interface-name packet-log-id message;
    severity LOG_INFO;
    flag auditable;
    edit "2010/10/01 mvr created";
}
```

The following security log shows the attributes for the IDP_ATTACK_LOG_EVENT_LS log for a device that is configured for logical systems (note that logical-system-name is the first attribute):

```
IDP_ATTACK_LOG_EVENT_LS {
    help "IDP attack log";
    description "IDP Attack log generated for attack";
    type event;
    args logical-system-name timestamp message-type source-address source-port
 destination-address destination-port protocol-name service-name
application-name rule-name rulebase-name policy-name repeat-count action
threat-severity attack-name nat-source-address nat-source-port
nat-destination-address nat-destination-port elapsed-time inbound-bytes
outbound-bytes inbound-packets outbound-packets source-zone-name
source-interface-name destination-zone-name destination-interface-name
packet-log-id message;
    severity LOG_INFO;
    flag auditable;
    edit "2010/10/01 mvr created";
}
```

If a device is configured for logical systems, log parsing scripts might need to be modified because the log name includes the **_LS** suffix and the logical-system-name attribute can be used to segregate logs by logical system.

If a device is not configured for logical systems, the security logs remain unchanged and scripts built to parse logs do not need any modification.

NOTE: Only the master administrator can configure logging at the [edit security log] hierarchy level. User logical system administrators cannot configure logging for their logical systems.

Understanding Data Path Debugging for Logical Systems

Data path debugging provides tracing and debugging at multiple processing units along the packet-processing path. Data path debugging can also be performed on traffic between logical systems.

NOTE: Only the master administrator can configure data path debugging for logical systems at the [edit security datapath-debug] level. User logical system administrators cannot configure data path debugging for their logical systems.

End-to-end event tracing traces the path of a packet from when it enters the device to when it leaves the device. When the master administrator configures end-to-end event tracing, the trace output contains logical system information.

The master administrator can also configure tracing for traffic between logical systems. The trace output shows traffic entering and leaving the logical tunnel between logical systems. When the **preserve-trace-order** option is configured, the trace message is sorted chronologically. In addition to the trace action, other actions such as packet-dump and packet-summary may be configured for traffic between logical systems.

Related Documentation
- Performing Tracing for Logical Systems (Master Administrators Only) on page 275

Performing Tracing for Logical Systems (Master Administrators Only)

> **NOTE:** Only the master administrator can configure data path debugging for logical systems at the root level.

To configure an action profile for a trace or packet capture:

1. Specify event types and trace actions. You can specify any combination of event types and trace actions. For example, the following statements configure multiple trace actions for each event type:

   ```
   [edit security datapath-debug]
   user@host# set action-profile p1 event lbt trace
   user@host# set action-profile p1 event lbt count
   user@host# set action-profile p1 event lbt packet-summary
   user@host# set action-profile p1 event lbt packet-dump
   user@host# set action-profile p1 event pot trace
   user@host# set action-profile p1 event pot count
   user@host# set action-profile p1 event pot packet-summary
   user@host# set action-profile p1 event pot packet-dump
   user@host# set action-profile p1 event np-ingress trace
   user@host# set action-profile p1 event np-ingress count
   user@host# set action-profile p1 event np-ingress packet-summary
   user@host# set action-profile p1 event np-ingress packet-dump
   user@host# set action-profile p1 event np-egress trace
   user@host# set action-profile p1 event np-egress count
   user@host# set action-profile p1 event np-egress packet-summary
   user@host# set action-profile p1 event np-egress packet-dump
   user@host# set action-profile p1 event jexec trace
   user@host# set action-profile p1 event jexec count
   user@host# set action-profile p1 event jexec packet-summary
   user@host# set action-profile p1 event jexec packet-dump
   user@host# set action-profile p1 event lt-enter trace
   user@host# set action-profile p1 event lt-enter count
   user@host# set action-profile p1 event lt-enter packet-summary
   user@host# set action-profile p1 event lt-enter packet-dump
   user@host# set action-profile p1 event lt-leave trace
   user@host# set action-profile p1 event lt-leave count
   user@host# set action-profile p1 event lt-leave packet-summary
   user@host# set action-profile p1 event lt-leave packet-dump
   ```

2. Specify action profile options.

   ```
   [edit security datapath-debug]
   user@host# set action-profile p1 record-pic-history
   user@host# set action-profile p1 preserve-trace-order
   ```

3. Configure packet filter options.

   ```
   [edit security datapath-debug]
   user@host# set packet-filter 1 action-profile p1
   user@host# set packet-filter 1 protocol udp
   ```

To capture trace messages for logical systems:

1. Configure the trace capture file.

   [edit security datapath-debug]
   user@host# **set traceoptions file e2e.trace**
   user@host# **set traceoptions file size 10m**

2. Display the captured trace in operational mode.

   user@host> **show log e2e.trace**
   ```
     Jul  7 09:49:56
     09:49:56.417578:CID-00:FPC-01:PIC-00:THREAD_ID-00:FINDEX:0:IIF:75:SEQ:0:TC:0
     PIC History: ->C0/F1/P0
     NP ingress channel 0 packet
     Meta: Src: F1/P0 Dst: F0/P0
     IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

     Jul  7 09:49:56
     09:49:55.1414031:CID-00:FPC-00:PIC-00:THREAD_ID-04:FINDEX:0:IIF:75:SEQ:0:TC:1
     PIC History: ->C0/F1/P0->C0/F0/P0
     LBT pkt, payload: DATA
     Meta: Src: F1/P0 Dst: F0/P0
     IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500


     ...
     (Some trace information omitted)
     ...

     .Jul  7 09:49:56
     09:49:55.1415649:CID-00:FPC-00:PIC-00:THREAD_ID-05:FINDEX:0:IIF:75:SEQ:0:TC:16
     PIC History: ->C0/F1/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0
     POT pkt, action: POT_SEND payload: DATA
     Meta: Src: F0/P0 Dst: F1/P0
     IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500

     Jul  7 09:49:56
     09:49:56.419274:CID-00:FPC-01:PIC-00:THREAD_ID-00:FINDEX:0:IIF:75:SEQ:0:TC:17
     PIC History: ->C0/F1/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F0/P0->C0/F1/P0
     NP egress channel 0 packet
     Meta: Src: F0/P0 Dst: F1/P0
     IP: saddr 10.1.1.2 daddr 30.1.1.2 proto 6 len 500
   ```

3. Clear the log.

   user@host> **clear log e2e.trace**

To perform packet capture for logical systems:

1. Configure the packet capture file.

   [edit security datapath-debug]
   user@host# **set capture-file e2e.pcap**
   user@host# **set capture-file format pcap**
   user@host# **set capture-file size 10m**
   user@host# **set capture-file world-readable**
   user@host# **set capture-file maximum-capture-size 1500**

2. Enter operational mode to start and then stop the packet capture.

   user@host> **request security datapath-debug capture start**
   user@host> **request security datapath-debug capture stop**

i NOTE: Packet capture files can be opened and analyzed offline with tcpdump or any packet analyzer that recognizes the libpcap format. You can also use FTP or the Session Control Protocol (SCP) to transfer the packet capture files to an external device.

3. Disable packet capture from configuration mode.

i NOTE: Disable packet capture before opening the file for analysis or transferring the file to an external device with FTP or SCP. Disabling packet capture ensures that the internal file buffer is flushed and all the captured packets are written to the file.

[edit forwarding-options]
user@host# **set packet-capture disable**

4. Display the packet capture.

- To display the packet capture with the tcpdump utility:

```
user@host# tcpdump -nr /var/log/e2e.pcap
 09:49:55.1413990 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414154 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415062 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415184 C0/F0/P0 event:11(lbt) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414093 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414638 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415011 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415129 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415511 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415649 C0/F0/P0 event:12(pot) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415249 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1415558 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414226 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414696 C0/F0/P0 event:18(jexec) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414828 C0/F0/P0 event:16(lt-enter) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:55.1414919 C0/F0/P0 event:15(lt-leave) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
 09:49:56.417560 C0/F1/P0 event:1(np-ingress) SEQ:0 IP 10.1.1.2.23451 >
 30.1.1.2.12345: S 0:460(460) win 0
```

```
                              09:49:56.419263 C0/F1/P0 event:2(np-egress) SEQ:0 IP 10.1.1.2.23451 >
                              30.1.1.2.12345: S 0:460(460) win 0
```

- To display the packet capture from CLI operational mode:

  user@host> **show security datapath-debug capture**
```
  Packet 1, len 568: (C0/F0/P0/SEQ:0:lbt)
  00 00 00 00 00 00 50 c5 8d 0c 99 4a 00 00 0a 01
  01 02 08 00 45 60 01 f4 00 00 00 00 40 06 4e 9f
  0a 01 01 02 1e 01 01 02 5b 9b 30 39 00 00 00 00
  00 00 00 00 50 02 00 00 f8 3c 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 ac 7a 00 04
  00 00 00 00 b3 e3 15 4e 66 93 15 00 04 22 38 02
  38 02 00 00 00 01 00 03 0b 00 00 00 50 d0 1a 08
  30 de be bf e4 f3 19 08
  Packet 2, len 624: (C0/F0/P0/SEQ:0:lbt)
  aa 35 00 00 00 00 00 00 00 00 00 00 00 03 00 00
  00 0a 00 00 00 00 00 00 05 bd 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 50 c5
  8d 0c 99 4a 00 00 0a 01 01 02 08 00 45 60 01 f4
  00 00 00 00 40 06 4e 9f 0a 01 01 02 ac 7a 00 04
  00 00 00 00 b3 e3 15 4e 0a 94 15 00 04 5a 70 02
  70 02 00 00 00 03 00 03 0b 00 00 00 50 d0 1a 08
  30 de be bf e4 f3 19 08

  ...
  (Packets 3 through 17 omitted)
  ...

  Packet 18, len 568: (C0/F1/P0/SEQ:0:np-egress)
  00 00 00 04 00 00 00 00 1e 01 01 02 50 c5 8d 0c
  99 4b 08 00 45 60 01 f4 00 00 00 00 3e 06 50 9f
  0a 01 01 02 1e 01 01 02 5b 9b 30 39 00 00 00 00
  00 00 00 00 50 02 00 00 f8 3c 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  00 00 00 00 00 00 00 00 00 00 00 00 ac 7a 04 00
  00 00 00 00 b4 e3 15 4e bf 65 06 00 04 22 38 02
  38 02 00 00 00 11 00 03 02 00 00 00 50 d0 1a 08
  30 de be bf e4 f3 19 08
```

  user@host> **show security datapath-debug counters**
```
  Datapath debug counters
  Packet Filter 1:
  lt-enter
  Chassis 0 FPC 0 PIC 1: 0
  lt-enter
  Chassis 0 FPC 0 PIC 0: 1
  lt-leave
  Chassis 0 FPC 0 PIC 1: 0
  lt-leave
  Chassis 0 FPC 0 PIC 0: 1
  np-egress
  Chassis 0 FPC 1 PIC 3: 0
  np-egress
  Chassis 0 FPC 1 PIC 1: 0
  np-egress
  Chassis 0 FPC 1 PIC 2: 0
  np-egress
  Chassis 0 FPC 1 PIC 0: 1
  pot
```

```
Chassis 0 FPC 0 PIC 1: 0
pot
Chassis 0 FPC 0 PIC 0: 6
np-ingress
Chassis 0 FPC 1 PIC 3: 0
np-ingress
Chassis 0 FPC 1 PIC 1: 0
np-ingress
Chassis 0 FPC 1 PIC 2: 0
np-ingress
Chassis 0 FPC 1 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 4
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 4
```

Related
Documentation

- Understanding Data Path Debugging for Logical Systems on page 274

Troubleshooting DNS Name Resolution in Logical System Security Policies (Master Administrators Only)

Problem   **Description:** The address of a hostname in an address book entry that is used in a security policy might fail to resolve correctly.

Cause   Normally, address book entries that contain dynamic hostnames refresh automatically for SRX Series devices. The TTL field associated with a DNS entry indicates the time after which the entry should be refreshed in the policy cache. Once the TTL value expires, the SRX Series device automatically refreshes the DNS entry for an address book entry.

However, if the SRX Series device is unable to obtain a response from the DNS server (for example, the DNS request or response packet is lost in the network or the DNS server cannot send a response), the address of a hostname in an address book entry might fail to resolve correctly. This can cause traffic to drop as no security policy or session match is found.

Solution   The master administrator can use the **show security dns-cache** command to display DNS cache information on the SRX Series device. If the DNS cache information needs to be refreshed, the master administrator can use the **clear security dns-cache** command.

> *i* NOTE: These commands are only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Related Documentation

- Understanding Logical System Security Policies

Configuration Statements and Operational Commands

Configuration Statements

address-book

Syntax
```
address-book (book-name | global) {
    address address-name {
        ip-prefix {
            description text;
        }
        description text;
        dns-name domain-name {
            ipv4-only;
            ipv6-only;
        }
        range-address lower-limit to upper-limit;
        wildcard-address ipv4-address/wildcard-mask;
    }
    address-set address-set-name {
        address address-name;
        address-set address-set-name;
        description text;
    }
    attach {
        zone zone-name;
    }
    description text;
}
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OSRelease 11.1. Statement moved under the security hierarchy in Junos OS Release 11.2. Support for address range added in Junos OS Release 12.1. The **description** option added in Junos OS Release 12.1.

Description   Define entries in the address book. Address book entries can include any combination of IPv4 addresses, IPv6 addresses, DNS names, wildcard addresses, and address range. You define addresses and address sets in an address book and then use them when configuring different features, such as security policies and NAT.

> (i) NOTE: IPv6 wildcard address configuration is not supported in this release.

Options
- **address-book** *book-name*—Name of the address book.

- **global**—An address book that is available by default. You can add any combination of IPv4 addresses, IPv6 addresses, wildcard addresses, DNS names, or address range to the global address book. You do not need to attach the global address book to a security zone; entries in the global address book are available to all security zones that are not attached to address books.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *Understanding Address Books*
- *Understanding Address Sets*

address-book (System)

Syntax
```
address-book {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level
[edit system security-profile *security-profile-name*]

Release Information
Statement introduced in Junos OS Release 11.2.

Description
Specify the number of address books that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- *Understanding Address Books*

appfw-profile (System)

Syntax	appfw-profile {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *profile-name*]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the application firewall profile quota of a logical system.
Options	• **maximum** *amount*—Specify the maximum allowed quota value. **Range**: 0 through 1024 • **reserved** *amount*—Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	• *Application Firewall Overview*

appfw-rule

Syntax	appfw-rule {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the number of application firewall rule configurations that a master administrator can configure for a master logical system or user logical system when the security profile is bound to the logical systems.

The master administrator:

- Uses security profiles to provision logical systems with resources

- Binds security profiles to the master logical system and the user logical systems

- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- *Application Firewall Overview*

appfw-rule-set

Syntax	appfw-rule-set {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the number of application firewall rule set configurations that a master administrator can configure for a master logical system or user logical system when the security profile is bound to the logical systems.

The master administrator:

- Uses security profiles to provision logical systems with resources

- Binds security profiles to the master logical system and the user logical systems

- Can configure more than one security profile, allocating different numbers of resources in various profiles

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can use resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- *Application Firewall Overview*

application-firewall

Syntax

```
application-firewall {
    rule-sets rule-set-name {
        default-rule {
            (deny | permit);
        }
        rule rule-name {
            match {
                dynamic-application [system-application];
                dynamic-application-group [system-application-group];
            }
            then {
                (deny | permit);
            }
        }
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
    }
}
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 11.1.

Description   Configure application firewall rule sets with rules defining match criteria and the action to be performed.

Options   The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *Application Firewall Overview*

application-tracking

Syntax
```
application-tracking {
    disable;
    (first-update | first-update-interval first-update-interval);
    session-update-interval session-update-interval;
}
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 10.2. Support for **disable** added in Junos OS Release 11.4.

Description   AppTrack, an application tracking tool, is a form of statistical profiling. Enabling this feature for a zone logs flow statistics (the byte count, packet count, and start and end times for a session) at session end. You can modify the logging time and log frequency with command options. Periodically, a network management tool, such as STRM, collects the logged statistics sent by each network device for bandwidth usage analysis of the network.

Options   The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *Example: Configuring AppTrack*

auth-entry

Syntax

```
auth-entry {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level   [edit system security-profile *security-profile-name*]

Release Information   Statement introduced in Junos OS Release 11.2.

Description   Specify the number of firewall authentication entries that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

cluster (Chassis)

Syntax

```
cluster {
    configuration-synchronize {
        no-secondary-bootup-auto;
    }
    control-link-recovery;
    heartbeat-interval milliseconds;
    heartbeat-threshold number;
    network-management {
        cluster-master;
    }
    redundancy-group group-number {
        gratuitous-arp-count number;
        hold-down-interval number;
        interface-monitor interface-name {
            weight number;
        }
        ip-monitoring {
            family {
                inet {
                    ipv4-address {
                        interface {
                            logical-interface-name;
                            secondary-ip-address ip-address;
                        }
                        weight number;
                    }
                }
            }
            global-threshold number;
            global-weight number;
            retry-count number;
            retry-interval seconds;
        }
        node (0 | 1 ) {
            priority number;
        }
        preempt;
    }
    reth-count number;
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            (world-readable | no-world-readable);
            size maximum-file-size;
        }
        flag flag;
        level {
            (alert | all | critical | debug | emergency | error | info | notice | warning);
        }
        no-remote-trace;
```

```
            }
          }
```

Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 9.0.
Description	Configure a chassis cluster.
Options	The remaining statements are explained separately. See CLI Explorer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	• *ip-monitoring*

cpu

Syntax
```
cpu {
    reserved percent;
}
```

Hierarchy Level   [edit system security-profile *security-profile-name*]

Release Information   Statement introduced in Junos OS Release 11.4.

Description   Specify the percentage of CPU utilization that is always available to a logical system. This value is configured in a security profile that is bound to a logical system.

Only the master administrator can create security profiles and bind them to logical systems.

> *i* NOTE: The cpu-control option at the [edit system security-profile resources] hierarchy level must be specified for the reserved value to take effect.

Options   reserved *percent*—A reserved quota that guarantees that the percentage of CPU specified is always available to the logical system.

Range: 0 through 100 percent (decimal point allowed).

Default: 1 percent for the master logical system and 0 percent for user logical systems.

> ⚠ CAUTION: The master logical system must not be bound to a security profile that is configured with a 0 percent reserved CPU quota as traffic loss could occur.

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

Related   • Understanding Logical System Security Profiles (Master Administrators Only) on
Documentation   page 51

datapath-debug

Syntax
```
datapath-debug {
    action-profile profile-name {
        event (jexec | lbt | lt-enter | lt-leave | mac-egress | mac-ingress | np-egress | np-ingress
            | pot) {
            count;
            packet-dump;
            packet-summary;
            trace;
        }
        module {
            flow {
                flag {
                    all;
                }
            }
        }
        preserve-trace-order;
        record-pic-history;
    }
    capture-file {
        filename;
        files number;
        format pacp-format;
        size maximum-file-size;
        (world-readable | no-world-readable);
    }
    maximum-capture-size value;
    packet-filter packet-filter-name {
        action-profile (profile-name | default);
        destination-port (port-range | protocol-name);
        destination-prefix destination-prefix;
        interface logical-interface-name;
        protocol (protocol-number | protocol-name;
        source-port (port-range | protocol- name);
        source-prefix source-prefix;
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        no-remote-trace;
    }
}
```

Hierarchy Level [edit security]

Release Information Command introduced in Junos OS Release 10.0.

Description   Configure the data path debugging options.

Options   The remaining statements are explained separately. See CLI Explorer.

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

Related   • Understanding Data Path Debugging for Logical Systems
Documentation

dslite-softwire-initiator

Syntax
```
dslite-softwire-initiator {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level   [edit system security-profile security-profile-name]

Release Information   Statement introduced in Junos OS Release 12.1.

Description   Specify the number of IPv6 dual-stack lite (DS-Lite) softwire initiators that can connect to the softwire concentrator configured in either a user logical system or the master logical system. This statement is configured in the security profile that is bound to the logical system.

Only the master administrator can create security profiles and bind them to logical systems. The master administrator:

• Uses security profiles to provision logical systems with resources

• Binds security profiles to user logical systems and the master logical system

• Configures more than one security profile, specifying different amounts of resource allocations in various profiles

Options   • maximum amount—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources. The default is the system maximum.

• reserved amount—A reserved quota that guarantees that the resource amount specified is always available to the logical system. The default is 0.

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

Related   • Understanding IPv6 Dual-Stack Lite in Logical Systems on page 238
Documentation

file (System Logging)

Syntax	file *filename* {
	    allow-duplicates;
	    any (alert | any | critical | emergency | error | info | none | notice | warning);
	    archive {
	      archive-sites {
	        url *password*;
	      }
	      (binary-data | no-binary-data);
	      files *number*;
	      size *size*;
	      start-time *start-time*;
	      transfer-interval *transfer-interval*;
	      (world-readable | no-world-readable);
	    }
	    authorization (alert | any | critical | emergency | error | info | none | notice | warning);
	    change-log (alert | any | critical | emergency | error | info | none | notice | warning);
	    conflict-log (alert | any | critical | emergency | error | info | none | notice | warning);
	    daemon (alert | any | critical | emergency | error | info | none | notice | warning);
	    dfc (alert | any | critical | emergency | error | info | none | notice | warning);
	    explicit-priority;
	    external (alert | any | critical | emergency | error | info | none | notice | warning);
	    firewall (alert | any | critical | emergency | error | info | none | notice | warning);
	    ftp (alert | any | critical | emergency | error | info | none | notice | warning);
	    interactive-commands (alert | any | critical | emergency | error | info | none | notice | warning);
	    kernel (alert | any | critical | emergency | error | info | none | notice | warning);
	    match "*regular-expression*";
	    ntp (alert | any | critical | emergency | error | info | none | notice | warning);
	    pfe (alert | any | critical | emergency | error | info | none | notice | warning);
	    security (alert | any | critical | emergency | error | info | none | notice | warning);
	    structured-data {
	      brief;
	    }
	    user (alert | any | critical | emergency | error | info | none | notice | warning);
	}

Hierarchy Level	[edit system syslog]
Release Information	Statement introduced before Junos OS Release 12.1X47 for SRX Series.
Description	Specify the file in which to log data.

Options

- *filename*—Specify the name of the file in which to log data.

- *allow-duplicates*—Do not suppress the repeated messages.

- *any*—Specify all facilities information.

  - *alert*—Specify the conditions that should be corrected immediately.

  - *critical*—Specify the critical conditions.

  - *emergency*—Specify the conditions that cause security functions to stop.

  - *error*—Specify the general error conditions.

- *info*—Specify the information about normal security operations.

- *none*—Do not specify any messages.

- *notice*—Specify the conditions that should be handled specifically.

- *warning*—Specify the general warning conditions.

- *archive*—Specify the archive file information.

  - *archive-sites*—Specify a list of destination URLs for the archived log files.

    - *url*—Specify the primary and failover URLs to receive archive files.

  - *binary-data*—Mark file such that it contains binary data.

  - *no-binary-data*—Do not mark the file such that it contains binary data.

  - *files*—Specify the number of files to be archived. Range: 1 through 1000 files.

  - *size*—Specify the size of files to be archived. Range: 65,536 through 1,073,741,824 bytes.

  - *world-readable*—Allow any user to read the log file.

  - *no-world-readable*—Do not allow any user to read the log file.

  - *start-time*—Specify the start time for file transmission. Enter the start time in the yyyy-mm-dd.hh:mm format.

  - *transfer-interval*—Specify the frequency at which to transfer the files to archive sites.

- *authorization*—Specify the authorization system.

- *change-log*—Specify the configuration change log.

- *conflict-log*—Specify the configuration conflict log.

- *daemon*—Specify the various system processes.

- *dfc*—Specify the dynamic flow capture.

- *explicit-priority*—Include the priority and facility in messages.

- *external*—Specify the local external applications.

- *firewall*—Specify the firewall filtering system.

- *ftp*—Specify the FTP process.

- *interactive-commands*—Specify the commands executed by the UI.

- *kernel*—Specify the kernel information.

- *match*—Specify the regular expression for lines to be logged.

- *ntp*—Specify the NTP process.

- *pfe*—Specify the Packet Forwarding Engine.

- *security*—Specify the security-related information.

- *structured-data*—Log the messages in structured log format.

  - *brief*—Omit English language text from the end of the logged message.

- *user*—Specify the user processes.

  - *info*—Specify the informational messages.

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

firewall-authentication (Security)

Syntax
```
firewall-authentication {
    traceoptions {
        flag flag;
    }
}
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 8.5.

Description   Define data-plane firewall authentication tracing options.

Options
- **flag**—Trace operation to perform. To specify more than one trace operation, include multiple flag statements.

  - **all**—Enable all tracing operations.

  - **authentication**—Trace data-plane firewall authentication events.

  - **proxy**—Trace data-plane firewall authentication proxy events.

- **detail**—Display moderate amount of data.

- **extensive**—Display extensive amount of data.

- **terse**—Display minimum amount of data.

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

Related Documentation
- *Firewall User Authentication Overview*

- Understanding Logical System Firewall Authentication on page 65

flow (Security Flow)

Syntax

```
flow {
    aging {
        early-ageout seconds;
        high-watermark percent;
        low-watermark percent;
    }
    allow-dns-reply;
    ethernet-switching {
        block-non-ip-all;
        bpdu-vlan-flooding;
        bypass-non-ip-unicast;
        no-packet-flooding {
            no-trace-route;
        }
    }
    force-ip-reassembly;
    ipsec-performance-acceleration;
    load distribution {
        session-affinity ipsec;
    }
    pending-sess-queue-length (high | moderate | normal);
    route-change-timeout seconds;
    syn-flood-protection-mode (syn-cookie | syn-proxy);
    tcp-mss {
        all-tcp mss value;
        gre-in {
            mss value;
        }
        gre-out {
            mss value;
        }
        ipsec-vpn {
            mss value;
        }
    }
    tcp-session {
        fin-invalidate-session;
        no-sequence-check;
        no-syn-check;
        no-syn-check-in-tunnel;
        rst-invalidate-session;
        rst-sequence-check;
        strict-syn-check;
        tcp-initial-timeout seconds;
        time-wait-state {
            (session-ageout | session-timeout seconds);
        }
    }
    traceoptions {
        file {
            filename;
            files number;
```

```
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
          }
          flag flag;
          no-remote-trace;
          packet-filter filter-name {
            destination-port port-identifier;
            destination-prefix address;
            interface interface-name;
            protocol protocol-identifier;
            source-port port-identifier;
            source-prefix address;
          }
          rate-limit messages-per-second;
      }
    }
```

Hierarchy Level   [edit security]

Release Information   Statement modified in Junos OS Release 9.5.

Description   Determine how the device manages packet flow. The device can regulate packet flow in the following ways:

- Enable or disable DNS replies when there is no matching DNS request.

- Set the initial session-timeout values.

Options   The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- Juniper Networks Devices Processing Overview

- Understanding Session Characteristics for SRX Series Services Gateways

- Understanding Flow in Logical Systems for SRX Series Devices

flow-gate

Syntax
```
flow-gate {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of flow gates, also known as **pinholes**, that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

flow-session

Syntax	flow-session {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the number of flow sessions that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

> **NOTE:** An IPv6 session consumes twice the memory of an IPv4 session. Therefore the number of sessions available for IPv6 is half the reserved and maximum quotas configured for the flow session resource in a security profile. Use the vty command **show usp flow resource usage cp-session** to check flow session usage.

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

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 56

idp (Security)

Syntax
```
idp {
    active-policy policy-name;
    custom-attack attack-name {
        attack-type {
            anomaly {
                direction (any | client-to-server | server-to-client);
                service service-name;
                shellcode (all | intel | no-shellcode | sparc);
                test test-condition;
            }
            chain {
                expression boolean-expression;
                member member-name {
                    attack-type {
                        (anomaly ...same statements as in [edit security idp custom-attack attack-name
                            attack-type anomaly] hierarchy level | signature ...same statements as in [edit
                            security idp custom-attack attack-name attack-type signature] hierarchy
                            level);
                    }
                }
                order;
                protocol-binding {
                    application application-name;
                    icmp;
                    icmpv6;
                    ip {
                        protocol-number transport-layer-protocol-number;
                    }
                    ipv6 {
                        protocol-number transport-layer-protocol-number;
                    }
                    rpc {
                        program-number rpc-program-number;
                    }
                    tcp {
                        minimum-port port-number <maximum-port port-number>;
                    }
                    udp {
                        minimum-port port-number <maximum-port port-number>;
                    }
                }
                reset;
                scope (session | transaction);
            }
            signature {
                context context-name;
                direction (any | client-to-server | server-to-client);
                negate;
                pattern signature-pattern;
                protocol {
                    icmp {
                        code {
```

```
            match (equal | greater-than | less-than | not-equal);
            value code-value;
         }
         data-length {
            match (equal | greater-than | less-than | not-equal);
            value data-length;
         }
         identification {
            match (equal | greater-than | less-than | not-equal);
            value identification-value;
         }
         sequence-number {
            match (equal | greater-than | less-than | not-equal);
            value sequence-number;
         }
         type {
            match (equal | greater-than | less-than | not-equal);
            value type-value;
         }
      }
      ipv4 {
         destination {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
         }
         identification {
            match (equal | greater-than | less-than | not-equal);
            value identification-value;
         }
         ip-flags {
            (df | no-df);
            (mf | no-mf);
            (rb | no-rb);
         }
         protocol {
            match (equal | greater-than | less-than | not-equal);
            value transport-layer-protocol-id;
         }
         source {
            match (equal | greater-than | less-than | not-equal);
            value ip-address-or-hostname;
         }
         tos {
            match (equal | greater-than | less-than | not-equal);
            value type-of-service-in-decimal;
         }
         total-length {
            match (equal | greater-than | less-than | not-equal);
            value total-length-of-ip-datagram;
         }
         ttl {
            match (equal | greater-than | less-than | not-equal);
            value time-to-live;
         }
      }
      ipv6 {
```

```
destination {
  match (equal | greater-than | less-than | not-equal);
  value ip-address-or-hostname;
}
flow-label {
  match (equal | greater-than | less-than | not-equal);
  value flow-label-value;
}
hop-limit {
  match (equal | greater-than | less-than | not-equal);
  value hop-limit-value;
}
next-header {
  match (equal | greater-than | less-than | not-equal);
  value next-header-value;
}
payload-length {
  match (equal | greater-than | less-than | not-equal);
  value payload-length-value;
}
source {
  match (equal | greater-than | less-than | not-equal);
  value ip-address-or-hostname;
}
traffic-class {
  match (equal | greater-than | less-than | not-equal);
  value traffic-class-value;
}
tcp {
  ack-number {
    match (equal | greater-than | less-than | not-equal);
    value acknowledgement-number;
  }
  data-length {
    match (equal | greater-than | less-than | not-equal);
    value tcp-data-length;
  }
  destination-port {
    match (equal | greater-than | less-than | not-equal);
    value destination-port;
  }
  header-length {
    match (equal | greater-than | less-than | not-equal);
    value header-length;
  }
  mss {
    match (equal | greater-than | less-than | not-equal);
    value maximum-segment-size;
  }
  option {
    match (equal | greater-than | less-than | not-equal);
    value tcp-option;
  }
  sequence-number {
    match (equal | greater-than | less-than | not-equal);
    value sequence-number;
```

```
                              }
                              source-port {
                                 match (equal | greater-than | less-than | not-equal);
                                 value source-port;
                              }
                              tcp-flags {
                                 (ack | no-ack);
                                 (fin | no-fin);
                                 (psh | no-psh);
                                 (r1 | no-r1);
                                 (r2 | no-r2);
                                 (rst | no-rst);
                                 (syn | no-syn);
                                 (urg | no-urg);
                              }
                              urgent-pointer {
                                 match (equal | greater-than | less-than | not-equal);
                                 value urgent-pointer;
                              }
                              window-scale {
                                 match (equal | greater-than | less-than | not-equal);
                                 value window-scale-factor;
                              }
                              window-size {
                                 match (equal | greater-than | less-than | not-equal);
                                 value window-size;
                              }
                           }
                           udp {
                              data-length {
                                 match (equal | greater-than | less-than | not-equal);
                                 value data-length;
                              }
                              destination-port {
                                 match (equal | greater-than | less-than | not-equal);
                                 value destination-port;
                              }
                              source-port {
                                 match (equal | greater-than | less-than | not-equal);
                                 value source-port;
                              }
                           }
                        }
                        protocol-binding {
                           application application-name;
                           icmp;
                           icmpv6;
                           ip {
                              protocol-number transport-layer-protocol-number;
                           }
                           ipv6 {
                              protocol-number transport-layer-protocol-number;
                           }
                           rpc {
                              program-number rpc-program-number;
                           }
```

```
                            tcp {
                              minimum-port port-number <maximum-port port-number>;
                            }
                            udp {
                              minimum-port port-number <maximum-port port-number>;
                            }
                          }
                          regexp regular-expression;
                          shellcode (all | intel | no-shellcode | sparc);
                        }
                      }
                      recommended-action (close | close-client | close-server | drop | drop-packet | ignore |
                        none);
                      severity (critical | info | major | minor | warning);
                      time-binding {
                        count count-value;
                        scope (destination | peer | source);
                      }
                    }
                    custom-attack-group custom-attack-group-name {
                      group-members [attack-or-attack-group-name];
                    }
                    dynamic-attack-group dynamic-attack-group-name {
                      filters {
                        category {
                          values [category-value];
                        }
                        direction {
                          expression (and | or);
                          values [any client-to-server exclude-any exclude-client-to-server
                            exclude-server-to-client server-to-client];
                        }
                        false-positives {
                          values [frequently occasionally rarely unknown];
                        }
                        performance {
                          values [fast normal slow unknown];
                        }
                        products {
                          values [product-value];
                        }
                        recommended;
                        service {
                          values [service-value];
                        }
                        severity {
                          values [critical info major minor warning];
                        }
                        type {
                          values [anomaly signature];
                        }
                      }
                    }
                    idp-policy policy-name {
                      rulebase-exempt {
                        rule rule-name {
```

```
                              description text;
                              match {
                                attacks {
                                  custom-attack-groups [attack-group-name];
                                  custom-attacks [attack-name];
                                  dynamic-attack-groups [attack-group-name];
                                  predefined-attack-groups [attack-group-name];
                                  predefined-attacks [attack-name];
                                }
                                destination-address ([address-name] | any | any-ipv4 | any-ipv6);
                                destination-except [address-name];
                                from-zone (zone-name | any );
                                source-address ([address-name] | any | any-ipv4 | any-ipv6);
                                source-except [address-name];
                                to-zone (zone-name | any);
                              }
                            }
                          }
                          rulebase-ips {
                            rule rule-name {
                              description text;
                              match {
                                application (application-name | any | default);
                                attacks {
                                  custom-attack-groups [attack-group-name];
                                  custom-attacks [attack-name];
                                  dynamic-attack-groups [attack-group-name];
                                  predefined-attack-groups [attack-group-name];
                                  predefined-attacks [attack-name];
                                }
                                destination-address ([address-name] | any | any-ipv4 | any-ipv6);
                                destination-except [address-name];
                                from-zone (zone-name | any );
                                source-address ([address-name] | any | any-ipv4 | any-ipv6);
                                source-except [address-name];
                                to-zone (zone-name | any);
                              }
                              terminal;
                              then {
                                action {
                                  class-of-service {
                                    dscp-code-point number;
                                    forwarding-class forwarding-class;
                                  }
                                  (close-client | close-client-and-server | close-server |drop-connection |
                                    drop-packet | ignore-connection | mark-diffserv value | no-action |
                                    recommended);
                                }
                                ip-action {
                                  (ip-block | ip-close | ip-notify);
                                  log;
                                  log-create;
                                  refresh-timeout;
                                  target (destination-address | service | source-address | source-zone |
                                    source-zone-address | zone-service);
                                  timeout seconds;
```

```
                      }
                      notification {
                        log-attacks {
                          alert;
                        }
                        packet-log {
                          post-attack number;
                          post-attack-timeout seconds;
                          pre-attack number;
                        }
                      }
                      severity (critical | info | major | minor | warning);
                  }
                }
              }
            }
            security-package {
              automatic {
                download-timeout minutes;
                enable;
                interval hours;
                start-time start-time;
              }
              install {
                ignore-version-check;
              }
              source-address address;
              url url-name;
            }
            sensor-configuration {
              application-identification {
                max-packet-memory value;
                max-tcp-session-packet-memory value;
                max-udp-session-packet-memory value;
              }
              detector {
                protocol-name protocol-name {
                  tunable-name  tunable-name {
                    tunable-value protocol-value;
                  }
                }
              }
              flow {
                (allow-icmp-without-flow | no-allow-icmp-without-flow);
                fifo-max-size value;
                hash-table-size value;
                (log-errors | no-log-errors);
                max-session-offset value;
                max-timers-poll-ticks value;
                reject-timeout value;
                (reset-on-policy | no-reset-on-policy);
                udp-anticipated-timeout value;
              }
              global {
                (enable-all-qmodules | no-enable-all-qmodules);
                (enable-packet-pool | no-enable-packet-pool);
```

```
            gtp (decapsulation | no-decapsulation);
            memory-limit-percent value;
            (policy-lookup-cache | no-policy-lookup-cache);
        }
        high-availability {
            no-policy-cold-synchronization;
        }
        ips {
            content-decompression-max-memory-kb value;
            content-decompression-max-ratio value;
            (detect-shellcode | no-detect-shellcode);
            fifo-max-size value;
            (ignore-regular-expression | no-ignore-regular-expression);
            log-supercede-min minimum-value;
            pre-filter-shellcode;
            (process-ignore-s2c | no-process-ignore-s2c);
            (process-override | no-process-override);
            process-port port-number;
        }
        log {
            cache-size size;
            suppression {
                disable;
                (include-destination-address | no-include-destination-address);
                max-logs-operate value;
                max-time-report value;
                start-log value;
            }
        }
        packet-log {
            host ip-address <port number>;
            max-sessions percentage;
            source-address ip-address;
            total-memory percentage;
        }
        re-assembler {
            action-on-reassembly-failure (drop | drop-session | ignore);
            (force-tcp-window-checks | no-force-tcp-window-checks);
            (ignore-memory-overflow | no-ignore-memory-overflow);
            (ignore-reassembly-memory-overflow | no-ignore-reassembly-memory-overflow);
            ignore-reassembly-overflow;
            max-flow-mem value;
            max-packet-mem value;
            (tcp-error-logging | no-tcp-error-logging);
        }
        ssl-inspection {
            cache-prune-chunk-size number;
            key-protection;
            maximum-cache-size number;
            session-id-cache-timeout seconds;
            sessions number;
        }
    }
    traceoptions {
        file {
            filename;
```

```
              files number;
              match regular-expression;
              size maximum-file-size;
              (world-readable | no-world-readable);
            }
            flag all;
            level (all | error | info | notice | verbose | warning);
            no-remote-trace;
          }
        }
```

Hierarchy Level	[edit security]
Release Information	Statement modified in Junos OS Release 9.3. The **expression** option added in Junos OS Release 11.4.
Description	Configure Intrusion Detection and Prevention (IDP) to selectively enforce various IDP attack detection and prevention techniques on the network.
Options	The remaining statements are explained separately. See CLI Explorer.
Required Privilege Level	security—To view this statement in the configuration. security-control—To add this statement to the configuration.
Related Documentation	• *Understanding Intrusion Detection and Prevention for SRX Series*

idp-policy

Syntax	idp-policy *idp-policy-name*;
Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the IDP policy for the security profile.
Options	*idp-policy-name*—Name of the IDP policy.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	• *Understanding Intrusion Detection and Prevention for SRX Series*

ike (Security)

Syntax
```
ike {
    gateway gateway-name {
        address [ip-address-or-hostname];
        advpn {
            suggester {
                disable;
            }
            partner {
                connection-limit <number>;
                idle-threshold <packets/sec>;
                idle-time <seconds>;
                disable;
            }
        }
        dead-peer-detection {
            (always-send | optimized | probe-idle-tunnel);
            interval seconds;
            threshold number;
        }
        dynamic {
            connections-limit number;
            (distinguished-name <container container-string> <wildcard wildcard-string> |
                hostname domain-name | inet ip-address | inet6 ipv6-address | user-at-hostname
                e-mail-address);
            ike-user-type (group-ike-id | shared-ike-id);
        }
        external-interface external-interface-name;
        fragmentation {
            enable:
            size bytes;
        }
        general-ikeid;
        ike-policy policy-name;
        local-address (ipv4-address | ipv6-address);
        local-identity {
            (distinguished-name | hostname hostname | inet ip-address | inet6 ipv6-address |
                user-at-hostname e-mail-address);
        }
        nat-keepalive seconds;
        no-nat-traversal;
        remote-identity {
            (distinguished-name <container container-string> <wildcard wildcard-string> |
                hostname hostname | inet ip-address | inet6 ipv6-address | user-at-hostname
                e-mail-address);
        }
        version (v1-only | v2-only);
        xauth {
            access-profile profile-name;
        }
    }
    policy policy-name {
        certificate {
```

```
            local-certificate certificate-id;
            peer-certificate-type (pkcs7 | x509-signature);
            policy-oids [ oid ];
         }
         description description;
         mode (aggressive | main);
         pre-shared-key (ascii-text key | hexadecimal key);
         proposal-set (basic | compatible | standard |suiteb-gcm-128 | suiteb-gcm-256);
         proposals [proposal-name];
         reauth-frequency number;
      }
      proposal proposal-name {
         authentication-algorithm (md5 | sha-256 | sha-384| sha1);
         authentication-method (dsa-signatures | ecdsa-signatures-256 | ecdsa-signatures-384
            | pre-shared-keys | rsa-signatures);
         description description;
         dh-group (group1 | group14 | group19 | group2 | group20 | group24 | group5);
         encryption-algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
         lifetime-seconds seconds;
      }
      respond-bad-spi <max-responses>;
      traceoptions {
         file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
         }
         flag flag;
         no-remote-trace;
         rate-limit messages-per-second;
      }
   }
}
```

Hierarchy Level	[edit security]
Release Information	Statement modified in Junos OS Release 8.5. Support for IPv6 addresses added in Junos OS Release 11.1. The **inet6** option added in Junos OS Release 11.1.
Description	Define Internet Key Exchange (IKE) configuration.
Options	The remaining statements are explained separately. See CLI Explorer.
Required Privilege Level	security—To view this statement in the configuration. security-control—To add this statement to the configuration.
Related Documentation	• *IPsec VPN Overview* • *ALG Overview* • Understanding Logical Systems for SRX Series Services Gateways on page 3

ipsec (Security)

Syntax   ipsec {

```
ipsec {
    policy policy-name {
        description description;
        perfect-forward-secrecy keys (group1 | group14 | group19 | group2 | group20 | group24 |
            group5);
        proposal-set (basic | compatible | standard | suiteb-gcm-128 | suiteb-gcm-256);
        proposals [proposal-name];
    }
    proposal proposal-name {
        authentication-algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
        description description;
        encryption-algorithm (3des-cbc | aes-128-cbc | aes-128-gcm | aes-192-cbc | aes-192-gcm
            | aes-256-cbc | aes-256-gcm | des-cbc);
        lifetime-kilobytes kilobytes;
        lifetime-seconds seconds;
        protocol (ah | esp);
    }
    security-association sa-name {
        manual {
            direction bidirectional {
                authentication {
                    algorithm (hmac-md5-96 | hmac-sha1-96);
                    key {
                        ascii-text key;
                        hexadecimal key;
                    }
                }
                auxiliary-spi auxiliary-spi-value;
                encryption {
                    algorithm (3des-cbc | des-cbc | null);
                    key {
                        ascii-text key;
                        hexadecimal key;
                    }
                }
                protocol (ah | esp);
                spi spi-value;
            }
        }
        mode transport;
    }
    traceoptions {
        flag flag;
    }
    vpn vpn-name {
        bind-interface interface-name;
        copy-outer-dscp;
        establish-tunnels (immediately | on-traffic);
        ike {
            gateway gateway-name;
            idle-time seconds;
            install-interval seconds;
```

```
                                    ipsec-policy ipsec-policy-name;
                                    no-anti-replay;
                                    proxy-identity {
                                        local ip-prefix;
                                        remote ip-prefix;
                                        service (any | service-name);
                                    }
                                }
                                manual {
                                    authentication {
                                        algorithm (hmac-md5-96 | hmac-sha-256-128 | hmac-sha1-96);
                                        key (ascii-text key | hexadecimal key);
                                    }
                                    encryption {
                                        algorithm (3des-cbc | aes-128-cbc | aes-192-cbc | aes-256-cbc | des-cbc);
                                        key (ascii-text key | hexadecimal key);
                                    }
                                    external-interface external-interface-name;
                                    gateway ip-address;
                                    protocol (ah | esp);
                                    spi spi-value;
                                }
                                traffic-selector traffic-selector-name {
                                    local-ip ip-address/netmask;
                                    remote-ip ip-address/netmask;
                                }
                                vpn-monitor {
                                    destination-ip ip-address;
                                    optimized;
                                    source-interface interface-name;
                                    verify-path {
                                        destination-ip ip-address;
                                    }
                                }
                            }
                            vpn-monitor-options {
                                interval seconds;
                                threshold number;
                            }
                        }
```

Hierarchy Level [edit security]

Release Information Statement modified in Junos OS Release 8.5.

Description Define IPsec configuration.

Options The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *IPsec VPN Overview*

log (Security)

Syntax
```
log {
    cache {
        exclude exclude-name {
            destination-address destination-address;
            destination-port destination-port;
            event-id event-id;
            failure;
            interface-name interface-name;
            policy-name policy-name;
            process process-name;
            protocol protocol;
            source-address source-address;
            source-port source-port;
            success;
            user-name user-name;
        }
        limit value;
    }
    disable;
    event-rate rate;
    facility-override;
    file {
        files max-file-number;
        name file-name;
        path binary-log-file-path;
        size maximum-file-size;
    }
    format (binary | sd-syslog | syslog);
    max-database-record
    mode (event | stream);
    rate-cap <rate-cap-value;rate-limit> (0.5000 logs per second);
    (source-address source-address | source-interface interface-name);
    stream stream-name {
        category (all | content-security | fw-auth | screen | alg | nat | flow | sctp | gtp | ipsec | idp
            | rtlog |pst-ds-lite | appqos |secintel);
        file {
            name file-name;
            size file-size;
            rotation max-rotation-number;
        }
        filter
        threat-attack
        format (binary | sd-syslog | syslog | welf);
        host {
            ip-address;
            port port-number;
        }
        rate-limit (1..65535 logs per second)
        severity (alert | critical | debug | emergency | error | info | notice | warning);
    }
    traceoptions {
        file {
```

```
            filename;
            files number;
            match regular-expression;
            size maximum-file-size (10240..1073741824);
            world-readable
            no-world-readable
        }
        flag (all | configuration | hpl | report | source);
        no-remote-trace (file | flag);
    }
    transport {
        protocol (udp | tcp | tls);
        tls-profile tls-profile-name;
        tcp-connections tcp-connections;
    }
    utc-timestamp;
}
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 9.2.

Description   You can set the mode of logging (event for traditional system logging or stream for streaming security logs through a revenue port to a server). You can also specify all the other parameters for security logging.

Options
- **disable**—Disable the security logging for the device.
- **event-rate** *rate*—Limits the rate (0 through 1500) at which logs will be streamed per second.
- **rate-cap** *rate-cap-value*—Works with event mode only. Limits the rate (0 through 5000) at which data plane logs will be generated per second.
- **stream**—Every stream can configure file or host.
- **file-name**—Specify the file name.
- **file-size**—Specify the file size.
  - SRX1500 - The default value is 25M and the range is 10M through 50M.
  - SRX4100 and SRX4200 - The default value is 25M and the range is 10M through 50M.
  - vSRX - The default value is 2M and the range is 1M through 3M.
- **rotation**—Configure the max file number for rotation. The default value is 10 and the range is 2 through 19.
- **max-database-record**—The following are the disk usage range limit for database:
  - SRX1500: 0 through 15,000,000
  - SRX4100 and SRX4200: 0 through 15,000,000
  - vSRX: 0 through 10,000,000

> *i* NOTE: Be sure there is enough free space in **/var/log/hostlogs/**, otherwise logs may be dropped when written into database.

- **source-address** *source-address*—Specify a source IP address or IP address used when exporting security logs, which is mandatory to configure **stream host**.
- **source-interface** *interface-name*—Specify a source interface name, which is mandatory to configure **stream host**.

> *i* NOTE: The **source-address** and **source-interface** are alternate values. Using one of the options is mandatory.

- **utc-time-stamp**—Specify to use UTC time for security log timestamps.

The remaining statements are explained separately. See CLI Explorer.

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

---

logical-system (System Security Profile)

Syntax	logical-system *logical-system-name*;
Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the user logical system to bind the security profile to.

The master administrator uses security profiles to provision logical systems with resources. You can bind security profiles to user logical systems and the master logical system. The master administrator can configure more than one security profile allocating different amounts of a resource in various ones.

Only the master administrator can create security profiles and bind them to logical systems.

Options	*logical-system-name*—Name of the logical system.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	• Understanding Logical Systems for SRX Series Services Gateways on page 3

logical-systems (All)

Syntax
```
logical-systems logical-system-name {
    ...
}
```

Hierarchy Level   [edit]

Release Information   Statement introduced in Junos OS Release 11.2.

Description   Configure a logical system. Only the master administrator can configure a logical system at this hierarchy level.

At this hierarchy level, you can include several of the hierarchies that can be included at the [**edit**] hierarchy level. For descriptions of the applicable statements, see the appropriate hierarchies.

> **NOTE:** The **logical-systems** configuration statement can only be used by the master administrator.

Options   *logical-system-name*—Name of the logical system.

Required Privilege   all—To view this statement in the configuration.
Level   all—To add this statement to the configuration.

Related Documentation   • Understanding Logical Systems for SRX Series Services Gateways on page 3

nat

Syntax

```
nat {
    destination {
        pool pool-name {
            address ip-address {
                (port port-number | to ip-address);
            }
            description text;
            routing-instance routing-instance-name;
        }
        rule-set rule-set-name {
            description text;
            from {
                interface [interface-name];
                routing-instance [routing-instance-name];
                zone [zone-name];
            }
            rule rule-name {
                description text;
                match {
                    (destination-address <ip-address> | destination-address-name <address-name>);
                    destination-port port-number;
                    protocol [protocol-name-or-number];
                    source-address [ip-address];
                    source-address-name [address-name];
                }
                then {
                    destination-nat (off | pool pool-name);
                }
            }
        }
    }
    proxy-arp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    proxy-ndp {
        interface interface-name {
            address ip-address {
                to ip-address;
            }
        }
    }
    source {
        address-persistent;
        interface {
            port-overloading {
                off;
            }
        }
```

```
pool pool-name {
  address ip-address {
    to ip-address;
  }
  description text;
  host-address-base ip-address;
  overflow-pool (interface | pool-name);
  port {
    (no-translation | port-overloading-factor number | range port-low <to port-high>);
  }
  routing-instance routing-instance-name;
}
pool-default-port-range lower-port-range to upper-port-range;
pool-utilization-alarm {
  clear-threshold value;
  raise-threshold value;
}
port-randomization {
  disable;
}
port-round-robin {
  disable;
}
rule-set rule-set-name {
  description text;
  from {
    interface [interface-name];
    routing-instance [routing-instance-name];
    zone [zone-name];
  }
  rule rule-name {
    description text;
    match {
      (destination-address <ip-address> | destination-address-name <address-name>);
      destination-port port-number;
      protocol [protocol-name-or-number];
      source-address [ip-address];
      source-address-name [address-name];
    }
    then {
      source-nat {
        interface {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number value;
            permit (any-remote-host | target-host | target-host-port);
          }
        }
        off;
        pool {
          persistent-nat {
            address-mapping;
            inactivity-timeout seconds;
            max-session-number number;
            permit (any-remote-host | target-host | target-host-port);
```

```
                    }
                    pool-name;
                  }
                }
              }
            }
          }
          to {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
          }
        }
      }
      static {
        rule-set rule-set-name {
          description text;
          from {
            interface [interface-name];
            routing-instance [routing-instance-name];
            zone [zone-name];
          }
          rule rule-name {
            description text;
            match {
              (destination-address ip-address | destination-address-name address-name);
            }
            then {
              static-nat {
                inet {
                  routing-instance (default | routing-instance-name);
                }
                prefix {
                  address-prefix;
                  routing-instance (default | routing-instance-name);
                }
                prefix-name {
                  address-prefix-name;
                  routing-instance (default | routing-instance-name);
                }
              }
            }
          }
        }
      }
      traceoptions {
        file {
          filename;
          files number;
          match regular-expression;
          size maximum-file-size;
          (world-readable | no-world-readable);
        }
        flag flag;
        no-remote-trace;
      }
    }
```

Hierarchy Level	[edit security]
Release Information	Statement modified in Junos OS Release 9.6. The **description** option added in Junos OS Release 12.1.
Description	Configure Network Address Translation (NAT) for SRX Series devices.
Options	The remaining statements are explained separately. See CLI Explorer.
Required Privilege Level	security—To view this statement in the configuration. security-control—To add this statement to the configuration.
Related Documentation	• *Introduction to NAT* • Understanding Logical System Network Address Translation on page 153

nat-cone-binding

Syntax

```
nat-cone-binding {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level

[edit system security-profile *security-profile-name*]

Release Information

Statement introduced in Junos OS Release 11.2.

Description

Specify the number of NAT cone binding configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

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

Related Documentation

- *Introduction to NAT*

nat-destination-pool

Syntax
```
nat-destination-pool {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level   [edit system security-profile *security-profile-name*]

Release Information   Statement introduced in Junos OS Release 11.2.

Description   Specify the number of NAT destination pool configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation   • *Introduction to NAT*

nat-destination-rule

Syntax	nat-destination-rule {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the number of NAT destination rule configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

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

Related Documentation

- *Introduction to NAT*

nat-interface-port-ol (System)

Syntax	nat-interface-port-ol {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *profile-name*]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the security NAT interface port overloading the quota of a logical system.
Options	• **maximum** *amount*—Specify the maximum allowed quota value. **Range**: 0 through 64 • **reserved** *amount*—Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	• *Introduction to NAT*

nat-nopat-address

Syntax	nat-nopat-address {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the number of NAT without port address translation configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

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

Related Documentation

- *Introduction to NAT*

nat-pat-address

Syntax
```
nat-pat-address {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of NAT with port address translation (PAT) configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical System Network Address Translation on page 153
- *Introduction to NAT*

nat-pat-portnum

Syntax
```
nat-pat-portnum {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the maximum quantity and the reserved quantity of ports for the logical system as part of its security profile. The total number of PAT pools must not exceed the configured maximum ports for the logical system.

Options maximum *amount*—Specify the maximum number of ports allowed for a logical system. The maximum quantity is not guaranteed and is shared among multiple logical systems.

reserved *amount*—Specify the number of resources guaranteed for a logical system.

Range: For SRX5600 and SRX5800 devices, up to 402,653,184 ports are supported. Pool specifications for logical systems can be viewed using the **show security nat source summary logical-system all** command.

Required Privilege Level system—To view this statement in the configuration.
system—To add this statement to the configuration.

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

nat-port-ol-ipnumber

Syntax	nat-port-ol-ipnumber {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the number of NAT port overloading IP number configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options

- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

Required Privilege Level

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

Related Documentation

- Understanding Logical Systems for SRX Series Services Gateways on page 3

nat-rule-referenced-prefix (System)

Syntax

```
nat-rule-referenced-prefix {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *profile-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Specify the security NAT rule referenced IP prefix quota of a logical system.

Options
- **maximum** *amount* —Specify the maximum allowed quota value.

  Range: 0 through 1,048,576

- **reserved** *amount* —Specify a reserved quota value that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

nat-source-pool

Syntax	nat-source-pool {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the NAT source pool configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options	• **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.
	• **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.
Related Documentation	• [Understanding Logical Systems for SRX Series Services Gateways on page 3](#)

nat-source-rule

Syntax	nat-source-rule {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the NAT source rule configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.
- binds security profiles to user logical systems and the master logical system.
- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

nat-static-rule

Syntax	nat-static-rule {
	    maximum *amount*;
	    reserved *amount*;
	}

Hierarchy Level	[edit system security-profile *security-profile-name*]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Specify the number of NAT static rule configurations that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

policies

Syntax

```
policies {
    default-policy (deny-all | permit-all);
    from-zone zone-name to-zone zone-name {
        policy policy-name {
            description description;
            match {
                application {
                    [application];
                    any;
                }
                destination-address {
                    [address];
                    any;
                    any-ipv4;
                    any-ipv6;
                }
                source-address {
                    [address];
                    any;
                    any-ipv4;
                    any-ipv6;
                }
                source-identity {
                    [role-name];
                    any;
                    authenticated-user;
                    unauthenticated-user;
                    unknown-user;
                }
            }
            scheduler-name scheduler-name;
            then {
                count {
                    alarm {
                        per-minute-threshold number;
                        per-second-threshold number;
                    }
                }
                deny;
                log {
                    session-close;
                    session-init;
                }
                permit {
                    application-services {
                        application-firewall {
                            rule-set rule-set-name;
                        }
                        application-traffic-control {
                            rule-set rule-set-name;
                        }
                        gprs-gtp-profile profile-name;
```

```
                        gprs-sctp-profile profile-name;
                        idp;
                        redirect-wx | reverse-redirect-wx;
                        ssl-proxy {
                            profile-name profile-name;
                        }
                        uac-policy {
                            captive-portal captive-portal;
                        }
                        utm-policy policy-name;
                    }
                    destination-address {
                        drop-translated;
                        drop-untranslated;
                    }
                    firewall-authentication {
                        pass-through {
                            access-profile profile-name;
                            client-match user-or-group-name;
                            ssl-termination-profile profile-name;
                            web-redirect;
                            web-redirect-to-https;
                        }
                        user-firewall {
                            access-profile profile-name;
                            domain domain-name
                            ssl-termination-profile profile-name;
                        }
                        web-authentication {
                            client-match user-or-group-name;
                        }
                    }
                    services-offload;
                    tcp-options {
                        sequence-check-required;
                        syn-check-required;
                    }
                    tunnel {
                        ipsec-group-vpn group-vpn;
                        ipsec-vpn vpn-name;
                        pair-policy pair-policy;
                    }
                }
            }
            reject;
        }
    }
}
global {
    policy policy-name {
        description description;
        match {
            application {
                [application];
                any;
            }
            destination-address {
```

```
              [address];
              any;
              any-ipv4;
              any-ipv6;
            }
            from-zone {
              [zone-name];
              any;
            }
            source-address {
              [address];
              any;
              any-ipv4;
              any-ipv6;
            }
            source-identity {
              [role-name];
              any;
              authenticated-user;
              unauthenticated-user;
              unknown-user;
            }
            to-zone {
              [zone-name];
              any;
            }
          }
          scheduler-name scheduler-name;
          then {
            count {
              alarm {
                per-minute-threshold number;
                per-second-threshold number;
              }
            }
            deny;
            log {
              session-close;
              session-init;
            }
            permit {
              application-services {
                application-firewall {
                  rule-set rule-set-name;
                }
                application-traffic-control {
                  rule-set rule-set-name;
                }
                gprs-gtp-profile profile-name;
                gprs-sctp-profile profile-name;
                idp;
                redirect-wx | reverse-redirect-wx;
                ssl-proxy {
                  profile-name profile-name;
                }
                uac-policy {
```

```
                          captive-portal captive-portal;
                        }
                        utm-policy policy-name;
                      }
                      destination-address {
                        drop-translated;
                        drop-untranslated;
                      }
                      firewall-authentication {
                        pass-through {
                          access-profile profile-name;
                          client-match user-or-group-name;
                          ssl-termination-profile profile-name;
                          web-redirect;
                          web-redirect-to-https;
                        }
                        web-authentication {
                          client-match user-or-group-name;
                        }
                      }
                      services-offload;
                      tcp-options {
                        initial-tcp-mss mss-value;
                        reverse-tcp-mss mss-value;
                        sequence-check-required;
                        syn-check-required;
                      }
                    }
                    reject;
                  }
                }
              }
            policy-rematch;
            policy-stats {
              system-wide (disable | enable) ;
            }
            traceoptions {
              file {
                filename;
                files number;
                match regular-expression;
                size maximum-file-size;
                (world-readable | no-world-readable);
              }
              flag flag;
              no-remote-trace;
            }
          }
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 8.5.

Support for the **services-offload** option added in Junos OS Release 11.4.

Support for the **source-identitiy** option added in Junos OS Release 12.1.

Support for the **description** option added in Junos OS Release 12.1.

Support for the **ssl-termination-profile** and **web-redirect-to-https** options added on high-end SRX Series Services Gateways starting from Junos OS Release 12.1X44-D10 and on vSRX, SRX300, SRX320, SRX340, SRX345, SRX550M, and SRX1500 Services Gateways starting from Junos OS Release 15.1X49-D40.

Support for the **user-firewall** option added in Junos OS Release 12.1X45-D10.

Support for the **domain** option, and for the **from-zone** and **to-zone** global policy match options, added in Junos OS Release 12.1X47-D10.

Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Support for the **extensive** option for **policy-rematch** added in Junos OS Release 15.1X49-D20.

Description   Configure network security policies.

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

Related Documentation   • *Security Policies Overview*

policy (System Security Profile)

Syntax

```
policy {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify the number of security policies that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

policy-with-count

Syntax

```
policy-with-count {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level   [edit system security-profile *security-profile-name*]

Release Information   Statement introduced in Junos OS Release 11.2.

Description   Specify the number of security policies with a count that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation   • Understanding Logical Systems for SRX Series Services Gateways on page 3

profile (Access)

Syntax   profile *profile-name* {
    accounting {
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
      coa-immediate-update;
      duplication;
      immediate-update;
      order [*accounting-method*];
      statistics (time | volume-time);
      update-interval *minutes*;
    }
    accounting-order [*accounting-method*];
    address-assignment pool *pool-name*;
    authentication-order [ldap | none | password | securid];
    authorization-order [jsrc];
    client *client-name* {
      chap-secret *chap-secret*;
      client-group [*group-names*];
      firewall-user {
        password *password*;
      }
      no-rfc2486;
      pap-password *pap-password*;
      x-auth *ip-address*;
    }
    client-name-filter {
      count *number*;
      domain-name *domain-name*;
      separator *special-character*;
    }
    ldap-options {
      assemble {
        common-name *common-name*;
      }
      base-distinguished-name *base-distinguished-name*;
      revert-interval *seconds*;
      search {
        admin-search {
          distinguished-name *distinguished-name*;
          password *password*;
        }
        search-filter *search-filter-name*;
      }
    }
    ldap-server *server-address* {
      port *port-number*;
      retry *attempts*;
      routing-instance *routing-instance-name*;
      source-address *source-address*;
      timeout *seconds*;
    }
    provisioning-order (gx-plus | jsrc);

```
          service {
            accounting-order {
              activation-protocol;
              radius;
            }
          }
          session-options {
            client-group [group-name];
            client-idle-timeout minutes;
            client-session-timeout minutes;
          }
        }
```

Hierarchy Level	[edit access]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Create a profile containing a set of attributes that define device management access.
Required Privilege Level	access—To view this statement in the configuration. access-control—To add this statement to the configuration.
Related Documentation	• *Understanding Interfaces* • *Understanding User Authentication for Security Devices* • *Ethernet Switching and Layer 2 Transparent Mode Overview*

purging

Syntax	purging;
Hierarchy Level	[edit system arp]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	Purge obsolete ARP entries from the cache when an interface or link goes offline.
Required Privilege Level	system—To view this statement in the configuration. system-control—To add this statement to the configuration.

root-authentication

Syntax	root-authentication {
	    encrypted-password *password*;
	    load-key-file *URL*;
	    plain-text-password;
	    ssh-dsa *public-key* {
	    <from *pattern-list*>;
	    }
	    ssh-rsa *public-key* {
	    <from *pattern-list*>;
	    }
	}

Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 8.5.
Description	Specify authentication information for the root login.

Options

- **encrypted-password** *password*—Specify the encrypted authentication password. You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

- **plain-text-password**—The CLI prompts you for a password encrypts it, and stores the encrypted version in its user database.

- **load-key-file***URL*—File URL containing one or more SSH keys.

- **ssh-dsa** *public-key*—SSH DSA public key string.

  - **from** *pattern-list*—Pattern list of allowed hosts.

- **ssh-rsa** *public-key*—SSH RSA public key string.

  - **from** *pattern-list*—Pattern list of allowed hosts.

Required Privilege Level

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

root-logical-system

Syntax root-logical-system;

Hierarchy Level [edit system security-profile *security-profile-name*]

Release Information Statement introduced in Junos OS Release 11.2.

Description Specify root-logical-system to bind the security profile to the master logical system.

The master administrator uses security profiles to provision logical systems with resources. The security profile containing this statement must be bound to root-logical-system only.

Only the master administrator can create security profiles and bind them to logical systems.

Options none

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

Related Documentation
- Understanding Logical Systems for SRX Series Services Gateways on page 3

scheduler (System Security Profile)

Syntax
```
scheduler {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level   [edit system security-profile *security-profile-name*]

Release Information   Statement introduced in Junos OS Release 11.2.

Description   Specify the number of schedulers that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- maximum *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- reserved *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation   • Understanding Logical Systems for SRX Series Services Gateways on page 3

screen (Security)

Syntax

```
screen {
    ids-option screen-name {
        alarm-without-drop;
        description text;
        icmp {
            flood {
                threshold number;
            }
            fragment;
            icmpv6-malformed;
            ip-sweep {
                threshold number;
            }
            large;
            ping-death;
        }
        ip {
            bad-option;
            block-frag;
            ipv6-extension-header {
                AH-header;
                ESP-header;
                HIP-header;
                }
                destination-header {
                    ILNP-nonce-option;
                    home-address-option;
                    line-identification-option;
                    tunnel-encapsulation-limit-option;
                    user-defined-option-type <type-low> to <type-high>;
                }
                fragment-header;
                hop-by-hop-header {
                    CALIPSO-option;
                    RPL-option;
                    SFM-DPD-option;
                    jumbo-payload-option;
                    quick-start-option;
                    router-alert-option;
                    user-defined-option-type <type-low> to <type-high>;
                }
                mobility-header;
                no-next-header;
                routing-header;
                shim6-header
                user-defined-option-type <type-low> to <type-high>;
            }
            ipv6-extension-header-limit limit;
            ipv6-malformed-header;
            loose-source-route-option;
            record-route-option;
            security-option;
```

```
                    source-route-option;
                    spoofing;
                    stream-option;
                    strict-source-route-option;
                    tear-drop;
                    timestamp-option;
                    unknown-protocol;
                    tunnel {
                      gre {
                        gre-4in4;
                        gre-4in6;
                        gre-6in4;
                        gre-6in6;
                      }
                      ip-in-udp {
                        teredo;
                      }
                      ipip {
                        ipip-4in4;
                        ipip-4in6;
                        ipip-6in4;
                        ipip-6in6;
                        ipip-6over4;
                        ipip-6to4relay;
                        isatap;
                        dslite;
                      }
                      bad-inner-header;
                    }
                  }
                  limit-session {
                    destination-ip-based number;
                    source-ip-based number;
                  }
                  tcp {
                    fin-no-ack;
                    land;
                    port-scan {
                      threshold number;
                    }
                    syn-ack-ack-proxy {
                      threshold number;
                    }
                    syn-fin;
                    syn-flood {
                      alarm-threshold number;
                      attack-threshold number;
                      destination-threshold number;
                      source-threshold number;
                      timeout seconds;
                      white-list name {
                        destination-address destination-address;
                        source-address source-address;
                      }
                    }
                    syn-frag;
```

```
                      tcp-no-flag;
                      tcp-sweep {
                        threshold threshold number;
                      }
                      winnuke;
                    }
                    udp {
                      flood {
                        threshold number;
                      }
                      port-scan {
                        threshold number;
                      }
                      udp-sweep {
                        threshold threshold number;
                      }
                    }
                  }
                }
                traceoptions {
                  file {
                    filename;
                    files number;
                    match regular-expression;
                    size maximum-file-size;
                    (world-readable | no-world-readable);
                  }
                  flag flag;
                  no-remote-trace;
                }
              }
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 8.5. The **description** option added in Junos OS Release 12.1.

Description   Configure security screen options.

Options   *screen-name*—Name of the screen configured at the **security screen ids-options** level.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation   • *Attack Detection and Prevention Overview*

security-profile

Syntax

```
security-profile security-profile-name {
    address-book {
        maximum amount;
        reserved amount;
    }
    appfw-profile {
        maximum amount;
        reserved amount;
    }
    appfw-rule {
        maximum amount;
        reserved amount;
    }
    appfw-rule-set {
        maximum amount;
        reserved amount;
    }
    auth-entry {
        maximum amount;
        reserved amount;
    }
    cpu {
        reserved percent;
    }
    dslite-softwire-initiator {
        maximum amount;
        reserved amount;
    }
    flow-gate {
        maximum amount;
        reserved amount;
    }
    flow-session {
        maximum amount;
        reserved amount;
    }
    idp-policy idp-policy-name;
    logical-system [logical-system-name];
    nat-cone-binding {
        maximum amount;
        reserved amount;
    }
    nat-destination-pool {
        maximum amount;
        reserved amount;
    }
    nat-destination-rule {
        maximum amount;
        reserved amount;
    }
    nat-interface-port-ol {
        maximum amount;
```

```
                                  reserved amount;
                               }
                               nat-nopat-address {
                                  maximum amount;
                                  reserved amount;
                               }
                               nat-pat-address {
                                  maximum amount;
                                  reserved amount;
                               }
                               nat-pat-portnum {
                                  maximum amount
                                  reserved amount
                               }
                               nat-port-ol-ipnumber {
                                  maximum amount;
                                  reserved amount;
                               }
                               nat-rule-referenced-prefix {
                                  maximum amount;
                                  reserved amount;
                               }
                               nat-source-pool {
                                  maximum amount;
                                  reserved amount;
                               }
                               nat-source-rule {
                                  maximum amount;
                                  reserved amount;
                               }
                               nat-static-rule {
                                  maximum amount;
                                  reserved amount;
                               }
                               policy {
                                  maximum amount;
                                  reserved amount;
                               }
                               policy-with-count {
                                  maximum amount;
                                  reserved amount;
                               }
                               root-logical-system;
                               scheduler {
                                  maximum amount;
                                  reserved amount;
                               }
                               zone {
                                  maximum amount;
                                  reserved amount;
                               }
```

Hierarchy Level   [edit system]

Release Information   Statement introduced in Junos OS Release 11.2. The **dslite-softwire-initiator** option
introduced in Junos OS Release 12.1.

Description   Create a security profile and specify the kinds and amounts of resources to allocate to
a logical system to which the security profile is bound.

As a master administrator, you can create a security profile and bind it to more than one
logical system if you want to allocate the same kinds and amounts of resources to them.
For details on how many security profiles you can create, see "Understanding Logical
System Security Profiles (Master Administrators Only)" on page 51. When you reach the
limit, you must delete a security profile and commit the configuration before you can
create and commit the configuration for another security profile.

Only the master administrator can create security profiles.

Options   • *security-profile-name*—Name of the security profile.

The remaining statements are explained separately. See CLI Explorer.

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

Related   • Understanding Logical System Security Profiles (Master Administrators Only) on
Documentation   page 51

• Example: Configuring Logical Systems Security Profiles (Master Administrators Only)
on page 56

security-profile-resources

Syntax	security-profile-resources {
	   cpu-control;
	   cpu-control-target *percent*;
	}

Hierarchy Level	[edit system]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Configure global settings that apply to all logical systems in the device.
Options	**cpu-control**—Enable CPU utilization control.

cpu-control-target *percent*—Specify the upper limit for CPU utilization on the device under normal operating conditions.

Range: 0 through 100 percent (decimal point allowed).

Default: 80 percent.

> *i*   NOTE: The **cpu-control** option must be specified for the **cpu-control-target** value to take effect.

Required Privilege Level

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

Related Documentation

- Understanding Logical System Security Profiles (Master Administrators Only) on page 51
- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 56

softwires

Syntax

```
softwires {
    softwire-name name {
        softwire-concentrator ipv6-address;
        softwire-type IPv4-in-IPv6;
    }
    traceoptions {
        file {
            filename;
            files number;
            match regular-expression;
            size maximum-file-size;
            (world-readable | no-world-readable);
        }
        flag (all | configuration | flow);
        no-remote-trace;
    }
}
```

Hierarchy Level   [edit security]

Release Information   Statement introduced before Junos OS Release 12.1.

Description   Configure softwires for IPv6 dual-stack lite (DS-Lite). DS-Lite allows migration to an IPv6 access network without changing end-user software. IPv4 users can continue to access IPv4 internet content using their current hardware, while IPv6 users are able to access IPv6 content.

Options   • **softwire-name** *name*—Name of the softwire configuration.

    • **softwire-concentrator** *ipv6-address*—IPv6 address of the concentrator.

    • **softwire-type**—Must be IPv4-in-IPv6.

The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation   • [edit security softwires] Hierarchy Level

zone (System Security Profile)

Syntax

```
zone {
    maximum amount;
    reserved amount;
}
```

Hierarchy Level   [edit system security-profile *security-profile-name*]

Release Information   Statement introduced in Junos OS Release 11.2.

Description   Specify the zones that user logical system administrators and master logical system administrators can configure for their logical systems if the security profile is bound to the logical systems.

The master administrator:

- uses security profiles to provision logical systems with resources.

- binds security profiles to user logical systems and the master logical system.

- can configure more than one security profile, specifying different amounts of resource allocations in various profiles.

Only the master administrator can create security profiles and bind them to logical systems.

Options
- **maximum** *amount*—A maximum allowed quota. If a logical system requires more of a resource than its reserved amount allows, it can utilize resources configured for the global maximum amount if they are available—that is, if they are not allocated to other logical systems. The maximum allowed quota specifies the portion of the free global resources that the logical system can use. The maximum allowed quota does not guarantee that the amount specified for the resource in the security profile is available. Logical systems compete for global resources.

- **reserved** *amount*—A reserved quota that guarantees that the resource amount specified is always available to the logical system.

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

Related Documentation
- Understanding Logical System Security Profiles (Master Administrators Only) on page 51

- Example: Configuring Logical Systems Security Profiles (Master Administrators Only) on page 56

zones

```
Syntax    zones {
              functional-zone {
                management {
                  description text;
                  host-inbound-traffic {
                    protocols protocol-name {
                      except;
                    }
                    system-services service-name {
                      except;
                    }
                  }
                  interfaces interface-name  {
                    host-inbound-traffic {
                      protocols protocol-name {
                        except;
                      }
                      system-services service-name {
                        except;
                      }
                    }
                  }
                  screen screen-name;
                }
              }
              security-zone zone-name {
                address-book {
                  address address-name {
                    ip-prefix {
                      description text;
                    }
                    description text;
                    dns-name domain-name {
                      ipv4-only;
                      ipv6-only;
                    }
                    range-address lower-limit to upper-limit;
                    wildcard-address ipv4-address/wildcard-mask;
                  }
                  address-set address-set-name {
                    address address-name;
                    address-set address-set-name;
                    description text;
                  }
                }
                advance-policy-based-routing;
                application-tracking;
                description text;
                host-inbound-traffic {
                  protocols protocol-name {
                    except;
                  }
```

```
            system-services service-name {
              except;
            }
          }
          interfaces interface-name {
            host-inbound-traffic {
              protocols protocol-name {
                except;
              }
              system-services service-name {
                except;
              }
            }
          }
          screen screen-name;
          tcp-rst;
        }
      }
```

Hierarchy Level   [edit security]

Release Information   Statement introduced in Junos OS Release 8.5. Support for wildcard addresses added in Junos OS Release 11.1. The **description** option added in Junos OS Release 12.1.

Description   A zone is a collection of interfaces for security purposes. All interfaces in a zone are equivalent from a security point of view. Configure the following zones:

- Functional zone—Special-purpose zone, such as a management zone that can host dedicated management interfaces.

- Security zone—Most common type of zone that is used as a building block in policies.

Options   The remaining statements are explained separately. See CLI Explorer.

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

Related Documentation
- *Security Zones and Interfaces Overview*

- *Supported System Services for Host Inbound Traffic*

Operational Commands

- clear security application-firewall rule-set statistics logical-system
- clear security dns-cache
- request security datapath-debug capture start
- request security datapath-debug capture stop
- set chassis cluster cluster-id node reboot
- show chassis cluster status
- show log
- show security application-firewall rule-set
- show security application-firewall rule-set logical-system
- show security application-tracking counters
- show security datapath-debug capture
- show security datapath-debug counter
- show security dns-cache
- show security firewall-authentication history
- show security firewall-authentication users
- show security flow session
- show security idp logical-system policy-association
- show security ike security-associations
- show security ipsec security-associations
- show security match-policies
- show security nat destination rule
- show security nat destination summary
- show security nat source rule
- show security nat source summary
- show security nat static rule
- show security policies
- show security screen statistics
- show system security-profile

- show security softwires
- show security zones

clear security application-firewall rule-set statistics logical-system

Syntax   The master, or root, administrator can issue the following statements:

clear security application-firewall rule-set statistics [logical-system *logical-system-name* |
all | root-logical-system]

The user logical system administrator can issue the following statement:

clear security application-firewall rule-set statistics all

Release Information   Command introduced in Junos OS Release 11.4.

Description   Clear all security application firewall rule set statistics.

> NOTE: User logical system administrators can clear statistics only for the
> logical systems they can access. For information about master and user
> administrator roles in logical systems, see "Understanding the Master Logical
> System and the Master Administrator Role" on page 19.

Options   *logical-system-name*—Name of a specific logical system.

all—(default) Clear all rule set statistics for a specific logical system or all logical systems.

root-logical-system—Clear application firewall rule set statistics on the root logical system
(master administrator only).

Required Privilege
Level   clear

Related
Documentation
- show security application-firewall rule-set logical-system on page 380

Output Fields   This command produces no output.

clear security dns-cache

Syntax	clear security dns-cache <dns-name *dns-name*>
Release Information	Command introduced in Junos OS Release 12.1X44-D10.
Description	Reset DNS cache information.

> **NOTE:** This command is only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options	• **dns-name**—Clear DNS cache information for the specified name.
Required Privilege Level	clear
Related Documentation	• show security dns-cache on page 386
	• Understanding the Master Logical System and the Master Administrator Role on page 19

request security datapath-debug capture start

Syntax   request security datapath-debug capture start

Release Information   Command introduced in Junos OS Release 10.0.

Description   Start the data path debugging capture.

Required Privilege   maintenance
Level

Related   • Understanding Data Path Debugging for Logical Systems on page 274
Documentation

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

Sample Output

request security datapath-debug capture start

```
user@host> request security datapath-debug capture start
datapath-debug capture started on file
```

request security datapath-debug capture stop

Syntax	request security datapath-debug capture stop
Release Information	Command introduced in Junos OS Release 10.0.
Description	Stop the data path debugging capture.
Required Privilege Level	maintenance
Related Documentation	• Understanding Data Path Debugging for Logical Systems on page 274
Output Fields	When you enter this command, you are provided feedback on the status of your request.

Sample Output

request security datapath-debug capture stop

```
user@host> request security datapath-debug capture stop
datapath-debug capture successfully stopped, use show security datapath-debug
capture to view
```

set chassis cluster cluster-id node reboot

Syntax	set chassis cluster cluster-id *cluster-id* node *node* reboot
Release Information	Support for extended cluster identifiers (more than 15 identifiers) added in Junos OS Release 12.1X45-D10.
Description	This operational mode command sets the chassis cluster identifier (ID) and node ID on each device, and reboots the devices to enable clustering. The system uses the chassis cluster ID and chassis cluster node ID to apply the correct configuration for each node (for example, when you use the **apply-groups** command to configure the chassis cluster management interface). The chassis cluster ID and node ID statements are written to the EPROM, and the statements take effect when the system is rebooted.

Setting a cluster ID to 0 is equivalent to disabling a cluster. A cluster ID greater than 15 can only be set when the fabric and control link interfaces are connected back-to-back.

> *i* **NOTE:** If you have a cluster set up and running with an earlier release of Junos OS, you can upgrade to Junos OS Release 12.1X45-D10 or later and re-create a cluster with cluster IDs greater than 16. If for any reason you decide to revert to the previous version of Junos OS that did not support extended cluster IDs, the system comes up with standalone devices after you reboot. If the cluster ID set is less than 16 and you roll back to a previous release, the system comes back with the previous setup.

Options	cluster-id *cluster-id* —Identifies the cluster within the Layer 2 domain.
	Range: 0 through 255
	node *node* —Identifies a node within a cluster.
	Range: 0 to 1
Required Privilege Level	maintenance
Related Documentation	• *Example: Setting the Chassis Cluster Node ID and Cluster ID for Branch SRX Series Devices*
	• *Example: Setting the Chassis Cluster Node ID and Cluster ID for High-End SRX Series Devices*
	• Understanding the Interconnect Logical System and Logical Tunnel Interfaces on page 8
	• Example: Configuring Logical Systems in an Active/Passive Chassis Cluster (Master Administrators Only) on page 168
Output Fields	When you enter this command, you are provided feedback on the status of your request.

show chassis cluster status

Syntax	show chassis cluster status <redundancy-group *group-number* >
Release Information	Command modified in Junos OS Release 9.2. Support for dual control ports added in Junos OS Release 10.0. Support for monitoring failures added in Junos OS Release 12.1X47-D10.
Description	Display the failover status of a chassis cluster.
Options	• none—Display the status of all redundancy groups in the chassis cluster. • redundancy-group *group-number* —(Optional) Display the status of the specified redundancy group.
Required Privilege Level	view
Related Documentation	• *redundancy-group (Chassis Cluster)* • *clear chassis cluster failover-count* • *request chassis cluster failover node* • *request chassis cluster failover reset*
List of Sample Output	show chassis cluster status on page 371 show chassis cluster status redundancy-group 1 on page 372
Output Fields	Table 24 on page 370 lists the output fields for the **show chassis cluster status** command. Output fields are listed in the approximate order in which they appear.

Table 24: show chassis cluster status Output Fields

Field Name	Field Description
Cluster ID	ID number (1-15) of a cluster is applicable for releases upto 12.1X45-D10. ID number (1-255) is applicable for releases 12.1X45-D10 and later. Setting a cluster ID to 0 is equivalent to disabling a cluster.
Redundancy-Group	ID number (1-128) of a redundancy group in the chassis cluster.
Node name	Node (device) in the chassis cluster (**node0** or **node1**).
Priority	Assigned priority for the redundancy group on that node.

Table 24: show chassis cluster status Output Fields *(continued)*

Field Name	Field Description
Status	State of the redundancy group (**Primary**, **Secondary**, **Lost**, or **Unavailable**). • **Primary**—Redundancy group is active and passing traffic. • **Secondary**—Redundancy group is passive and not passing traffic. • **Lost**—Node loses contact with the other node through the control link. Most likely to occur when both nodes are in a cluster and due to control link failure, one node cannot exchange heartbeats, or when the other node is rebooted. • **Unavailable**—Node has not received a single heartbeat over the control link from the other node since the other node booted up. Most likely to occur when one node boots up before the other node, or if only one node is present in the cluster.
Preempt	• **Yes**: Mastership can be preempted based on priority. • **No**: Change in priority will not preempt the mastership.
Manual failover	• **Yes**: If the Mastership is set manually through the CLI with the **request chassis cluster failover node** or **request chassis cluster failover redundancy-group** command. This overrides **Priority** and **Preempt**. • **No**: Mastership is not set manually through the CLI.
Monitor-failures	• **None**: Cluster working properly. • **Monitor Failure code**: Cluster is not working properly and the respective failure code is displayed.

Sample Output

Displays chassis cluster status with all redundancy groups.

show chassis cluster status

```
user@host> show chassis cluster status

Monitor Failure codes:
  CS  Cold Sync monitoring        FL  Fabric Connection monitoring
   GR  GRES monitoring            HW  Hardware monitoring
   IF  Interface monitoring       IP  IP monitoring
   LB  Loopback monitoring        MB  Mbuf monitoring
   NH  Nexthop monitoring         NP  NPC monitoring
   SP  SPU monitoring             SM  Schedule monitoring
   CF  Config Sync monitoring

Cluster ID: 1
Node    Priority Status         Preempt Manual    Monitor-failures

Redundancy group: 0 , Failover count: 1
node0   200      primary        no      no        None
node1   1        secondary      no      no        None

Redundancy group: 1 , Failover count: 1
node0   101      primary        no      no        None
node1   1        secondary      no      no        None
```

Sample Output

Displays chassis cluster status with redundancy group 1 only.

show chassis cluster status redundancy-group 1

```
user@host> show chassis cluster status redundancy-group 1

Monitor Failure codes:
    CS  Cold Sync monitoring        FL  Fabric Connection monitoring
    GR  GRES monitoring             HW  Hardware monitoring
    IF  Interface monitoring        IP  IP monitoring
    LB  Loopback monitoring         MB  Mbuf monitoring
    NH  Nexthop monitoring          NP  NPC monitoring
    SP  SPU monitoring              SM  Schedule monitoring
    CF  Config Sync monitoring

Cluster ID: 1
Node    Priority Status          Preempt Manual    Monitor-failures

Redundancy group: 1 , Failover count: 1
node0   101      primary         no      no        None
node1   1        secondary       no      no        None
```

show log

List of Syntax	
Syntax	show log *<filename	user <username>>*
Syntax (QFX Series and OCX Series)	show log *filename* *<device-type (device-id | device-alias)>*
Syntax (TX Matrix Router)	show log <all-lcc	lcc *number*	scc> *<filename	user <username>>*
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series. Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series. Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Description	List log files, display log file contents, or display information about users who have logged in to the router or switch.

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

Options	**none**—List all log files.
	**<all-lcc	lcc *number*	scc>**—(Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace *number* with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).
	device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.

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

- **interconnect-device**—Display logs for Interconnect devices.

- **node-device**—Display logs for Node devices.

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

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

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

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

Required Privilege Level   trace

List of Sample Output

Sample Output

show log

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
```

show log filename

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

show log filename (QFabric System)

```
user@qfabric> show log messages
Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC:  @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default___NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default___NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492
file: UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
 chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)
```

show log user

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

show security application-firewall rule-set

Syntax	show security application-firewall rule-set (*<rule-set-name>* | all)
Release Information	Command introduced in Junos OS Release 11.1. Updated in Junos OS Release 12.1X44-D10 with output format changes. Updated in Junos OS Release 12.1X45-D10 with redirection counters.
Description	Display information about the specified rule set defined in the application firewall.
Options	*rule-set-name*—Name of the rule set.
	all—Display information about all the application firewall rule sets.
Required Privilege Level	view
Related Documentation	• *clear security application-firewall rule-set statistics*
List of Sample Output	show security application-firewall rule-set my_ruleset1 on page 378 show security application-firewall rule-set all on page 378
Output Fields	Table 25 on page 377 lists the output fields for the **show security application-firewall rule-set** command. Output fields are listed in the approximate order in which they appear.

Table 25: show security application-firewall rule-set Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system of the rule set.
Profile	The redirect profile to be used for rules requiring redirection for reject or deny actions.

Rule	Name of the rule
	
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **SSL-Encryption**—Setting for SSL traffic.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  ◦ **permit**
	  ◦ **deny**
	  ◦ **reject**
	  ◦ **redirect**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.
	• **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.

Table 25: show security application-firewall rule-set Output Fields *(continued)*

Field Name	Field Description
Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule. • **Number of sessions redirected**—Number of sessions redirected by the application firewall rule.
Number of sessions with appid pending	Number of sessions that are pending application identification processing

Sample Output

show security application-firewall rule-set my_ruleset1

```
user@host>show security application-firewall rule-set my_ruleset1
Rule-set: my_ruleset1
    Rule: rule1
        Dynamic Applications: junos:facebook, junos:messenger
        Dynamic Application Groups: junos:web, junos:chat
        SSL-Encryption: any
        Action: deny or redirect
        Number of sessions matched: 10
        Number of sessions redirected: 10
    Default rule: permit
        Number of sessions matched: 200
        Number of sessions redirected: 0
    Number of sessions with appid pending: 2
```

Sample Output

show security application-firewall rule-set all

```
user@host> show security application-firewall rule-set all
Rule-set: appfw
    Logical system: root-logical-system
    Profile: lsy2_pf555
    Rule: 2
        Dynamic Applications: junos:HTTP
        SSL-Encryption: any
        Action:deny or redirect
        Number of sessions matched: 2
        Number of sessions redirected: 2
    Rule: 1
        Dynamic Applications: junos:FTP
        SSL-Encryption: any
        Action:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
    Default rule:permit
        Number of sessions matched: 0
        Number of sessions redirected: 0
    Number of sessions with appid pending: 0
```

show security application-firewall rule-set logical-system

Syntax
The master, or root, administrator can issue the following statements:

show security application-firewall rule-set all
show security application-firewall rule-set *rule-set-name* | all | logical-system
logical-system-name | all | root-logical-system [*logical-system-name* | all]

The user logical system administrator can issue the following statement:

show security application-firewall rule-set all

Release Information
Command introduced in Junos OS Release 11.4.

Description
Display information about application firewall rule set(s) associated with a specific logical system, all logical systems, or the root logical system configured on a device.

> *i* NOTE: The master administrator can configure and view application firewall rule sets for the root logical system and all user logical systems configured on the device. User logical system administrators can configure and view application firewall rule set information only for the user logical systems for which they have access. For information about master and user administrator roles in logical systems, see "Understanding Logical Systems for SRX Series Services Gateways" on page 3.

Options
rule-set-name—Name of a specific rule set.

logical-system-name—Name of a specific logical system.

all—(default) Display all rule sets for all logical systems. The user logical system administrator can display all rule sets only for the logical system they can access.

root-logical-system—Display application firewall rule set information for the root logical system (master administrator only).

Required Privilege
Level
view

Related
Documentation
• clear security application-firewall rule-set statistics logical-system on page 365

List of Sample Output
show security application-firewall rule-set logical-system all on page 381
show security application-firewall rule-set all on page 382

Output Fields
Table 26 on page 381 lists the output fields for the **show security application-firewall rule-set logical-system** command. Output fields are listed in the approximate order in which they appear.

Table 26: show security application-firewall rule-set logical-system Output Fields

Field Name	Field Description
Rule-set	Name of the rule set.
Logical system	Name of the logical system.

Rule	Name of the rule.
	
	• **Dynamic applications**—Name of the applications.
	• **Dynamic application groups**—Name of the application groups.
	• **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	  ◦ **permit**
	  ◦ **deny**
	• **Number of sessions matched**—Number of sessions matched with the application firewall rule.

Default rule	The default rule applied when the identified application is not specified in any rules of the rule set. • **Number of sessions matched**—Number of sessions matched with the application firewall default rule.
Number of sessions with appid pending	Number of sessions that are pending with the application ID processing.

Sample Output

show security application-firewall rule-set logical-system all

```
root@host> show security application-firewall rule-set logical-system all

Rule-set: root_rs1
    Logical system: root-logical-system
    Rule: r1
        Dynamic Applications: junos:FTP
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 4

Rule-set: root-rs2
    Logical system: root-logical-system
    Rule: r1
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 10
```

show security application-firewall rule-set all

```
root@host> show security application-firewall rule-set all

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:TELNET
        Action:permit
        Number of sessions matched: 10
Default rule:deny
        Number of sessions matched: 100
Number of sessions with appid pending: 2

Rule-set: ls-product-design-rs1
    Logical system: ls-product-design
    Rule: r2
        Dynamic Application Groups: junos:web
        Action:permit
        Number of sessions matched: 20
Default rule:deny
        Number of sessions matched: 200
Number of sessions with appid pending: 4

Rule-set: ls-product-design-rs2
    Logical system: ls-product-design
    Rule: r1
        Dynamic Applications: junos:FACEBOOK
        Action:deny
        Number of sessions matched: 40
Default rule:permit
        Number of sessions matched: 400
Number of sessions with appid pending: 10
```

show security application-tracking counters

Syntax	show security application-tracking counters
Release Information	Command introduced in Junos OS Release 10.2.
Description	Display the status of AppTrack counters.
Required Privilege Level	view
Related Documentation	• *Example: Configuring AppTrack*
Output Fields	Table 27 on page 383 lists the output fields for the **show security application-tracking counters** command. Output fields are listed in the approximate order in which they appear.

Table 27: show security application-tracking counters

Field Name	Field Description
Session create messages	The number of log messages generated when a session was created.
Session close messages	The number of log messages generated when a session was closed.
Session volume updates	The number of log messages generated when an update interval was exceeded.
Failed messages	The number of messages that were not generated due to memory or session constraints.

Sample Output

show security application-tracking counters

```
user@host> show security application-tracking counters
AVT counters:                    Value
 Session create messages          0
 Session close messages           0
 Session volume updates           0
 Failed messages                  0
```

show security datapath-debug capture

Syntax	show security datapath-debug capture
Release Information	Command introduced in Junos OS Release 10.0.
Description	Display details of the data path debugging capture file.
Required Privilege Level	view
Related Documentation	• show security datapath-debug counter on page 385
	• Understanding Data Path Debugging for Logical Systems
List of Sample Output	show security datapath—debug capture on page 384
Output Fields	Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath—debug capture

```
user@host> show security datapath-debug capture
Packet 1, len 120: (C0/F0/P0/SEQ:71:lbt)
91 00 00 47 11 00 10 00 9a 14 00 19 03 00 00 00
00 00 00 00 00 01 00 47 10 00 00 00 00 00 00 00
00 1f 12 f8 dd 29 00 21 59 84 f4 01 81 00 02 1e
08 00 45 60 01 f4 00 00 00 00 3f 06 73 9f 01 01
01 02 03 01 01 02 d4 31 d4 31 00 00 00 00 00 00
00 00 50 02 00 00 ff ad 00 00 00 00
Packet 2, len 120: (C0/F0/P0/SEQ:71:lbt)
90 00 00 47 04 00 00 00 00 00 00 00 02 02 00 47
10 00 00 00 00 00 00 00 50 00 a6 1c 00 00 00 00
00 00 00 0a 00 00 00 00 00 00 09 d9 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 1f 12 f8
dd 29 00 21 59 84 f4 01 81 00 02 1e
```

show security datapath-debug counter

Syntax	show security datapath-debug counter
Release Information	Command introduced in Junos OS Release 10.0.
Description	Display details of the data path debugging counter.
Required Privilege Level	view
Related Documentation	• show security datapath-debug capture
	• Understanding Data Path Debugging for Logical Systems
List of Sample Output	show security datapath-debug counter on page 385
Output Fields	Output fields are listed in the approximate order in which they appear.

Sample Output

show security datapath-debug counter

```
user@host> show security datapath-debug counter
Datapath debug counters
Packet Filter 1:
np-ingress
Chassis 0 FPC 4 : 1
np-ingress
Chassis 0 FPC 3 : 0
np-egress
Chassis 0 FPC 4 : 1
np-egress
Chassis 0 FPC 3 : 0
jexec
Chassis 0 FPC 0 PIC 1: 0
jexec
Chassis 0 FPC 0 PIC 0: 1
lbt
Chassis 0 FPC 0 PIC 1: 0
lbt
Chassis 0 FPC 0 PIC 0: 2
pot
Chassis 0 FPC 0 PIC 1: 0
pot
```

show security dns-cache

Syntax	show security dns-cache <dns-name *dns-name*>
Release Information	Command introduced in Junos OS Release 12.1X44-D10.
Description	Display DNS cache information.

> **(i) NOTE:** This command is only available to the master administrator on devices that are configured for logical systems. This command is not available in user logical systems or on devices that are not configured for logical systems.

Options	• **dns-name**—Display DNS cache information for the specified name.
Required Privilege Level	view
Related Documentation	•
List of Sample Output	
Output Fields	lists the output fields for the **show security dns-cache** command. Output fields are listed in the approximate order in which they appear.

Table 28: show security dns-cache Output Fields

Field Name	Field Description
DNS Name	DNS name.
Address Family	IPv4 or IPv6.
TTL	Time-to-live value.
IP Address	IP address for the DNS name.

Sample Output

show security dns-cache

```
user@host> show security dns-cache
DNS Name: dns1.test.com:
  Address Family: IPv4, TTL: 10
    IP Address: 1.1.1.1
  Address Family: IPv6: TTL = 15
    IP Address: 2001:1.1.1.1
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
```

```
                                    IP Address: 2.2.2.2
                                    IP Address: 2.2.2.3
```

Sample Output

show security dns-cache dns-name dns2.test.com

```
user@host> show security dns-cache dns-name dns2.test.com
DNS Name: dns2.test.com:
  Address Family: IPv4, TTL: 20
      IP Address: 2.2.2.2
      IP Address: 2.2.2.3
```

show security firewall-authentication history

Syntax	show security firewall-authentication history <node (*node-id* | all | local | primary)>
Release Information	Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.
Description	Display security firewall authentication history information.

Options	• none—Display history of firewall authentication information.
	
	• **node**—(Optional) For chassis cluster configurations, display all firewall authentication history on a specific node (device) in the cluster.
	
	  • *node-id* —Identification number of the node. It can be 0 or 1.
	
	  • **all**—Display information about all nodes.
	
	  • **local**—Display information about the local node.
	
	  • **primary**—Display information about the primary node.

Required Privilege Level	view
Related Documentation	• Understanding Logical System Firewall Authentication on page 65 • *Firewall User Authentication Overview*
List of Sample Output	show security firewall-authentication history on page 389 show security firewall-authentication history node all on page 389
Output Fields	Table 29 on page 388 lists the output fields for the **show security firewall-authentication history** command. Output fields are listed in the approximate order in which they appear.

Table 29: show security firewall-authentication history Output Fields

Field Name	Field Description
Authentications	Number of authentications.
Id	Identification number.
Source IP	IP address of the authentication source.
Date	Authentication date.
Time	Authentication time.
Duration	Authentication duration.
Status	Authentication status success or failure.

Table 29: show security firewall-authentication history Output Fields *(continued)*

Field Name	Field Description
User	Name of the user.

Sample Output

show security firewall-authentication history

```
user@host> show security firewall-authentication history
History of firewall authentication data:
  Authentications: 1
         Id Source Ip        Date       Time     Duration  Status    User
          1 203.0.113.1      2007-04-03 11:43:06 00:00:45  Success   hello
```

Sample Output

show security firewall-authentication history node all

```
user@host> show security firewall-authentication history node all
node0:
--------------------------------------------------------------------------
History of firewall authentication data:
Authentications: 2
Id Source Ip        Date       Time     Duration   Status    User
1 203.0.113.1      2008-01-04 12:00:10 0:05:49     Success   local1
2 203.0.113.1      2008-01-04 14:36:52 0:01:03     Success   local1
node1:
--------------------------------------------------------------------------
History of firewall authentication data:
  Authentications: 1
         Id Source Ip        Date       Time     Duration   Status    User
            203.0.113.1      2008-01-04 14:59:43 1193046:06: Success local1
```

show security firewall-authentication users

Syntax	show security firewall-authentication users <node (*node-id* | all | local | primary) >
Release Information	Command introduced in Junos OS Release 8.5. The **node** options added in Junos OS Release 9.0.
Description	Display firewall authentication details about all users.

Options	• none—Display details about all firewall authentication users.
	
	• **node**—(Optional) For chassis cluster configurations, display firewall authentication details for all users on a specific node.
	
	  ◦ *node-id*—Identification number of the node. It can be 0 or 1.
	
	  ◦ **all**—Display information about all nodes.
	
	  ◦ **local**—Display information about the local node.
	
	  ◦ **primary**—Display information about the primary node.

Required Privilege Level	view
Related Documentation	• *Firewall User Authentication Overview*
List of Sample Output	show security firewall-authentication users on page 391 show security firewall-authentication users node 0 on page 391 show security firewall-authentication users node all on page 391
Output Fields	Table 30 on page 390 lists the output fields for the **show security firewall-authentication users** command. Output fields are listed in the approximate order in which they appear.

Table 30: show security firewall-authentication users Output Fields

Field Name	Field Description
Total users in table	Gives count of how many entries/users the command will display.
Id	Identification number.
Source IP	IP address of the authentication source.
Src zone	User traffic received from the zone.
Dst zone	User traffic destined to the zone.
Profile	Name of profile used for authentication.
Age	Idle timeout for the user.

Table 30: show security firewall-authentication users Output Fields *(continued)*

Field Name	Field Description
Status	Authentication status success or failure.
User	Name of the user.

Sample Output

show security firewall-authentication users

```
user@host> show security firewall-authentication users
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           1 192.0.2.5/24         z1       z2        p1      0 Success  local1
```

Sample Output

show security firewall-authentication users node 0

```
user@host> show security firewall-authentication users node 0
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User

           3 192.0.2.5/24        z1       z2        p1        1 Success  local1
```

Sample Output

show security firewall-authentication users node all

```
user@host> show security firewall-authentication users node all
node0:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           3 192.0.2.5           z1       z2        p1        1 Success  local1

node1:
--------------------------------------------------------------------------
Firewall authentication data:
  Total users in table: 1
          Id Source Ip       Src zone Dst zone Profile    Age Status   User
           2 192.0.2.5           z1       z2        p1        1 Success  local1
```

show security flow session

Syntax	show security flow session [*filter*] [brief | extensive | summary]
Release Information	Command introduced in Junos OS Release 8.5. Support for filter and view options added in Junos OS Release 10.2. Application firewall, dynamic application, and logical system filters added in Junos OS Release 11.2. Policy ID filter added in Junos OS Release 12.3X48-D10. Support for connection tag added in Junos OS Release 15.1X49-D40.
Description	Display information about all currently active security sessions on the device.

> *i* NOTE: For the normal flow sessions, the show security flow session command displays bytes counters based on IP header length. However for sessions in Express Path mode, the statistics is collected from IOC2 and IOC3 ASIC hardware engine, and includes full packet length with L2 headers. Because of this, the output displays slightly larger bytes counters for sessions in Express Path mode than the normal flow session.

Options

- *filter*—Filter the display by the specified criteria.

  The following filters reduce the display to those sessions that match the criteria specified by the filter. Refer to the specific **show** command for examples of the filtered output.

  application—Predefined application name

  application-firewall—Application firewall enabled

  application-firewall-rule-set—Application firewall enabled with the specified rule set

  application-traffic-control—Application traffic control session

  application-traffic-control-rule-set—Application traffic control rule set name and rule name

  conn-tag—Session connection tag

  destination-port—Destination port

  destination-prefix—Destination IP prefix or address

  dynamic-application—Dynamic application

  dynamic-application-group—Dynamic application

  encrypted—Encrypted traffic

  extensive—Display detailed output

family—Display session by family

idp—IDP enabled sessions

interface—Name of incoming or outgoing interface

logical-system (all | *logical-system-name*)—Name of a specific logical system or all to
display all logical systems

nat—Display sessions with network address translation

policy-id—Display session information based on policy ID; the range is 1 through
4,294,967,295

protocol—IP protocol number

resource-manager—Resource manager

root-logical-system—Display root logical system as default

security-intelligence—Display security intelligence sessions

services-offload—Display services offload sessions

session-identifier—Display session with specified session identifier

source-port—Source port

source-prefix—Source IP prefix

summary—Display output summary

tunnel—Tunnel sessions

- brief | extensive | summary—Display the specified level of output.

- none—Display information about all active sessions.

Required Privilege Level	view

Related Documentation

- *Juniper Networks Devices Processing Overview*

- *clear security flow session all*

List of Sample Output

Output Fields

lists the output fields for the **show security flow session** command.
Output fields are listed in the approximate order in which they appear.

Table 31: show security flow session Output Fields

Field Name	Field Description
Session ID	Number that identifies the session. Use this ID to get more information about the session.
Conn Tag	A 32-bit connection tag that uniquely identifies the GPRS tunneling protocol, user plane (GTP-U) and the Stream Control Transmission Protocol (STCP) sessions. The connection tag for GTP-U is the tunnel endpoint identifier (TEID) and for SCTP is the vTag. The connection ID remains 0 if the connection tag is not used by the sessions.
CP Session ID	Number that identifies the central point session. Use this ID to get more information about the central point session.
Policy name	Policy that permitted the traffic.
Timeout	Idle timeout after which the session expires.
In	Incoming flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Out	Reverse flow (source and destination IP addresses, application protocol, interface, session token, route, gateway, tunnel, port sequence, FIN sequence, FIN state, packets and bytes).
Total sessions	Total number of sessions.
Status	Session status.
Flag	Internal flag depicting the state of the session, used for debugging purposes. The three available flags are: • flag • natflag • natflag2
Policy name	Name and ID of the policy that the first packet of the session matched.
Source NAT pool	The name of the source pool where NAT is used.
Dynamic application	Name of the application.
Application traffic control rule-set	AppQoS rule set for this session.
Rule	AppQoS rule for this session.
Forwarding class	The AppQoS forwarding class name for this session that distinguishes the transmission priority
DSCP code point	Differentiated Services (DiffServ) code point (DSCP) value remarked by the matching rule for this session.

Table 31: show security flow session Output Fields *(continued)*

Field Name	Field Description
Loss priority	One of four priority levels set by the matching rule to control discarding a packet during periods of congestion. A high loss priority means a high probability that the packet could be dropped during a period of congestion.
Rate limiter client to server	The rate-limiter profile assigned to the client-to-server traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Rate limiter server to client	The rate-limiter profile assigned to the server-to-client traffic defining a unique combination of **bandwidth-limit** and **burst-size-limit** specifications.
Maximum timeout	Maximum session timeout.
Current timeout	Remaining time for the session unless traffic exists in the session.
Session State	Session state.
Start time	Time when the session was created, offset from the system start time.
Unicast-sessions	Number of unicast sessions.
Multicast-sessions	Number of multicast sessions.
Failed-sessions	Number of failed sessions.
Sessions-in-use	Number of sessions in use. • Valid sessions • Pending sessions • Invalidated sessions • Sessions in other states
Maximum-sessions	Maximum number of sessions permitted.

Sample Output

show security flow session

```
root> show security flow session
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 56, Valid
  In: 203.0.113.1/1000 --> 203.0.113.11/2000;udp, Conn Tag: 0x0, If: reth1.0,
Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.11/2000 --> 203.0.113.1/1000;udp, Conn Tag: 0x0, If: reth0.0,
Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session brief

```
root> show security flow session brief
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Policy name: SG/4, State: Active, Timeout: 62, Valid
  In: 203.0.113.11/1000 --> 203.0.113.1/2000;udp, Conn Tag: 0x0, If: reth1.0,
Pkts: 1, Bytes: 86, CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp, Conn Tag: 0x0, If: reth0.0,
Pkts: 0, Bytes: 0, CP Session ID: 10320276

Total sessions: 1
```

show security flow session extensive

```
root> show security flow session extensive
Flow Sessions on FPC0 PIC1:

Session ID: 10115977, Status: Normal, State: Active
Flags: 0x8000040/0x18000000/0x12000003
Policy name: SG/4
Source NAT pool: Null, Application: junos-gprs-gtp-v0-udp/76
Dynamic application: junos:UNKNOWN,
Encryption:  Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 90, Current timeout: 54
Session State: Valid
Start time: 6704, Duration: 35
  In: 203.0.113.11/1000 --> 201.11.0.100/2000;udp,
  Conn Tag: 0x0, Interface: reth1.0,
    Session token: 0x6, Flag: 0x40000021
    Route: 0x86053c2, Gateway: 201.10.0.100, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 86
    CP Session ID: 10320276
  Out: 203.0.113.1/2000 --> 203.0.113.11/1000;udp,
  Conn Tag: 0x0, Interface: reth0.0,
    Session token: 0x7, Flag: 0x50000000
    Route: 0x86143c2, Gateway: 203.0.113.11, Tunnel: 0
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
    CP Session ID: 10320276
Total sessions: 1
```

show security flow session summary

```
root> show security flow session summary
Flow Sessions on FPC10 PIC1:
Unicast-sessions: 1
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 1
  Valid sessions: 1
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

```
Flow Sessions on FPC10 PIC2:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456

Flow Sessions on FPC10 PIC3:
Unicast-sessions: 0
Multicast-sessions: 0
Services-offload-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 6291456
```

show security idp logical-system policy-association

Syntax	show security idp logical-system policy-association
Release Information	Command introduced in Junos OS Release 11.3.
Description	Display the IDP policy assigned to a logical system. The IDP policy is assigned to a logical system through the security profile.
Required Privilege Level	view
Related Documentation	• security-profile on page 355
List of Sample Output	show security idp logical-system policy-association on page 398
Output Fields	Table 32 on page 398 lists the output fields for the **show security idp logical-system policy-association** command.

Table 32: show security idp logical-system policy-association Output Fields

Field Name	Field Description
Logical system	Name of the logical system to which an IDP policy is assigned.
IDP policy	Name of the IDP policy that is specified in the security profile that is bound to the logical system.

Sample Output

show security idp logical-system policy-association

```
user@host> show security idp logical-system policy-association
Logical system        IDP policy
root-logical-system    idp-policy1
lsys1                  idp-policy2
```

show security ike security-associations

Syntax	show security ike security-associations *peer-address* brief | detail family (inet | inet6) fpc *slot-number* index *SA-index-number* kmd-instance (all | *kmd-instance-name*) pic *slot-number* sa-type shortcut <detail>
Release Information	Command introduced in Junos OS Release 8.5 . Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for Auto Discovery VPN added in Junos OS Release 12.3X48-D10. Support for IKEv2 reauthentication and IKEv2 fragmentation added in Junos OS Release 15.1X49-D60.
Description	Display information about Internet Key Exchange security associations (IKE SAs).

Options
- none—Display standard information about existing IKE SAs, including index numbers.

- *peer-address*—(Optional) Display details about a particular SA based on the IPv4 or IPv6 address of the destination peer. This option and **index** provide the same level of output.

- **brief**—(Optional) Display standard information about all existing IKE SAs. (Default)

- **detail**—(Optional) Display detailed information about all existing IKE SAs.

- **family**—(Optional) Display IKE SAs by family. This option is used to filter the output.

  - **inet**—IPv4 address family.

  - **inet6**—IPv6 address family.

- **fpc** *slot-number*—(Optional) Display information about existing IKE SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.

- **index** *SA-index-number*—(Optional) Display information for a particular SA based on the index number of the SA. For a particular SA, display the list of existing SAs by using the command with no options. This option and *peer-address* provide the same level of output.

- **kmd-instance** —(Optional) Display information about existing IKE SAs in the key management process (in this case, it is KMD) identified by FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.

  - **all**—All KMD instances running on the Services Processing Unit (SPU).

  - *kmd-instance-name*—Name of the KMD instance running on the SPU.

- **pic** *slot-number* —(Optional) Display information about existing IKE SAs in this PIC slot. This option is used to filter the output.

- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.

Required Privilege Level
view

Related Documentation
- Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 133

List of Sample Output
show security ike security-associations (IPv4) on page 402
show security ike security-associations (IPv6) on page 403
show security ike security-associations detail (Branch SRX Series Devices) on page 403
show security ike security-associations detail (High-End SRX Series Devices) on page 403
show security ike security-associations family inet6 on page 404
show security ike security-associations index 8 detail on page 404
show security ike security-associations 192.168.1.2 on page 405
show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices) on page 405
show security ike security-associations detail (ADVPN Suggester, Static Tunnel) on page 405
show security ike security-associations detail (ADVPN Partner, Static Tunnel) on page 405
show security ike security-associations detail (ADVPN Partner, Shortcut) on page 405
show security ike security-associations sa-type shortcut (ADVPN) on page 406
show security ike security-associations sa-type shortcut detail (ADVPN) on page 406
show security ike security-associations detail (IKEv2 Reauthentication) on page 406

Output Fields
Table 33 on page 400 lists the output fields for the **show security ike security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 33: show security ike security-associations Output Fields

Field Name	Field Description
IKE Peer or Remote Address	IP address of the destination peer with which the local peer communicates.
Index	Index number of an SA. This number is an internally generated number you can use to display information about a single SA.
Gateway Name	Name of the IKE gateway.
Location	• **FPC**—Flexible PIC Concentrator (FPC) slot number. • **PIC**—PIC slot number. • **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances are running on each SPU, and any particular IKE negotiation is carried out by a single KMD instance.
Role	Part played in the IKE session. The device triggering the IKE negotiation is the initiator, and the device accepting the first IKE exchange packets is the responder.

Table 33: show security ike security-associations Output Fields *(continued)*

Field Name	Field Description
State	State of the IKE SAs: • **DOWN**—SA has not been negotiated with the peer. • **UP**—SA has been negotiated with the peer.
Initiator cookie	Random number, called a cookie, which is sent to the remote node when the IKE negotiation is triggered.
Responder cookie	Random number generated by the remote node and sent back to the initiator as a verification that the packets were received. A cookie is aimed at protecting the computing resources from attack without spending excessive CPU resources to determine the cookie's authenticity.
Mode or Exchange type	Negotiation method agreed on by the two IPsec endpoints, or peers, used to exchange information between one another. Each exchange type determines the number of messages and the payload types that are contained in each message. The modes, or exchange types, are: • **main**—The exchange is done with six messages. This mode or exchange type encrypts the payload, protecting the identity of the neighbor. The authentication method used is displayed: preshared keys or certificate. • **aggressive**—The exchange is done with three messages. This mode or exchange type does not encrypt the payload, leaving the identity of the neighbor unprotected. NOTE: IKEv2 protocol does not use the mode configuration for negotiation. Therefore, the mode displays the version number of the security association.
Local	Address of the local peer.
Remote	Address of the remote peer.
Lifetime	Number of seconds remaining until the IKE SA expires.
Reauth Lifetime	When enabled, number of seconds remaining until reauthentication triggers a new IKEv2 SA negotiation.

Algorithms	IKE algorithms used to encrypt and secure exchanges between the peers during the IPsec Phase 2 process:
	
	• **Authentication**—Type of authentication algorithm used:
	  • **sha1**—Secure Hash Algorithm 1 authentication.
	  • **md5**—MD5 authentication.
	• **Encryption**—Type of encryption algorithm used:
	  • **aes-256-cbc**—Advanced Encryption Standard (AES) 256-bit encryption.
	  • **aes-192-cbc**— AES192-bit encryption.
	  • **aes-128-cbc**—AES 128-bit encryption.
	  • **3des-cbc**—3 Data Encryption Standard (DES) encryption.
	  • **des-cbc**—DES encryption.

Table 33: show security ike security-associations Output Fields *(continued)*

Field Name	Field Description
Diffie-Hellman group	Specifies the IKE Diffie-Hellman group.
Traffic statistics	• **Input bytes**—Number of bytes received. • **Output bytes**—Number of bytes transmitted. • **Input packets**—Number of packets received. • **Output packets**—Number of packets transmitted.
Flags	Notification to the key management process of the status of the IKE negotiation: • **caller notification sent**—Caller program notified about the completion of the IKE negotiation. • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire. • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation. • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.
IPSec security associations	• *number* **created**: The number of SAs created. • *number* **deleted**: The number of SAs deleted.

Phase 2 negotiations in progress	Number of Phase 2 IKE negotiations in progress and status information:
	
	• **Negotiation type**—Type of Phase 2 negotiation. Junos OS currently supports quick mode.
	• **Message ID**—Unique identifier for a Phase 2 negotiation.
	• **Local identity**—Identity of the local Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• **Remote identity**—Identity of the remote Phase 2 negotiation. The format is *id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)*.
	• Flags—Notification to the key management process of the status of the IKE negotiation:
	  • **caller notification sent**—Caller program notified about the completion of the IKE negotiation.
	  • **waiting for done**—Negotiation is done. The library is waiting for the remote end retransmission timers to expire.
	  • **waiting for remove**—Negotiation has failed. The library is waiting for the remote end retransmission timers to expire before removing this negotiation.
	  • **waiting for policy manager**—Negotiation is waiting for a response from the policy manager.

Sample Output

show security ike security-associations (IPv4)

```
user@host> show security ike security-associations
Index  Remote Address  State  Initiator cookie       Responder cookie Mode
8  192.168.1.2  UP  3a895f8a9f620198 9040753e66d700bb Main
Index  Remote Address  State  fInitiator cookie Responder cookie Mode
9  192.168.1.3   UP  5ba96hfa9f65067  70890755b65b80b  Main
```

show security ike security-associations (IPv6)

```
user@host> show security ike security-associations
Index   State  Initiator cookie  Responder cookie  Mode        Remote Address

5       UP     e48efd6a444853cf  0d09c59aafb720be  Aggressive  2001:db8::1112
```

show security ike security-associations detail (Branch SRX Series Devices)

```
user@host> show security ike security-associations detail
IKE peer 192.168.134.245, Index 2577565, Gateway Name: tropic
  Role: Initiator, State: UP
  Initiator cookie: b869b3424513340a, Responder cookie: 4cb3488cb19397c3
  Exchange type: Main, Authentication method: Pre-shared-keys
  Local: 192.168.134.241:500, Remote: 192.168.134.245:500
  Lifetime: Expires in 169 seconds
  Peer ike-id: 192.168.134.245
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes128-cbc
   Pseudo random function: hmac-sha1
  Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 1012
   Output bytes  :                 1196
   Input  packets:                    4
   Output packets:                    5
  Flags: IKE SA is created
  IPSec security associations: 1 created, 0 deleted
  Phase 2 negotiations in progress: 0

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 192.168.134.241:500, Remote: 192.168.134.245:500
    Local identity: 192.168.134.241
    Remote identity: 192.168.134.245
    Flags: IKE SA is created
```

show security ike security-associations detail (High-End SRX Series Devices)

```
user@host> show security ike security-associations detail
IKE peer 192.168.2, Index 914039858, Gateway Name: tropic
  Location: FPC 3, PIC 1, KMD-Instance 3
  Role: Initiator, State: UP
  Initiator cookie: 219a697652bdde37, Responder cookie: b49c30b229d36bcd
  Exchange type: Aggressive, Authentication method: Pre-shared-keys
  Local: 192.168.1.1:500, Remote: 192.168.1.2:500
  Lifetime: Expires in 26297 seconds
  Peer ike-id: 192.168.1.2
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : 3des-cbc
   Pseudo random function: hmac-sha1
  Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                    0
   Output bytes  :                    0
   Input  packets:                    0
   Output packets:                    0
```

```
                                IPSec security associations: 0 created, 0 deleted
                                Phase 2 negotiations in progress: 1
```

show security ike security-associations family inet6

```
                    user@host> show security ike security-associations family inet6
                      IKE peer 2001:db8:1212::1112, Index 5, Gateway Name: tropic
                      Role: Initiator, State: UP
                      Initiator cookie: e48efd6a444853cf, Responder cookie: 0d09c59aafb720be
                      Exchange type: Aggressive, Authentication method: Pre-shared-keys
                      Local: 2001:db8:1212::1111:500, Remote: 2001:db8:1212::1112:500
                      Lifetime: Expires in 19518 seconds
                      Peer ike-id: not valid
                      Xauth assigned IP: 0.0.0.0
                      Algorithms:
                       Authentication        : sha1
                       Encryption            : 3des-cbc
                       Pseudo random function: hmac-sha1
                       Diffie-Hellman group  : DH-group-5
                      Traffic statistics:
                       Input  bytes  :                  1568
                       Output bytes  :                  2748
                       Input  packets:                     6
                       Output packets:                    23
                      Flags: Caller notification sent
                      IPSec security associations: 5 created, 0 deleted
                      Phase 2 negotiations in progress: 1

                        Negotiation type: Quick mode, Role: Initiator, Message ID: 2900338624
                        Local: 2001:db8:1212::1111:500, Remote: 2001:db8:1212::1112:500
                        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)
                        Flags: Caller notification sent, Waiting for done
```

show security ike security-associations index 8 detail

```
                    user@host> show security ike security-associations index 8 detail
                    IKE peer 192.168.1.2, Index 8, Gateway Name: tropic
                     Role: Responder, State:UP
                     Initiator cookie: 3a895f8a9f620198, Responder cookie: 9040753e66d700bb
                     Exchange type; main, Authentication method: Pre-shared-keys
                     Local: 192.168.1.1:500, Remote: 192.168.1.2:500
                     Lifetime: Expired in 381 seconds
                     Algorithms:
                      Authentication:         md5
                      Encryption:         3des-cbc
                      Pseudo random function         hmac-md5
                        Diffie-Hellman group   : DH-group-5
                       Traffic statistics:
                      Input bytes:         11268
                      Output bytes:         6940
                      Input packets:         57
                      Output packets:         57
                     Flags: Caller notification sent
                     IPsec security associations: 0 created, 0 deleted
                     Phase 2 negotiations in progress: 1

                      Negotiation type: Quick mode, Role: Responder, Message ID: 1765792815
                      Local: 192.168.1.1:500, Remote: 192.168.1.2:500
                      Local identity: No Id
```

```
         Remote identity: No Id
         Flags: Caller notification sent, Waiting for remove
```

show security ike security-associations 192.168.1.2

```
user@host> show security ike security-associations 192.168.1.2
Index     State   Initiator cookie  Responder cookie  Mode Remote Address
    8       UP     3a895f8a9f620198  9040753e66d700bb  Main 192.168.1.2
```

show security ike security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
user@host> show security ike security-associations fpc 6 pic 1 kmd-instance all
Index      Remote Address  State  Initiator cookie  Responder cookie  Mode

1728053250 192.168.1.2      UP     fc959afd1070d10b  bdeb7e8c1ea99483  Main
```

show security ike security-associations detail (ADVPN Suggester, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.105, Index 13563297, Gateway Name: zth_hub_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent         :  12
     Suggestion response accepted:  12
     Suggestion response declined:   0
  Role: Responder, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.154:500, Remote: 192.168.0.105:500
  Lifetime: Expires in 26429 seconds
  Peer ike-id: DC=example, CN=host02, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.154, Index 4980720, Gateway Name: zth_spoke_gw
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
 Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:  12
     Suggestions accepted:  12
     Suggestions declined:   0
  Role: Initiator, State: UP
  Initiator cookie: 4d3f4e4b2e75d727, Responder cookie: 81ab914e13cecd21
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.0.105:500, Remote: 192.168.0.154:500
  Lifetime: Expires in 26252 seconds
  Peer ike-id: DC=example, CN=host01, OU=SBU, O=example, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations detail (ADVPN Partner, Shortcut)

```
user@host> show security ike security-associations detail
IKE peer 192.168.0.106, Index 4980737, Gateway Name:
GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173323
  Location: FPC 0, PIC 0, KMD-Instance 1
  Auto Discovery VPN:
   Type: Shortcut, Local Capability: Partner, Peer Capability: Partner
  Role: Responder, State: UP
  Initiator cookie: e1ed0c655929debc, Responder cookie: 437de6ed784ba63e
```

```
                   Exchange type: IKEv2, Authentication method: RSA-signatures
                   Local: 192.168.0.105:500, Remote: 192.168.0.106:500
                   Lifetime: Expires in 28796 seconds
                   Peer ike-id: DC=example, CN=paulyd, L=Sunnyvale, ST=CA, C=US
```

show security ike security-associations sa-type shortcut (ADVPN)

```
         user@host> show security ike security-associations sa-type shortcut
         Index   State  Initiator cookie  Responder cookie  Mode          Remote Address

         4980742 UP     vb56fbe694eaee5b6 064dbccbfa3b2aab  IKEv2          192.168.0.106
```

show security ike security-associations sa-type shortcut detail (ADVPN)

```
         user@host> show security ike security-associations sa-type shortcut detail
         IKE peer 192.168.0.106, Index 4980742, Gateway Name:
         GW-ADVPN-GT-ADVPN-zth_spoke_vpn-268173327
           Location: FPC 0, PIC 0, KMD-Instance 1
           Auto Discovery VPN:
            Type: Shortcut, Local Role: Partner, Peer Role: Partner
           Role: Responder, State: UP
```

show security ike security-associations detail (IKEv2 Reauthentication)

```
         user@host> show security ike security-associations detail
         IKE peer 10.1.2.11, Index 6009224, Gateway Name: GW
           Role: Responder, State: UP
           Initiator cookie: 2c74d14c798a9d70, Responder cookie: 83cbb49bfbcb80cb
           Exchange type: IKEv2, Authentication method: RSA-signatures
           Local: 10.1.1.11:500, Remote: 10.1.2.11:500
           Lifetime: Expires in 173 seconds
           Reauth Lifetime: Expires in 600 seconds
           Peer ike-id: vsrx@example.net
           Xauth assigned IP: 0.0.0.0
           Algorithms:
            Authentication         : hmac-sha1-96
            Encryption             : aes128-cbc
            Pseudo random function: hmac-sha1
            Diffie-Hellman group   : DH-group-2
           Traffic statistics:
            Input  bytes  :                1782
            Output bytes  :                1743
            Input  packets:                   2
```

show security ipsec security-associations

Syntax   show security ipsec security-associations
brief | detail
family (inet | inet6)
fpc *slot-number*
index *SA-index-number*
kmd-instance (all | *kmd-instance-name*)
pic *slot-number>*
sa-type shortcut
vpn-name *vpn-name* <traffic-selector *traffic-selector-name*>

Release Information   Command introduced in Junos OS Release 8.5. Support for the **fpc**, **pic**, and **kmd-instance** options added in Junos OS Release 9.3. Support for the **family** option added in Junos OS Release 11.1. Support for the **vpn-name** option added in Junos OS Release 11.4R3. Support for the **traffic-selector** option and traffic selector field added in Junos OS Release 12.1X46-D10. Support for Auto Discovery VPN (ADVPN) added in Junos OS Release 12.3X48-D10. Support for IPsec datapath verification added in Junos OS Release 15.1X49-D70.

Description   Display information about the IPsec security associations (SAs).

Options
- none—Display information about all SAs.

- **brief | detail**—(Optional) Display the specified level of output.

- **family**—(Optional) Display SAs by family. This option is used to filter the output.

  - **inet**—IPv4 address family.

  - **inet6**—IPv6 address family.

- **fpc** *slot-number*—(Optional) Display information about existing IPsec SAs in this Flexible PIC Concentrator (FPC) slot. This option is used to filter the output.

- **index** *SA-index-number*—(Optional) Display detailed information about the specified SA identified by this index number. To obtain a list of all SAs that includes their index numbers, use the command with no options.

- **kmd-instance**—(Optional) Display information about existing IPsec SAs in the key management process (in this case, it is KMD) identified by the FPC *slot-number* and PIC *slot-number*. This option is used to filter the output.

  - **all**—All KMD instances running on the Services Processing Unit (SPU).

  - *kmd-instance-name*—Name of the KMD instance running on the SPU.

- **pic** *slot-number*—(Optional) Display information about existing IPsec SAs in this PIC slot. This option is used to filter the output.

- **sa-type**—(Optional for ADVPN) Type of SA. **shortcut** is the only option for this release.

- **vpn-name** *vpn-name*—Name of the VPN. If configured, **traffic-selector** *traffic-selector-name* can optionally be specified.

Required Privilege Level	view

Related Documentation	• *clear security ipsec security-associations*
	• Example: Configuring a Route-Based VPN Tunnel in a User Logical System on page 133

Output Fields Table 34 on page 408 lists the output fields for the **show security ipsec security-associations** command. Output fields are listed in the approximate order in which they appear.

Table 34: show security ipsec security-associations

Field Name	Field Description
Total active tunnels	Total number of active IPsec tunnels.
ID	Index number of the SA. You can use this number to get additional information about the SA.
VPN name	IPsec name for VPN.
Gateway	IP address of the remote gateway.
Port	If Network Address Translation (NAT) is used, this value is 4500. Otherwise, it is the standard IKE port, 500.
Algorithm	Cryptography used to secure exchanges between peers during the IKE Phase 2 negotiations includes: • An authentication algorithm used to authenticate exchanges between the peers. Options are **hmac-md5-95**, **hmac-sha1-96**, or **ESP**. • An encryption algorithm used to encrypt data traffic. Options are **3des-cbc**, **aes-128-cbc**, **aes-192-cbc**, **aes-256-cbc**, or **des-cbc.**

Table 34: show security ipsec security-associations *(continued)*

Field Name	Field Description
SPI	Security parameter index (SPI) identifier. An SA is uniquely identified by an SPI. Each entry includes the name of the VPN, the remote gateway address, the SPIs for each direction, the encryption and authentication algorithms, and keys. The peer gateways each have two SAs, one resulting from each of the two phases of negotiation: Phase 1 and Phase 2.
Life: sec/kb	The lifetime of the SA, after which it expires, expressed either in seconds or kilobytes.
Sta	State has two options, **Installed** and **Not Installed**. • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.
Mon	The Mon field refers to VPN monitoring status. If VPN monitoring is enabled, then this field displays **U** (up) or **D** (down). A hyphen (-) means VPN monitoring is not enabled for this SA. A **V** means that IPsec datapath verification is in progress.
vsys or Virtual-system	The root system.
Tunnel index	Numeric identifier of the specific IPsec tunnel for the SA.
Local gateway	Gateway address of the local system.
Remote gateway	Gateway address of the remote system.
Traffic selector	Name of the traffic selector.
Local identity	Identity of the local peer so that its partner destination gateway can communicate with it. The value is specified as an IP address, fully qualified domain name, e-mail address, or distinguished name (DN).
Remote identity	IP address of the destination peer gateway.
DF-bit	State of the don't fragment bit: **set** or **cleared**.
Policy-name	Name of the applicable policy.
Location	**FPC**—Flexible PIC Concentrator (FPC) slot number. **PIC**—PIC slot number. **KMD-Instance**—The name of the KMD instance running on the SPU, identified by FPC *slot-number* and PIC *slot-number*. Currently, 4 KMD instances running on each SPU, and any particular IPsec negotiation is carried out by a single KMD instance.
Tunnel events	Tunnel event and the number of times the event has occurred. See *Tunnel Events* for descriptions of tunnel events and the action you can take.
Direction	Direction of the SA; it can be inbound or outbound.

Table 34: show security ipsec security-associations *(continued)*

Field Name	Field Description
AUX-SPI	Value of the auxiliary security parameter index(SPI). • When the value is **AH** or **ESP**, **AUX-SPI** is always 0. • When the value is **AH+ESP**, **AUX-SPI** is always a positive integer.
Mode	Mode of the SA: • **transport**—Protects host-to-host connections. • **tunnel**—Protects connections between security gateways.
Type	Type of the SA: • **manual**—Security parameters require no negotiation. They are static and are configured by the user. • **dynamic**—Security parameters are negotiated by the IKE protocol. Dynamic SAs are not supported in transport mode.
State	State of the SA: • **Installed**—The SA is installed in the SA database. • **Not Installed**—The SA is not installed in the SA database. For transport mode, the value of State is always **Installed**.

Protocol	Protocol supported.
	
	• Transport mode supports Encapsulation Security Protocol (ESP) and Authentication Header (AH).
	• Tunnel mode supports ESP and AH.
	  • **Authentication**—Type of authentication used.
	  • **Encryption**—Type of encryption used.

Soft lifetime	The soft lifetime informs the IPsec key management system that the SA is about to expire. Each lifetime of an SA has two display options, hard and soft, one of which must be present for a dynamic SA. This allows the key management system to negotiate a new SA before the hard lifetime expires. • **Expires in seconds**—Number of seconds left until the SA expires.
Hard lifetime	The hard lifetime specifies the lifetime of the SA. • **Expires in seconds**—Number of seconds left until the SA expires.
Lifesize Remaining	The lifesize remaining specifies the usage limits in kilobytes. If there is no lifesize specified, it shows unlimited. • **Expires in kilobytes**—Number of kilobytes left until the SA expires.
Anti-replay service	State of the service that prevents packets from being replayed. It can be **Enabled** or **Disabled**.

Table 34: show security ipsec security-associations *(continued)*

Field Name	Field Description
Replay window size	Configured size of the antireplay service window. It can be 32 or 64 packets. If the replay window size is 0, the antireplay service is disabled. The antireplay window size protects the receiver against replay attacks by rejecting old or duplicate packets.
Bind-interface	The tunnel interface to which the route-based VPN is bound.
Copy-Outer-DSCP	Indicates if copying outer IP header DSCP and ECN to inner IP header is enabled or disabled.

Sample Output

show security ipsec security-associations (IPv4)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID      Gateway            Port  Algorithm      SPI        Life:sec/kb  Mon vsys
  131075 192.168.28.241    500   ESP:3des/sha1  86758ff0 6918/ unlim   -   0
  131075 192.168.28.241    500   ESP:3des/sha1  3183ff26 6918/ unlim   -   0
```

show security ipsec security-associations (IPv6)

```
user@host> show security ipsec security-associations
Total active tunnels: 1
ID      Algorithm     SPI        Life:sec/kb  Mon  vsys Port  Gateway
131074 ESP:3des/sha1 14caf1d9 3597/ unlim    -    root 500   2001:db8::1112
131074 ESP:3des/sha1 9a4db486 3597/ unlim    -    root 500   2001:db8::1112
```

show security ipsec security-associations index 131073

```
user@host> show security ipsec security-associations index 131073
  ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
  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)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
```

```
                    Hard lifetime: Expires in 1774 seconds
                    Lifesize Remaining:  Unlimited
                    Soft lifetime: Expires in 1151 seconds
                    Mode: Tunnel(0 0), Type: dynamic, State: installed
                    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
                    Anti-replay service: counter-based enabled

                    , Replay window size: 64
                    Location: FPC 0, PIC 1, KMD-Instance 1
                    Direction: outbound, SPI: 727f629d, AUX-SPI: 0
                    Hard lifetime: Expires in 1774 seconds
                    Lifesize Remaining:  Unlimited
                    Soft lifetime: Expires in 1151 seconds
                    Mode: Tunnel(0 0), Type: dynamic, State: installed
                    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
                    Anti-replay service: counter-based enabled

                    , Replay window size: 64
```

show security ipsec security-associations brief

```
                    user@host> show security ipsec security-associations brief
                    Total active tunnels: 2
                    ID      Gateway     Port Algorithm    SPI        Life:sec/kb Mon vsys
                    <16384 192.168.1.1 500  ESP:3des/sha1 af88baa   28795/unlim D   0
                    >16384 192.168.1.1 500  ESP:3des/sha1 f4e3e5f4 28795/unlim D   0
```

show security ipsec security-associations detail

```
                    user@host> show security ipsec security-associations detail
                    ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
                    Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
                    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)
                    Version: IKEv1
                    DF-bit: clear
                    , Copy-Outer-DSCP Enabled
                    Bind-interface: st0.99

                    Port: 500, Nego#: 8, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
                    Tunnel events:
                    Mon Oct 26 2015 22:27:50 -0700: IPSec SA rekey successfully completed (7 times)
                    Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
                    times)
                    Mon Oct 26 2015 16:41:07 -0700: IKE SA negotiation successfully completed (1
                    times)
                    Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
                    peer to trigger negotiation (1 times)
                    Mon Oct 26 2015 16:40:56 -0700: External interface's address received. Information
                     updated (1 times)
                    Location: FPC 0, PIC 1, KMD-Instance 1
                    Direction: inbound, SPI: 81ed9998, AUX-SPI: 0
                    Hard lifetime: Expires in 2296 seconds
                    Lifesize Remaining:  Unlimited
                    Soft lifetime: Expires in 1688 seconds
                    Mode: Tunnel(0 0), Type: dynamic, State: installed
                    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
                    Anti-replay service: counter-based enabled

                    , Replay window size: 64
                    Location: FPC 0, PIC 1, KMD-Instance 1
```

```
                              Direction: outbound, SPI: 80565248, AUX-SPI: 0
                              Hard lifetime: Expires in 2296 seconds
                              Lifesize Remaining:  Unlimited
                              Soft lifetime: Expires in 1688 seconds
                              Mode: Tunnel(0 0), Type: dynamic, State: installed
                              Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
                              Anti-replay service: counter-based enabled

                            , Replay window size: 64
```

show security ipsec security-associations family inet6

```
                         user@host> show security ipsec security-associations family inet6
                           Virtual-system: root
                           Local Gateway: 2001:db8:1212::1111, Remote Gateway: 2001:db8:1212::1112
                           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)
                             DF-bit: clear
                             Direction: inbound, SPI: 14caf1d9, AUX-SPI: 0
                                                 , VPN Monitoring: -
                             Hard lifetime: Expires in 3440 seconds
                             Lifesize Remaining:  Unlimited
                             Soft lifetime: Expires in 2813 seconds
                             Mode: tunnel, Type: dynamic, State: installed
                             Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
                             Anti-replay service: counter-based enabled, Replay window size: 64

                             Direction: outbound, SPI: 9a4db486, AUX-SPI: 0
                                                 , VPN Monitoring: -
                             Hard lifetime: Expires in 3440 seconds
                             Lifesize Remaining:  Unlimited
                             Soft lifetime: Expires in 2813 seconds
                             Mode: tunnel, Type: dynamic, State: installed
                             Protocol: ESP, Authentication: hmac-sha1-96, Encryption: 3des-cbc
                             Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations fpc 6 pic 1 kmd-instance all (SRX Series Devices)

```
                         user@host> show security ipsec security-associations fpc 6 pic 1 kmd-instance all
                           Total active tunnels: 1

                           ID      Gateway         Port  Algorithm     SPI        Life:sec/kb  Mon vsys

                           <2      192.168.1.2     500   ESP:3des/sha1 67a7d25d  28280/unlim   -   0

                           >2      192.168.1.2     500   ESP:3des/sha1 a23cbcdc  28280/unlim   -   0
```

show security ipsec security-associations detail (ADVPN Suggester, Static Tunnel)

```
                         user@host> show security ipsec security-associations detail
                         ID: 70516737 Virtual-system: root, VPN Name: ZTH_HUB_VPN
                           Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
                           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)
                           Version: IKEv2
                           DF-bit: clear
                           Bind-interface: st0.1

                           Port: 500, Nego#: 5, Fail#: 0, Def-Del#: 0 Flag: 0x608a29
                           Tunnel events:
                           Tue Nov 03 2015 01:24:27 -0800: IPSec SA negotiation successfully completed (1
                         times)
```

```
      Tue Nov 03 2015 01:24:27 -0800: IKE SA negotiation successfully completed (4
    times)
      Tue Nov 03 2015 01:23:38 -0800: User cleared IPSec SA from CLI (1 times)
      Tue Nov 03 2015 01:21:32 -0800: IPSec SA negotiation successfully completed (1
    times)
      Tue Nov 03 2015 01:21:31 -0800: IPSec SA delete payload received from peer,
    corresponding IPSec SAs cleared (1 times)
      Tue Nov 03 2015 01:21:27 -0800: IPSec SA negotiation successfully completed (1
    times)
      Tue Nov 03 2015 01:21:13 -0800: Tunnel configuration changed. Corresponding
    IKE/IPSec SAs are deleted (1 times)
      Tue Nov 03 2015 01:19:27 -0800: IPSec SA negotiation successfully completed (1
    times)
      Tue Nov 03 2015 01:19:27 -0800: Tunnel is ready. Waiting for trigger event or
    peer to trigger negotiation (1 times)
      Location: FPC 0, PIC 3, KMD-Instance 2
      Direction: inbound, SPI: 43de5d65, AUX-SPI: 0
      Hard lifetime: Expires in 1335 seconds
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expires in 996 seconds
      Mode: Tunnel(0 0), Type: dynamic, State: installed
      Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
      Anti-replay service: counter-based enabled

      , Replay window size: 64
      Location: FPC 0, PIC 3, KMD-Instance 2
      Direction: outbound, SPI: 5b6e157c, AUX-SPI: 0
      Hard lifetime: Expires in 1335 seconds
      Lifesize Remaining:  Unlimited
      Soft lifetime: Expires in 996 seconds
      Mode: Tunnel(0 0), Type: dynamic, State: installed
      Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
      Anti-replay service: counter-based enabled

      , Replay window size: 64
```

show security ike sa index 222075191 detail

```
user@host> show security ike sa index 222075191 detail
node0:
--------------------------------------------------------------------------
IKE peer 192.168.1.2, Index 222075191, Gateway Name: ZTH_HUB_GW
  Location: FPC 0, PIC 3, KMD-Instance 2
  Auto Discovery VPN:
   Type: Static, Local Capability: Suggester, Peer Capability: Partner
   Suggester Shortcut Suggestions Statistics:
     Suggestions sent   :    2
     Suggestions accepted:    4
     Suggestions declined:    1
  Role: Responder, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.1.1:500, Remote: 192.168.1.2:500
  Lifetime: Expires in 828 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering,
CN=cssvk36-d
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
```

```
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :                 20474
   Output bytes  :                 21091
   Input  packets:                   237
   Output packets:                   237
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Responder, Message ID: 0
    Local: 192.168.1.1:500, Remote: 192.168.1.2:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example,
OU=engineering, CN=host1
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example,
OU=engineering, CN=host2
    Flags: IKE SA is created
```

show security ipsec security-associations detail (ADVPN Partner, Static Tunnel)

```
user@host> show security ipsec security-associations detail
ID: 67108872 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 192.168.1.2, Remote Gateway: 192.168.1.1
  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)
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x8608a29
  Tunnel events:
  Tue Nov 03 2015 01:24:26 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:24:26 -0800: IKE SA negotiation successfully completed (4
times)
  Tue Nov 03 2015 01:23:37 -0800: IPSec SA delete payload received from peer,
corresponding IPSec SAs cleared (1 times)
  Tue Nov 03 2015 01:21:31 -0800: IPSec SA negotiation successfully completed (1
 times)
  Tue Nov 03 2015 01:21:31 -0800: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Tue Nov 03 2015 01:18:26 -0800: Key pair not found for configured local
certificate. Negotiation failed (1 times)
  Tue Nov 03 2015 01:18:13 -0800: CA certificate for configured local certificate
 not found. Negotiation not initiated/successful (1 times)
  Direction: inbound, SPI: 5b6e157c, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: 43de5d65, AUX-SPI: 0
  Hard lifetime: Expires in 941 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 556 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
  Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ike sa index 788674 detail

```
user@host> show security ike sa index 788674 detail
```

```
IKE peer 192.168.1.1, Index 788674, Gateway Name: ZTH_SPOKE_GW
  Auto Discovery VPN:
   Type: Static, Local Capability: Partner, Peer Capability: Suggester
   Partner Shortcut Suggestions Statistics:
     Suggestions received:    2
     Suggestions accepted:    2
     Suggestions declined:    0
  Role: Initiator, State: UP
  Initiator cookie: 7b996b4c310d2424, Responder cookie: 5724c5882a212157
  Exchange type: IKEv2, Authentication method: RSA-signatures
  Local: 192.168.1.2:500, Remote: 192.168.1.1:500
  Lifetime: Expires in 734 seconds
  Peer ike-id: C=US, DC=example, ST=CA, L=Sunnyvale, O=example, OU=engineering,
CN=test
  Xauth user-name: not available
  Xauth assigned IP: 0.0.0.0
  Algorithms:
   Authentication        : hmac-sha1-96
   Encryption            : aes256-cbc
   Pseudo random function: hmac-sha1
   Diffie-Hellman group  : DH-group-5
  Traffic statistics:
   Input  bytes  :               22535
   Output bytes  :               21918
   Input  packets:                 256
   Output packets:                 256
  IPSec security associations: 2 created, 0 deleted
  Phase 2 negotiations in progress: 1

    Negotiation type: Quick mode, Role: Initiator, Message ID: 0
    Local: 192.168.1.2:500, Remote: 192.168.1.1:500
    Local identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example,
OU=engineering, CN=host1
    Remote identity: C=US, DC=example, ST=CA, L=Sunnyvale, O=example,
OU=engineering, CN=host2
    Flags: IKE SA is created
```

show security ipsec security-associations sa-type shortcut (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut
Total active tunnels: 1
ID    Algorithm      SPI     Life:sec/kb  Mon lsys Port  Gateway
<268173318 ESP:aes-cbc-256/sha1 6f164ee0 3580/ unlim - root 500 192.168.0.111

>268173318 ESP:aes-cbc-256/sha1 e6f29cb0 3580/ unlim - root 500 192.168.0.111
```

show security ipsec security-associations sa-type shortcut detail (ADVPN)

```
user@host> show security ipsec security-associations sa-type shortcut detail
node0:
--------------------------------------------------------------------------

ID: 67108874 Virtual-system: root, VPN Name: ZTH_SPOKE_VPN
  Local Gateway: 192.168.1.2, Remote Gateway: 192.168.1.2
  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)
  Auto Discovery VPN:
    Type: Shortcut, Shortcut Role: Initiator
  Version: IKEv2
  DF-bit: clear, Bind-interface: st0.1
  Port: 4500, Nego#: 0, Fail#: 0, Def-Del#: 0 Flag: 0x40608a29
```

```
   Tunnel events:
     Tue Nov 03 2015 01:47:26 -0800: IPSec SA negotiation successfully completed
(1 times)
     Tue Nov 03 2015 01:47:26 -0800: Tunnel is ready. Waiting for trigger event
or peer to trigger negotiation (1 times)
     Tue Nov 03 2015 01:47:26 -0800: IKE SA negotiation successfully completed (1
 times)
  Direction: inbound, SPI: b7a5518, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
  Direction: outbound, SPI: b7e0268, AUX-SPI: 0
    Hard lifetime: Expires in 1766 seconds
    Lifesize Remaining:  Unlimited
    Soft lifetime: Expires in 1381 seconds
    Mode: Tunnel(0 0), Type: dynamic, State: installed
    Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (192 bits)
    Anti-replay service: counter-based enabled, Replay window size: 64
```

show security ipsec security-associations family inet detail

```
user@host> show security ipsec security-associations family inet detail
ID: 131073 Virtual-system: root, VPN Name: ike-vpn-chicago
  Local Gateway: 192.168.1.1, Remote Gateway: 192.168.1.2
  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)
  Version: IKEv1
  DF-bit: clear
  , Copy-Outer-DSCP Enabled
  Bind-interface: st0.99

  Port: 500, Nego#: 116, Fail#: 0, Def-Del#: 0 Flag: 0x600a29
  Tunnel events:
  Fri Oct 30 2015 15:47:21 -0700: IPSec SA rekey successfully completed (115
times)
  Fri Oct 30 2015 11:38:35 -0700: IKE SA negotiation successfully completed (12
times)
  Mon Oct 26 2015 16:41:07 -0700: IPSec SA negotiation successfully completed (1
 times)
  Mon Oct 26 2015 16:40:56 -0700: Tunnel is ready. Waiting for trigger event or
peer to trigger negotiation (1 times)
  Mon Oct 26 2015 16:40:56 -0700: External interface's address received.
Information updated (1 times)
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: inbound, SPI: 81b9fc17, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
  Mode: Tunnel(0 0), Type: dynamic, State: installed
  Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
  Anti-replay service: counter-based enabled

  , Replay window size: 64
  Location: FPC 0, PIC 1, KMD-Instance 1
  Direction: outbound, SPI: 727f629d, AUX-SPI: 0
  Hard lifetime: Expires in 1713 seconds
  Lifesize Remaining:  Unlimited
  Soft lifetime: Expires in 1090 seconds
```

```
Mode: Tunnel(0 0), Type: dynamic, State: installed
Protocol: ESP, Authentication: hmac-sha1-96, Encryption: aes-cbc (128 bits)
Anti-replay service: counter-based enabled

, Replay window size: 64
```

show security match-policies

Syntax show security match-policies
destination-ip *ip-address*
destination-port *port-number*
from-zone *zone-name*
protocol *protocol-name* | *protocol-number*
result-count *number*
source-identity *role-name*
source-ip *ip-address*
source-port *port-number*
to-zone *zone-name*

Release Information Command introduced in Junos OS Release 10.3. Command updated in Junos OS Release 10.4 . Command updated in Junos OS Release 12.1. Command updated to include optional from-zone and to-zone global match options in Junos OS Release 12.1X47-D10.

Description The **show security match-policies** command allows you to troubleshoot traffic problems using the match criteria: source port, destination port, source IP address, destination IP address, and protocol. For example, if your traffic is not passing because either an appropriate policy is not configured or the match criteria is incorrect, then the **show security match-policies** command allows you to work offline and identify where the problem actually exists. It uses the search engine to identify the problem and thus enables you to use the appropriate match policy for the traffic.

The **result-count** option specifies how many policies to display. The first enabled policy in the list is the policy that is applied to all matching traffic. Other policies below it are "shadowed" by the first and are never encountered by matching traffic.

> ⓘ NOTE: The **show security match-policies** command is applicable only to security policies; IDP policies are not supported.

Options • destination-ip *destination-ip*—Destination IP address of the traffic.

• destination-port *destination-port*—Destination port number of the traffic. Range is 1 through 65,535

• from-zone *from-zone*—Name or ID of the source zone of the traffic.

• protocol *protocol-name* | *protocol-number*—Protocol name or numeric value of the traffic.

  • **ah** or **51**

  • **egp** or **8**

  • **esp** or **50**

  • **gre** or **47**

  • **icmp** or **1**

  • **igmp** or **2**

- **igp** or **9**

- **ipip** or **94**

- **ipv6** or **41**

- **ospf** or **89**

- **pgm** or **113**

- **pim** or **103**

- **rdp** or **27**

- **rsvp** or **46**

- **sctp** or **132**

- **tcp** or **6**

- **udp** or **17**

- **vrrp** or **112**

- **result-count** *number*—(Optional) The number of policy matches to display. Valid range is from 1 through 16. The default value is 1.

- **source-identity** *role-name*—Source identity of the traffic determined by the user role.

- **source-ip** *source-ip*—Source IP address of the traffic.

- **source-port** *source-port*—Source port number of the traffic. Range is 1 through 65,535.

- **to-zone** *to-zone*—Name or ID of the destination zone of the traffic.

Required Privilege Level   view

Related Documentation
- *clear security policies statistics*

- *Security Policies Overview*

- *Understanding Security Policy Rules*

- *Understanding Security Policy Elements*

List of Sample Output

Output Fields   lists the output fields for the **show security match-policies** command. Output fields are listed in the approximate order in which they appear.

Table 35: show security match-policies Output Fields

Field Name	Field Description
Policy	Name of the applicable policy.

Table 35: show security match-policies Output Fields *(continued)*

Field Name	Field Description
Action or **Action-type**	The action to be taken for traffic that matches the policy's match criteria. Actions include the following: • permit • firewall-authentication • tunnel ipsec-vpn *vpn-name* • pair-policy *pair-policy-name* • source-nat pool *pool-name* • pool-set *pool-set-name* • interface • destination-nat *name* • deny • reject
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	An internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, and 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, and 4.
From zone	Name of the source zone.
To zone	Name of the destination zone.
Source addresses	The names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	The names and corresponding IP addresses of the destination addresses (or address sets) for a policy as entered in the destination zone's address book. A packet's destination address must match one of these addresses for the policy to apply to it.
Application	Name of a preconfigured or custom application, or **any** if no application is specified.
IP protocol	Numeric value for the IP protocol used by the application, such as 6 for TCP or 1 for ICMP.
ALG	If an ALG is associated with the session, the name of the ALG. Otherwise, 0.
Inactivity timeout	Elapsed time without activity after which the application is terminated.
Source-port range	Range of matching source ports defined in the policy.

Table 35: show security match-policies Output Fields *(continued)*

Field Name	Field Description
Destination-port range	Range of matching destination ports defined in the policy.
Source identities	One or more user roles defined in the matching policy.

Sample Output

Example 1: show security match-policies

```
user@host> show security match-policies from-zone z1 to-zone z2 source-ip 10.10.10.1
destination-ip 192.0.2.1 source-port 1 destination-port 21 protocol tcp
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: z1, To zone: z2
  Source addresses:
    a2: 198.51.100.0/24
    a3: 10.10.10.1/32
  Destination addresses:
    d2: 203.0.113.0/24
    d3: 192.0.2.1/32
  Application: junos-ftp
    IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
      Source port range: [0-0]
      Destination port range: [21-21]
```

Example 2: show security match policies ... result-count

```
user@host> show security match-policies source-ip 10.10.10.1 destination-ip 192.0.2.5 source_port
1004 destination_port 80 protocol tcp result_count 5
Policy: p1, action-type: permit, State: enabled, Index: 4
  Sequence number: 1
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa1: 10.10.0.0/16
  Destination addresses:
    da5: 192.0.2.0/24
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
      Source port range: [1000-1030]
    Destination port range: [80-80]

Policy: p15, action-type: deny, State: enabled, Index: 18
  Sequence number: 15
  From zone: zone-A, To zone: zone-B
  Source addresses:
    sa11: 10.10.10.1/32
  Destination addresses:
    da15: 192.0.2.5/32
  Application: any
    IP protocol: 1, ALG: 0, Inactivity timeout: 0
      Source port range: [1000-1030]
    Destination port range: [80-80]
```

Example 3: show security match policies ... source-identity

```
user@host> show security match-policies from-zone untrust to-zone trust source-ip 10.10.10.1
destination-ip 192.0.2.1 destination_port 21 protocol 6 source-port 1234 source-identity role1
Policy: p1, action-type: permit, State: enabled, Index: 40
      Policy Type: Configured
      Sequence number: 1
      From zone: untrust, To zone: trust
      Source addresses:
        a1: 10.0.0.0/8
      Destination addresses:
        d1: 192.0.2.0/24
      Application: junos-ftp
        IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
          Source port range: [0-0]
          Destination port range: [21-21]
      Source identities: role1
      Per policy TCP Options: SYN check: No, SEQ check: No
```

show security nat destination rule

<table>
<tr>
<td>Syntax</td>
<td>show security nat destination rule
<i>rule-name</i>
all
logical-system (<i>logical-system-name</i> | all)
root-logical-system</td>
</tr>
<tr>
<td>Release Information</td>
<td>Command introduced in Junos OS Release 9.2. The Description output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the Successful sessions, Failed sessions, and Number of sessions output fields added in Junos OS Release 12.1X45-D10. Output for multiple destination ports and the application option field added in Junos OS Release 12.1X47-D10.</td>
</tr>
<tr>
<td>Description</td>
<td>Display information about the specified destination Network Address Translation (NAT) rule.</td>
</tr>
<tr>
<td>Options</td>
<td><i>rule-name</i>—Display information about the specified destination NAT rule.

all—Display information about all the destination NAT rules.

logical-system (<i>logical-system-name</i> | all)—Display information about the destination NAT rules for the specified logical system or for all logical systems.

root-logical-system—Display information about the destination NAT rules for the master (root) logical system.</td>
</tr>
<tr>
<td>Required Privilege Level</td>
<td>view</td>
</tr>
<tr>
<td>Related Documentation</td>
<td>• <i>rule (Security Destination NAT)</i></td>
</tr>
<tr>
<td>List of Sample Output</td>
<td>show security nat destination rule dst2-rule on page 425
show security nat destination rule all on page 426</td>
</tr>
<tr>
<td>Output Fields</td>
<td>Table 36 on page 424 lists the output fields for the show security nat destination rule command. Output fields are listed in the approximate order in which they appear.</td>
</tr>
</table>

Table 36: show security nat destination rule Output Fields

Field Name	Field Description
Total destination-nat rules	Number of destination NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly as address names and as address set names in the rule.
Destination NAT rule	Name of the destination NAT rule.
Description	Description of the destination NAT rule.

Table 36: show security nat destination rule Output Fields *(continued)*

Field Name	Field Description
Rule-Id	Rule identification number.
Rule position	Position of the destination NAT rule.
From routing instance	Name of the routing instance from which the packets flow.
From interface	Name of the interface from which the packets flow.
From zone	Name of the zone from which the packets flow.
Source addresses	Name of the source addresses that match the rule. The default value is any.
Destination addresses	Name of the destination addresses that match the rule. The default value is any.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: • **destination NAT pool**—Use user-defined destination NAT pool to perform destination NAT. • **off**—Do not perform destination NAT.
Destination ports	Destination ports number that match the rule. The default value is any.
Application	Indicates whether the application option is configured.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat destination rule dst2-rule

```
user@host>show security nat destination rule dst2-rule

 Destination NAT rule: dst2-rule              Rule-set: dst2
  Description                 : The destination rule dst2-rule is for the sales
team
  Rule-Id                     : 1
  Rule position               : 1
  From routing instance       : ri1
                              : ri2
  Match
    Source addresses          : add1
                                add2
    Destination addresses     : add9
 Action                       : off
```

```
   Destination port           : 0
Translation hits              : 68
   Successful sessions        : 25
   Failed sessions            : 43
Number of sessions            : 2
```

Sample Output

show security nat destination rule all

```
user@host> show security nat destination rule all

Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 2/0

Destination NAT rule: r4                    Rule-set: rs4
  Rule-Id                     : 2
  Rule position               : 2
  From zone                   : untrust
  Match
    Source addresses          : 192.0.2.0 - 192.0.2.255
    Destination addresses     : 198.51.100.0 - 198.51.100.255
    Application               : configured
  Action                      : off
  Translation hits            : 0
    Successful sessions       : 0
    Failed sessions           : 0
  Number of sessions          : 0
```

show security nat destination summary

Syntax	show security nat destination summary <logical-system (*logical-system-name* | all)> <root-logical-system>
Release Information	Command introduced in Junos OS Release 9.2. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.
Description	Display a summary of Network Address Translation (NAT) destination pool information.
Options	**none**—Display summary information about the destination NAT pool. **logical-system (*logical-system-name* | all)**—Display summary information about the destination NAT for the specified logical system or for all logical systems. **root-logical-system**—Display summary information about the destination NAT for the master (root) logical system.
Required Privilege Level	view
Related Documentation	• *pool (Security Destination NAT)* • *rule (Security Destination NAT)*
List of Sample Output	show security nat destination summary on page 428
Output Fields	Table 37 on page 427 lists the output fields for the **show security nat destination summary** command. Output fields are listed in the approximate order in which they appear.

Table 37: show security nat destination summary Output Fields

Field Name	Field Description
Total destination nat pool number	Number of destination NAT pools.
Pool name	Name of the destination address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
Port	Port number.
Total	Number of IP addresses that are in use.
Available	Number of IP addresses that are free for use.
Total destination nat rule number	Number of destination NAT rules.

Table 37: show security nat destination summary Output Fields *(continued)*

Field Name	Field Description
Total hit times	Number of times a translation in the translation table is used for all the destination NAT rules.
Total fail times	Number of times a translation in the translation table failed to translate for all the destination NAT rules.

Sample Output

show security nat destination summary

```
user@host> show security nat destination summary

Total pools: 2
Pool name          Address                   Routing      Port  Total
                   Range                     Instance           Address
dst-p1             203.0.113.1 -203.0.113.1    default      0     1
dst-p2             2001:db8::1 - 2001:db8::1   default      0     1

Total rules: 171
Rule name          Rule set      From                          Action
dst2-rule          dst2          ri1
                                 ri2
                                 ri3
                                 ri4
                                 ri5
                                 ri6
                                 ri7
dst3-rule          dst3          ri9                           off
                                 ri1
                                 ri2
                                 ri3
                                 ri4
                                 ri5

...
```

show security nat source rule

Syntax   show security nat source rule
rule-name
all
logical-system (*logical-system-name* | all)
root-logical-system

Release Information   Command introduced in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 11.2. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Source port**, **Successful sessions**, **Failed sessions**, and **Number of sessions** output fields added in Junos OS Release 12.1X45-D10. Output for multiple destination ports and the **application** output field added in Junos OS Release 12.1X47-D10.

Description   Display information about the specified source Network Address Translation (NAT) rule.

Options   *rule-name*—Name of the rule.

all—Display information about all the source NAT rules.

logical-system (*logical-system-name* | all)—Display information about the source NAT rules for the specified logical system or for all logical systems source NAT rules.

root-logical-system—Display information about the source NAT rules for the master (root) logical system.

Required Privilege Level   view

Related Documentation   • *rule (Security Source NAT)*

List of Sample Output   show security nat source rule r2 on page 431
show security nat source rule all on page 431

Output Fields   Table 38 on page 429 lists the output fields for the **show security nat source rule** command. Output fields are listed in the approximate order in which they appear

Table 38: show security nat source rule Output Fields

Field Name	Field Description
Source NAT rule	Name of the source NAT rule.
Total rules	Number of source NAT rules.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Description	Description of the source NAT rule.

Table 38: show security nat source rule Output Fields *(continued)*

Field Name	Field Description
Rule-Id	Rule identification number.
Rule position	Position of the source NAT rule.
From zone	Name of the zone from which the packets flow.
To zone	Name of the zone to which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
To routing instance	Name of the routing instance to which the packets flow.
From interface	Name of the interface from which the packets flow.
To interface	Name of the interface to which the packets flow.
Source addresses	Name of the source addresses that match the rule.
Source port	Source port numbers that match the rule.
Destination address	Name of the destination addresses that match the rule.
Destination ports	Destination port numbers that match the rule.
Application	Indicates whether the application option is configured.
Action	The action taken in regard to a packet that matches the rule's tuples. Actions include the following: • **off**—Do not perform source NAT. • **source NAT pool**—Use user-defined source NAT pool to perform source NAT • **interface**—Use egress interface's IP address to perform source NAT.
Persistent NAT type	Persistent NAT type.
Persistent NAT mapping type	Persistent NAT mapping type.
Inactivity timeout	Inactivity timeout for persistent NAT binding.
Max session number	Maximum number of sessions.
Translation hits	Number of translation hits.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.

Table 38: show security nat source rule Output Fields *(continued)*

Field Name	Field Description
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat source rule r2

```
user@host> show security nat source rule r2

source NAT rule: r2        Rule-set: src-nat
 Description              : The source rule r2 is for the sales team
 Rule-Id                 : 1
 Rule position           : 1
 From zone               : zone1
 To zone                 : zone9
 Match
  Source addresses       : add1
                           add2
  Destination addresses  : add9
                           add10
  Destination port       : 1002            - 1002
 Action                        : off
   Persistent NAT type         : N/A
   Persistent NAT mapping type : address-port-mapping
   Inactivity timeout          : 0
   Max session number          : 0
 Translation hits              : 4719
  Successful sessions    : 2000
  Failed sessions        : 2719
Number of sessions       : 5
```

Sample Output

show security nat source rule all

```
user@host> show security nat source rule all
Logical system: root
Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 3/0

source NAT rule: r2                    Rule-set: rs2
 Rule-Id                 : 2
 Rule position           : 1
 From zone               : trust
 To zone                 : untrust
 Match
   Source addresses      : 192.0.2.0 - 192.0.2.255
   Destination addresses : 203.0.113.0 - 203.0.113.255
                           198.51.100.0 - 198.51.100.255
   Application           : configured
 Action                        : off
   Persistent NAT type         : N/A
   Persistent NAT mapping type : address-port-mapping
   Inactivity timeout          : 0
   Max session number          : 0
 Translation hits              : 0
```

```
                        Successful sessions     : 0
                        Failed sessions         : 0
                Number of sessions              : 0
```

show security nat source summary

Syntax	show security nat source summary <logical-system (*logical-system-name* | all)> <root-logical-system>
Release Information	Command introduced in Junos OS Release 9.2. Support for IPv6 logical systems added in Junos OS Release 12.1X45-D10.
Description	Display a summary of Network Address Translation (NAT) source information.
Options	**none**—Display summary source NAT information. **logical-system (*logical-system-name* | all)**—Display summary information about the source NAT for the specified logical system or for all logical systems. **root-logical-system**—Display summary information about the source NAT for the master (root) logical system.
Required Privilege Level	view
Related Documentation	• *pool (Security Source NAT)* • *rule (Security Source NAT)*
List of Sample Output	show security nat source summary on page 434
Output Fields	Table 39 on page 433 lists the output fields for the **show security nat source summary** command. Output fields are listed in the approximate order in which they appear.

Table 39: show security nat source summary Output Fields

Field Name	Field Description
Total source nat pool number	Number of source NAT pools.
Pool name	Name of the source address pool.
Address range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
PAT	Whether Port Address Translation (PAT) is enabled (yes or no).
Total Address	Number of IP addresses that are in use.
Total source nat rule number	Number of source NAT rules.
Total port number usage for port translation pool	Number of ports assigned to the pool.

Table 39: show security nat source summary Output Fields *(continued)*

Field Name	Field Description
Maximum port number for port translation pool	Maximum number of NAT or PAT transactions done at any given time.

Sample Output

show security nat source summary

```
user@host> show security nat source summary logical-system all


Logical system: root-logical-system
Total port number usage for port translation pool: 67108864
Maximum port number for port translation pool: 134217728

Logical system: lsys1
Total port number usage for port translation pool: 193536
Maximum port number for port translation pool: 134217728
Total pools: 2

Logical system: root-logical-system
Pool                    Address                     Routing    PAT    Total
Name                    Range                       Instance          Address
pool1                   10.1.1.0-10.1.4.255-        default    yes    2048
                        10.1.5.0-10.1.8.255


Logical system: lsys1
Pool                    Address                     Routing    PAT    Total
Name                    Range                       Instance          Address
pool2                   203.0.113.1-203.0.113.3      default   yes    3


Total rules: 1

Logical system: root-logical-system
Rule name          Rule set      From        To            Action
rule 1             ruleset1      ge-2/2/2.0  ge-2/2/3.0    pool1
rule 1                           ge-2/2/4.0  ge-2/2/5.0
```

show security nat static rule

Syntax	show security nat static rule *rule-name* all logical-system (*logical-system-name* | all) root-logical-system
Release Information	Command introduced in Junos OS Release 9.3. The **Description** output field added in Junos OS Release 12.1. Support for IPv6 logical systems and the **Successful sessions**, **Failed sessions**, **Number of sessions**, **Source addresses,** and **Source ports** output fields added in Junos OS Release 12.1X45-D10. The **Destination NPTv6 addr** and **Destination NPTv6 Netmask** output fields added in Junos OS Release 12.3X48-D25.
Description	Display information about the specified static Network Address Translation (NAT) rule.
Options	*rule-name*—Name of the rule. **all**—Display information about all the static NAT rules. **logical-system (*logical-system-name* | all)**—Display information about the static NAT rules for the specified logical system or for all logical systems. **root-logical-system**—Display information about the static NAT rules for the master (root) logical system.
Required Privilege Level	view
Related Documentation	• *rule (Security Static NAT)*
List of Sample Output	show security nat static rule on page 436 show security nat static rule (IPv6) on page 437 show security nat static rule all on page 437
Output Fields	Table 40 on page 435 lists the output fields for the **show security nat static rule** command. Output fields are listed in the approximate order in which they appear.

Table 40: show security nat static rule Output Fields

Field Name	Field Description
Static NAT rule	Name of the static NAT rule.
Total referenced IPv4/IPv6 ip-prefixes	Number of IP prefixes referenced in source, destination, and static NAT rules. This total includes the IP prefixes configured directly, as address names, and as address set names in the rule.
Rule-set	Name of the rule set. Currently, you can configure 8 rules within the same rule set.
Description	Description of the static NAT rule.

Table 40: show security nat static rule Output Fields *(continued)*

Field Name	Field Description
Rule-Id	Rule identification number.
Rule position	Position of the rule that indicates the order in which it applies to traffic.
From interface	Name of the interface from which the packets flow.
From routing instance	Name of the routing instance from which the packets flow.
From zone	Name of the zone from which the packets flow.
Destination addresses	Name of the destination addresses that match the rule.
Destination NPTv6 addr	Destination address that matches the rule.
Source addresses	Name of the source addresses that match the rule.
Host addresses	Name of the host addresses that match the rule.
Netmask	Subnet IP address.
Destination NPTv6 Netmask	Subnet IPv6 address.
Host routing-instance	Name of the host routing instance.
Destination port	Destination port numbers that match the rule. The default value is any.
Source port	Source port numbers that match the rule.
Total static-nat rules	Number of static NAT rules.
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

show security nat static rule

```
user@host> show security nat static rule sta-r2

Static NAT rule: sta-r2                    Rule-set: sta-nat
 Description                : The static rule sta-r2 is for the sales team
 Rule-Id                    : 1
 Rule position             : 1
```

```
                        From zone                 : zone9
                        Destination addresses     : add3
                        Host addresses            : add4
                        Netmask                   : 24
                        Host routing-instance     : N/A
                        Translation hits          : 2
                         Successful sessions      : 2
                         Failed sessions          : 0
                      Number of sessions          : 2
```

Sample Output

show security nat static rule (IPv6)

```
user@host> show security nat static rule r1

Static NAT rule: r1                        Rule-set: rs1
   Rule-Id                   : 1
   Rule position             : 1
   From zone                 : trust
   Destination NPTv6 addr    : 2001:db8::
   Destination NPTv6 Netmask : 48
   Host addresses            : 2001:db8::3000
   Netmask                   : 48
   Host routing-instance     : N/A
   Translation hits          : 0
      Successful sessions    : 0
      Failed sessions        : 0
   Number of sessions        : 0
```

Sample Output

show security nat static rule all

```
user@host> show security nat static rule all

Static NAT rule: r1                        Rule-set: rs1
   Rule-Id                   : 1
   Rule position             : 1
   From zone                 : trust
   Source addresses          : 192.0.2.0 -192.0.2.3
                             : addr1
   Source ports              : 200 - 300
   Destination addresses     : 198.51.100.0
   Host addresses            : 203.0.113.0
   Netmask                   : 24
   Host routing-instance     : N/A
   Translation hits          : 4
    Successful sessions      : 4
    Failed sessions          : 0
 Number of sessions          : 4
Static NAT rule: r2                        Rule-set: rs1
   Rule-Id                   : 2
   Rule position             : 2
   From zone                 : trust
   Source addresses          : 192.0.2.0 -192.0.2.255
   Destination addresses     : 203.0.113.1
   Destination ports         : 100 - 200
   Host addresses            : 192.0.2.1
   Host ports                : 300 - 400
   Netmask                   : 32
```

```
                Host routing-instance     : N/A
                Translation hits          : 4
                 Successful sessions      : 4
                 Failed sessions          : 0
            Number of sessions            : 4
```

show security policies

Syntax   show security policies
<detail>
<none>
policy-name *policy-name*
<detail>
<global>

Release Information   Command modified in Junos OS Release 9.2. Support for IPv6 addresses added in Junos OS Release 10.2. Support for IPv6 addresses in active/active chassis cluster configurations in addition to the existing support of active/passive chassis cluster configurations added in Junos OS Release 10.4. Support for wildcard addresses added in Junos OS Release 11.1. Support for global policy added in Junos OS Release 11.4. Support for services offloading added in Junos OS Release 11.4. Support for source-identities added in Junos OS Release 12.1. The **Description** output field added in Junos OS Release 12.1. Support for negated address added in Junos OS Release 12.1X45-D10. The output fields for Policy Statistics expanded, and the output fields for the **global** and **policy-name** options expanded to include from-zone and to-zone global match criteria in Junos OS Release 12.1X47-D10. Support for the **initial-tcp-mss** and **reverse-tcp-mss** options added in Junos OS Release 12.3X48-D20. Output field and description for **source-end-user-profile** option added in Junos OS Release 15.1x49-D70.

Description   Display a summary of all security policies configured on the device. If a particular policy is specified, display information particular to that policy.

Options   • **none**—Display basic information about all configured policies.

• **detail**—(Optional) Display a detailed view of all of the policies configured on the device.

• **policy-name** *policy-name*—(Optional) Display information about the specified policy.

• **global**—Display information about global policies.

Required Privilege Level   view

Related Documentation   • *Security Policies Overview*

• *Understanding Security Policy Rules*

• *Understanding Security Policy Elements*

List of Sample Output

Output Fields   Table 41 on page 440 lists the output fields for the **show security policies** command. Output fields are listed in the approximate order in which they appear.

Table 41: show security policies Output Fields

Field Name	Field Description
From zone	Name of the source zone.
To zone	Name of the destination zone.
Policy	Name of the applicable policy.
Description	Description of the applicable policy.
State	Status of the policy: • **enabled**: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • **disabled**: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1, 2, 3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1, 2, 3, 4.
Source addresses	For standard display mode, the names of the source addresses for a policy. Address sets are resolved to their individual names. For detail display mode, the names and corresponding IP addresses of the source addresses for a policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
source-end-user-profile	Name of the device identity profile (also referred to as the end user profile) that contains attributes, or characteristics of a device. Specification of the device identity profile in the **source-end-user-profile** field is part of the device identity feature. If a device matches the attributes specified in the profile and other security policy parameters, then the security policy's action is applied to traffic issuing from the device.
Source addresses (excluded)	Name of the source address excluded from the policy.
Destination addresses (excluded)	Name of the destination address excluded from the policy.
Source identities	One or more user roles specified for a policy.

Table 41: show security policies Output Fields *(continued)*

Field Name	Field Description
Applications	Name of a preconfigured or custom application whose type the packet matches, as specified at configuration time. • **IP protocol**: The Internet protocol used by the application—for example, TCP, UDP, ICMP. • **ALG**: If an ALG is explicitly associated with the policy, the name of the ALG is displayed. If **application-protocol ignore** is configured, ignore is displayed. Otherwise, 0 is displayed. However, even if this command shows ALG: 0, ALGs might be triggered for packets destined to well-known ports on which ALGs are listening, unless ALGs are explicitly disabled or when **application-protocol ignore** is not configured for custom applications. • **Inactivity timeout**: Elapsed time without activity after which the application is terminated. • **Source port range**: The low-high source port range for the session application.
Destination Address Translation	Status of the destination address translation traffic: • **drop translated**—Drop the packets with translated destination addresses. • **drop untranslated**—Drop the packets without translated destination addresses.

Application Firewall	An application firewall includes the following:
	
	• **Rule-set**—Name of the rule set.
	• **Rule**—Name of the rule.
	  ◦ **Dynamic applications**—Name of the applications.
	  ◦ **Dynamic application groups**—Name of the application groups.
	  ◦ **Action**—The action taken with respect to a packet that matches the application firewall rule set. Actions include the following:
	    ◦ **permit**
	    ◦ **deny**
	• **Default rule**—The default rule applied when the identified application is not specified in any rules of the rule set.

Action or Action-type	• The action taken in regard to a packet that matches the policy's tuples. Actions include the following:
	  ◦ **permit**
	  ◦ **firewall-authentication**
	  ◦ **tunnel ipsec-vpn** *vpn-name*
	  ◦ **pair-policy** *pair-policy-name*
	  ◦ **source-nat pool** *pool-name*
	  ◦ **pool-set** *pool-set-name*
	  ◦ **interface**
	  ◦ **destination-nat** *name*
	  ◦ **deny**
	  ◦ **reject**
	  ◦ **services-offload**

Session log	Session log entry that indicates whether the **at-create** and **at-close** flags were set at configuration time to log session information.

Table 41: show security policies Output Fields *(continued)*

Field Name	Field Description
Scheduler name	Name of a preconfigured scheduler whose schedule determines when the policy is active and can be used as a possible match for traffic.

Policy statistics	• **Input bytes**—The total number of bytes presented for processing by the device.
	  • **Initial direction**—The number of bytes presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of bytes presented for processing by the device from the reply direction.
	• **Output bytes**—The total number of bytes actually processed by the device.
	  • **Initial direction**—The number of bytes from the initial direction actually processed by the device.
	  • **Reply direction**—The number of bytes from the reply direction actually processed by the device.
	• **Input packets**—The total number of packets presented for processing by the device.
	  • **Initial direction**—The number of packets presented for processing by the device from the initial direction.
	  • **Reply direction**—The number of packets presented for processing by the device from the reply direction.
	• **Output packets**—The total number of packets actually processed by the device.
	  • **Initial direction**—The number of packets actually processed by the device from the initial direction.
	  • **Reply direction**—The number of packets actually processed by the device from the reply direction.
	• **Session rate**—The total number of active and deleted sessions.
	• **Active sessions**—The number of sessions currently present because of access control lookups that used this policy.
	• **Session deletions**—The number of sessions deleted since system startup.
	• **Policy lookups**—The number of times the policy was accessed to check for a match.
	
	NOTE: Configure the Policy P1 with the **count** option to display policy statistics.

Per policy TCP Options	Configured sync and sequence checks, and the configured TCP MSS value for the initial direction and /or the reverse direction.

Sample Output

show security policies

```
user@host> show security policies
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Sequence number: 1
    Source addresses:
    sa-1-ipv4: 198.51.100.11/24
    sa-2-ipv6: 2001:db8:a0b:12f0::1/32
    sa-3-ipv6: 2001:db8:a0b:12f0::22/32
    sa-4-wc:   203.0.113.1/255.255.0.255
    Destination addresses:
    da-1-ipv4: 2.2.2.2/24
    da-2-ipv6: 2001:db8:a0b:12f0::8/32
```

```
          da-3-ipv6: 2001:db8:a0b:12f0::9/32
          da-4-wc:   192.168.22.11/255.255.0.255
          Source identities: role1, role2, role4
          Applications: any
          Action: permit, application services, log, scheduled
          Application firewall : my_ruleset1
        Policy: p2, State: enabled, Index: 5, Sequence number: 2
          Source addresses:
          sa-1-ipv4:  198.51.100.11/24
          sa-2-ipv6:  2001:db8:a0b:12f0::1/32
          sa-3-ipv6:  2001:db8:a0b:12f0::22/32
          Destination addresses:
          da-1-ipv4:  2.2.2.2/24
          da-2-ipv6:  2001:db8:a0b:12f0::1/32
          da-3-ipv6:  2001:db8:a0b:12f0::9/32
          Source identities: role1, role4
          Applications: any
          Action: deny, scheduled
```

show security policies policy-name p1 detail

```
user@host> show security policies policy-name p1 detail
Policy: p1, action-type: permit, State: enabled, Index: 4
  Description: The policy p1 is for the sales team
  Sequence number: 1
  From zone: trust, To zone: untrust
  Source addresses:
    sa-1-ipv4:  198.51.100.11/24
    sa-2-ipv6:  2001:db8:a0b:12f0::1/32
    sa-3-ipv6:  2001:db8:a0b:12f0::9/32
    sa-4-wc:    203.0.113.1/255.255.0.255
  Destination addresses:
    da-1-ipv4: 2.2.2.2/24
    da-2-ipv6: 2001:db8:a0b:12f0::1/32
    da-3-ipv6: 2001:db8:a0b:12f0::9/32
    da-4-wc:   192.168.22.11/255.255.0.255
 Source identities:
 role1
 role2
 role4
 Application: any
   IP protocol: 0, ALG: 0, Inactivity timeout: 0
     Source port range: [0-0]
     Destination port range: [0-0]
 Destination Address Translation: drop translated
 Application firewall :
 Rule-set: my_ruleset1
   Rule: rule1
     Dynamic Applications: junos:FACEBOOK, junos:YSMG
     Dynamic Application groups: junos:web, junos:chat
     Action: deny
   Default rule: permit
 Session log: at-create, at-close
 Scheduler name: sch20
 Per policy TCP Options: SYN check: No, SEQ check: No
 Policy statistics:
   Input  bytes      :                18144              545 bps
     Initial direction:                9072              272 bps
     Reply direction  :                9072              272 bps
   Output bytes       :                18144              545 bps
     Initial direction:                9072              272 bps
```

```
            Reply direction  :                9072              272 bps
      Input  packets    :                 216                6 pps
         Initial direction:                 108                3 bps
         Reply direction  :                 108                3 bps
      Output packets    :                 216                6 pps
         Initial direction:                 108                3 bps
         Reply direction  :                 108                3 bps
      Session rate      :                 108                3 sps
      Active sessions    :                  93
      Session deletions  :                  15
      Policy lookups    :                 108
```

show security policies (services-offload)

```
user@host> show security policies
Default policy: deny-all
From zone: trust, To zone: untrust
  Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload, count
From zone: untrust, To zone: trust
  Policy: p2, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
    Source addresses: any
    Destination addresses: any
    Source identities: role1, role2, role4
    Applications: any
    Action: permit, services-offload
```

show security policies (device identity)

```
user@host > show security policies
From zone: trust, To zone: untrust
        Policy: dev-id-marketing, State: enabled, Index: 5, Scope Policy: 0,
Sequence number: 1
        Source addresses: any
        Destination addresses: any
        source-end-user-profile: marketing-profile
        Applications: any
        Action: permit
```

show security policies detail

```
user@host> show security policies detail
Default policy: deny-all
Policy: p1, action-type: permit, services-offload:enabled , State: enabled, Index:
 4, Scope Policy: 0
 Policy Type: Configured
 Description: The policy p1 is for the sales team
 Sequence number: 1
 From zone: trust, To zone: untrust
 Source addresses:
   any-ipv4(global):  0.0.0.0/0
   any-ipv6(global): ::/0
 Destination addresses:
   any-ipv4(global):  0.0.0.0/0
   any-ipv6(global): ::/0
 Source identities:
 role1
 role2
```

```
                    role4
                    Application: any
                      IP protocol: 0, ALG: 0, Inactivity timeout: 0
                        Source port range: [0-0]
                        Destination port range: [0-0]
                    Per policy TCP Options: SYN check: No, SEQ check: No
                    Policy statistics:
                      Input  bytes      :                18144              545 bps
                        Initial direction:                9072              272 bps
                        Reply direction  :                9072              272 bps
                      Output bytes      :                18144              545 bps
                        Initial direction:                9072              272 bps
                        Reply direction  :                9072              272 bps
                      Input  packets    :                  216                6 pps
                        Initial direction:                  108                3 bps
                        Reply direction  :                  108                3 bps
                      Output packets    :                  216                6 pps
                        Initial direction:                  108                3 bps
                        Reply direction  :                  108                3 bps
                      Session rate      :                  108                3 sps
                      Active sessions   :                   93
                      Session deletions :                   15
                      Policy lookups    :                  108
                  Policy: p2, action-type: permit, services-offload:enabled , State: enabled, Index:
                   5, Scope Policy: 0
                    Policy Type: Configured
                    Description: The policy p2 is for the sales team
                    Sequence number: 1
                    From zone: untrust, To zone: trust
                    Source addresses:
                      any-ipv4(global):  0.0.0.0/0
                      any-ipv6(global): ::/0
                    Destination addresses:
                      any-ipv4(global):  0.0.0.0/0
                      any-ipv6(global): ::/0
                   Source identities:
                    role1
                    role2
                    role4
                    Application: any
                      IP protocol: 0, ALG: 0, Inactivity timeout: 0
                        Source port range: [0-0]
                        Destination port range: [0-0]
                    Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies detail (TCP Options)

```
user@host> show security policies policy-name policy1 detail
node0:
--------------------------------------------------------------------------
Policy: policy1, action-type: permit, State: enabled, Index: 7, Scope Policy: 0
  Policy Type: Configured
  Sequence number: 2
  From zone: trust, To zone: untrust
  Source addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
  Destination addresses:
    any-ipv4(global):  0.0.0.0/0
    any-ipv6(global): ::/0
  Application: any
```

```
                            IP protocol: 0, ALG: 0, Inactivity timeout: 0
                              Source port range: [0-0]
                              Destination port range: [0-0]
                          Per policy TCP Options: SYN check: No, SEQ check: No
                          Per policy TCP MSS: initial: 800, reverse: 900
```

show security policies policy-name p1 (Negated Address)

```
            user@host>show security policies policy-name p1
            node0:
            --------------------------------------------------------------------------
            From zone: trust, To zone: untrust
              Policy: p1, State: enabled, Index: 4, Scope Policy: 0, Sequence number: 1
                Source addresses(excluded): as1
                Destination addresses(excluded): as2
                Applications: any
                Action: permit
```

show security policies policy-name p1 detail (Negated Address)

```
            user@host>show security policies policy-name p1 detail
            node0:
            --------------------------------------------------------------------------
            Policy: p1, action-type: permit, State: enabled, Index: 4, Scope Policy: 0
              Policy Type: Configured
              Sequence number: 1
              From zone: trust, To zone: untrust
              Source addresses(excluded):
                ad1(ad): 255.255.255.255/32
                ad2(ad):  198.51.100.1/24
                ad3(ad):  198.51.100.6  ~ 198.51.100.56
                ad4(ad):  192.0.2.8/24
                ad5(ad):  198.51.100.99 ~ 198.51.100.199
                ad6(ad):  203.0.113.9/24
                ad7(ad):  203.0.113.23/24
              Destination addresses(excluded):
                ad13(ad2): 198.51.100.76/24
                ad12(ad2): 198.51.100.88/24
                ad11(ad2): 192.0.2.23 ~ 192.0.2.66
                ad10(ad2): 192.0.2.93
                ad9(ad2):  203.0.113.76 ~ 203.0.113.106
                ad8(ad2):  203.0.113.199
              Application: any
                IP protocol: 0, ALG: 0, Inactivity timeout: 0
                  Source port range: [0-0]
                  Destination port range: [0-0]
              Per policy TCP Options: SYN check: No, SEQ check: No
```

show security policies global

```
            user@host>show security policies global policy-name Pa
            node0:
            --------------------------------------------------------------------------
              Global policies:
              Policy: Pa, State: enabled, Index: 5, Scope Policy: 0, Sequence number: 1
                From zones: zone1, zone2
                To zones: zone3, zone4    Source addresses: any
                Destination addresses: any
                Applications: any
                Action: permit
```

show security screen statistics

Syntax   show security screen statistics (zone *zone-name* | interface *interface-name*)
<logical-system (*logical-system-name* | all)>
<node (*node-id* | all | local | primary)>
<root-logical-system>

Release Information   Command introduced in Junos OS Release 8.5. **node** options added in Junos OS Release 9.0. **logical-system all** option added in Junos OS Release 11.2R6. Support for IPv6 extension header screens added in Junos OS Release 12.1X46-D10.

Description   Display intrusion detection service (IDS) security screen statistics.

Options

- **zone** *zone-name* —Display screen statistics for this security zone.

- **interface** *interface-name*—Display screen statistics for this interface.

- **logical-system**—(Optional) Display screen statistics for configured logical systems.

  - *logical-system-name*—Display screen statistics for the named logical system.

  - **all**—Display screen statistics for all logical systems, including the master (root) logical system.

- **node**—(Optional) For chassis cluster configurations, display screen statistics on a specific node.

  - *node-id*—Identification number of a node. It can be 0 or 1.

  - **all**—Display information about all nodes.

  - **local**—Display information about the local node.

  - **primary**—Display information about the primary node.

- **root-logical-system**—(Optional) Display screen statistics for the master logical system only.

Required Privilege Level   view

Related Documentation

- *clear security screen statistics*

- *clear security screen statistics interface*

- *clear security screen statistics zone*

- *Example: Configuring Multiple Screening Options*

List of Sample Output

Output Fields   Table 42 on page 449 lists the output fields for the **show security screen statistics** command. Output fields are listed in the approximate order in which they appear.

Table 42: show security screen statistics Output Fields

Field Name	Field Description
ICMP flood	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.
UDP flood	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP port scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP address sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
IP tear drop	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN flood	Number of TCP SYN attacks.
IP spoofing	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP ping of death	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP source route option	Number of IP source route attacks.
TCP address sweep	Number of TCP address sweeps.
TCP land attack	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN fragment	Number of TCP SYN fragments.
TCP no flag	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
IP unknown protocol	Number of IPs.
IP bad options	Number of invalid options.
IP record route option	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.

Table 42: show security screen statistics Output Fields *(continued)*

Field Name	Field Description
IP timestamp option	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP security option	Number of IP security option attacks.
IP loose source route option	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
IP strict source route option	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
IP stream option	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP fragment	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP large packet	Number of large ICMP packets.
TCP SYN FIN	Number of TCP SYN FIN packets.
TCP FIN no ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.
Source session limit	Number of concurrent sessions that can be initiated from a source IP address.
TCP SYN-ACK-ACK proxy	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series devices running Junos OS reject further connection requests from that IP address.
IP block fragment	Number of IP block fragments.
Destination session limit	Number of concurrent sessions that can be directed to a single destination IP address.
UDP address sweep	Number of UDP address sweeps.
IPv6 extension header	Number of packets filtered for the defined IPv6 extension headers.
IPv6 extension hop by hop option	Number of packets filtered for the defined IPv6 hop-by-hop option types.
IPv6 extension destination option	Number of packets filtered for the defined IPv6 destination option types.
IPv6 extension header limit	Number of packets filtered for crossing the defined IPv6 extension header limit.
IPv6 malformed header	Number of IPv6 malformed headers defined for the intrusion detection service (IDS).

Table 42: show security screen statistics Output Fields *(continued)*

ICMPv6 malformed packet	Number of ICMPv6 malformed packets defined for the IDS options.

Sample Output

show security screen statistics zone scrzone

```
user@host> show security screen statistics zone scrzone
Screen statistics:
IDS attack type                         Statistics
  ICMP flood                            0
  UDP flood                             0
  TCP winnuke                           0
  TCP port scan                         91
  ICMP address sweep                    0
  TCP sweep                             0
  UDP sweep                             0
  IP tear drop                          0
  TCP SYN flood                         0
  IP spoofing                           0
  ICMP ping of death                    0
  IP source route option                0
  TCP land attack                       0
  TCP SYN fragment                      0
  TCP no flag                           0
  IP unknown protocol                   0
  IP bad options                        0
  IP record route option                0
  IP timestamp option                   0
  IP security option                    0
  IP loose source route option          0
  IP strict source route option         0
  IP stream option                      0
  ICMP fragment                         0
  ICMP large packet                     0
  TCP SYN FIN                           0
  TCP FIN no ACK                        0
  Source session limit                  0
  TCP SYN-ACK-ACK proxy                 0
  IP block fragment                     0
  Destination session limit             0
```

Sample Output

show security screen statistics zone untrust (IPv6)

```
user@host>show security screen statistics zone untrust
Screen statistics:

IDS attack type                         Statistics
  ICMP flood                                  0
  UDP  flood                                  0
  TCP  winnuke                                0
  ……
  IPv6   extension header                     0
  IPv6   extension hop by hop option          0
  IPv6   extension destination option         0
  IPv6   extension header limit               0
  IPv6   malformed header                     0
```

```
          ICMPv6 malformed packet                    0
```

Sample Output

show security screen statistics interface ge-0/0/3

```
user@host> show security screen statistics interface ge-0/0/3
Screen statistics:
IDS attack type                          Statistics
  ICMP flood                             0
  UDP flood                              0
  TCP winnuke                            0
  TCP port scan                          91
  ICMP address sweep                     0
  TCP sweep                              0
  UDP sweep                              0
  IP tear drop                           0
  TCP SYN flood                          0
  IP spoofing                            0
  ICMP ping of death                     0
  IP source route option                 0
  TCP land attack                        0
  TCP SYN fragment                       0
  TCP no flag                            0
  IP unknown protocol                    0
  IP bad options                         0
  IP record route option                 0
  IP timestamp option                    0
  IP security option                     0
  IP loose source route option           0
  IP strict source route option          0
  IP stream option                       0
  ICMP fragment                          0
  ICMP large packet                      0
  TCP SYN FIN                            0
  TCP FIN no ACK                         0
  Source session limit                   0
  TCP SYN-ACK-ACK proxy                  0
  IP block fragment                      0
  Destination session limit              0
```

Sample Output

show security screen statistics interface ge-0/0/1 (IPv6)

```
user@host> show security screen statistics interface ge-0/0/1

Screen statistics:

IDS attack type                          Statistics
  ICMP flood                                    0
  UDP flood                                     0
  ……
  IPv6 extension header                         0
  IPv6 extension hop by hop option              0
  IPv6 extension destination option             0
  IPv6 extension header limit                   0
  IPv6 malformed header                         0
  ICMPv6 malformed packet                       0
```

Sample Output

show security screen statistics interface ge-0/0/1 node primary

```
user@host> show security screen statistics interface ge-0/0/1 node primary
node0:
--------------------------------------------------------------------------
Screen statistics:
IDS attack type                             Statistics
  ICMP flood                                1
  UDP flood                                 1
  TCP winnuke                               1
  TCP port scan                             1
  ICMP address sweep                        1
  TCP sweep                                 1
  UDP sweep                                 1
  IP tear drop                              1
  TCP SYN flood                             1
  IP spoofing                               1
  ICMP ping of death                        1
  IP source route option                    1
  TCP land attack                           1
  TCP SYN fragment                          1
  TCP no flag                               1
  IP unknown protocol                       1
  IP bad options                            1
  IP record route option                    1
  IP timestamp option                       1
  IP security option                        1
  IP loose source route option             1
  IP strict source route option            1
  IP stream option                          1
  ICMP fragment                             1
  ICMP large packet                         1
  TCP SYN FIN                               1
  TCP FIN no ACK                            1
  Source session limit                      1
  TCP SYN-ACK-ACK proxy                     1
  IP block fragment                         1
  Destination session limit                 1
```

Sample Output

show security screen statistics zone trust logical-system all

```
user@host> show security screen statistics zone trust logical-system all
Logical system: root-logical-system
Screen statistics:

IDS attack type                             Statistics
  ICMP flood                                0
  UDP flood                                 0
  TCP winnuke                               0
  TCP port scan                             0
  ICMP address sweep                        0
  TCP sweep                                 0
  UDP sweep                                 0
  IP tear drop                              0
  TCP SYN flood                             0
  IP spoofing                               0
  ICMP ping of death                        0
```

```
                    IP source route option                  0
                    TCP land attack                         0
                    TCP SYN fragment                        0
                    TCP no flag                             0
                    IP unknown protocol                     0
                    IP bad options                          0
                    IP record route option                  0
                    IP timestamp option                     0
                    IP security option                      0
                    IP loose source route option            0
                    IP strict source route option           0
                    IP stream option                        0
                    ICMP fragment                           0
                    ICMP large packet                       0
                    TCP SYN FIN                              0
                    TCP FIN no ACK                           0
                    Source session limit                    0
                    TCP SYN-ACK-ACK proxy                    0
                    IP block fragment                       0
                    Destination session limit               0

                Logical system: ls1
                Screen statistics:

                IDS attack type                         Statistics
                    ICMP flood                              0
                    UDP flood                               0
                    TCP winnuke                             0
                    TCP port scan                           0
                    ICMP address sweep                      0
                    TCP sweep                               0
                    UDP sweep                               0
                    IP tear drop                            0
                    TCP SYN flood                           0
                    IP spoofing                             0
                    ICMP ping of death                      0
                    IP source route option                  0
                    TCP land attack                         0
                    TCP SYN fragment                        0
                    TCP no flag                             0
                    IP unknown protocol                     0
                    IP bad options                          0
                    IP record route option                  0
                    IP timestamp option                     0
                    IP security option                      0
                    IP loose source route option            0
                    IP strict source route option           0
                    IP stream option                        0
                    ICMP fragment                           0
                    ICMP large packet                       0
                    TCP SYN FIN                              0
                    TCP FIN no ACK                          0
                    Source session limit                    0
                    TCP SYN-ACK-ACK proxy                    0
                    IP block fragment                       0
                    Destination session limit               0

                Logical system: ls2
                Screen statistics:

                IDS attack type                         Statistics
```

```
                    ICMP flood                              0
                    UDP flood                               0
                    TCP winnuke                             0
                    TCP port scan                           0
                    ICMP address sweep                      0
                    TCP sweep                               0
                    UDP sweep                               0
                    IP tear drop                            0
                    TCP SYN flood                           0
                    IP spoofing                             0
                    ICMP ping of death                      0
                    IP source route option                  0
                    TCP land attack                         0
                    TCP SYN fragment                        0
                    TCP no flag                             0
                    IP unknown protocol                     0
                    IP bad options                          0
                    IP record route option                  0
                    IP timestamp option                     0
                    IP security option                      0
                    IP loose source route option            0
                    IP strict source route option           0
                    IP stream option                        0
                    ICMP fragment                           0
                    ICMP large packet                       0
                    TCP SYN FIN                              0
                    TCP FIN no ACK                           0
                    Source session limit                    0
                    TCP SYN-ACK-ACK proxy                    0
                    IP block fragment                        0
                    Destination session limit               0
```

show system security-profile

Syntax
show system security-profile (all-resource | *resource*) <detail | terse> <logical-system (all | *logical-system-name*)> <root-logical-system> <summary>

Release Information
Command introduced in Junos OS Release 11.2. Support for application firewall added in Junos OS Release 11.3. Option to display all resources for a logical system added in Junos OS Release 11.. Resource information for ports in source NAT pools with port translation added in Release Junos OS 11.4.

Description
Display information about a resource allocated to the logical system in a security profile. For each resource specified, the number used by the logical system and the configured maximum and reserved values are displayed.

This command can be used by the master administrator to display resource information for the master logical system or user logical system. This command can also be used by the user logical system administrator to display resource information for a user logical system.

Options
Either specify **all-resource** to display information about all resources allocated for the logical system, or specify one of the following resources:

- address-book—Address books.

- appfw-rule-set—Application firewall rule set entries.

- appfw-rule—Application firewall rule entries.

- auth-entry—Firewall authentication entries.

- cpu—CPU utilization.

- flow-gate—Flow gates, also known as pinholes.

- flow-session—Flow sessions.

- nat-cone-binding—Network Address Translation (NAT) cone bindings.

- nat-destination-pool—NAT destination pools.

- nat-destination-rule—NAT destination rules.

- nat-nopat-address—NAT without port address translations.

- nat-pat-address—NAT with port address translations.

- nat-pat-portnum—NAT source port numbers for port translation

- nat-port-ol-ipnumber—NAT port overloading IP numbers.

- nat-rule-referenced-prefix—NAT rule referenced IP-prefixes.

- nat-source-pool—NAT source pools.

- nat-source-rule—NAT source rules.

- nat-static-rule—NAT static rules.

- policy—Security policies.

- policy-with-count—Security policies with a count.

- scheduler—Schedulers.

- zone—Security zones.

detail | terse—(Optional) Display the specified level of output.

The following options are available only to the master administrator:

- logical-system—Display resource information for a specified user logical system. Specify **all** to display resource information for all logical systems, including the master logical system.

- root-logical-system—Display resource information for the master (root) logical system.

- summary—Display summary information about the resource for all logical systems.

Required Privilege Level
view

Related Documentation
- security-profile-resources on page 358

List of Sample Output
show system security-profile all-resource on page 458
show system security-profile policy on page 458
show system security-profile cpu on page 458
show system security-profile cpu logical-system all on page 459
show system security-profile cpu summary on page 459
show system security-profile nat-pat-portnum on page 459
show system security-profile nat-pat-portnum summary on page 460

Output Fields
Table 43 on page 457 lists the output fields for the **show system security-profile** command. Output fields are listed in the approximate order in which they appear.

Table 43: show system security-profile Output Fields

Field Name	Field Description
logical system name	Name of the logical system.
security profile name	Name of the security profile bound to the logical system.
usage	Number of resources that are currently being used by the logical system.
reserved	Number of resources that are guaranteed to be available to the logical system.
maximum	Number of resources that the logical system can use. The maximum does not guarantee that the amount specified for the resource in the security profile is available. The maximum is not applicable for CPU resources.
CPU control	**TRUE** if CPU control is enabled or **FALSE** if CPU control is not enabled.

Table 43: show system security-profile Output Fields *(continued)*

Field Name	Field Description
CPU control target	Upper limit for CPU utilization on the device. The default value is 80 percent.
CPU name	Central point (CP) or services processing unit (SPU). CP utilization and average utilization of all SPUs is shown. The **detail** option shows CPU utilization on each SPU.
drop rate	Packets dropped for CPU control.

Sample Output

show system security-profile all-resource

```
user@host> show system security-profile all-resource

resource                              usage    reserved    maximum

[logical system name:    root-logical-system]
[security profile name: Default-Profile]
address-book                              0           0        512
auth-entry                                0           0 2147483647
cpu on CP                             0.00%       1.00%     80.00%
cpu on SPU                            0.00%       1.00%     80.00%
flow-gate                                 0           0     524288
flow-session                              2           0    6291456
nat-cone-binding                          0           0      65536
nat-destination-pool                      0           0       4096
nat-destination-rule                      0           0       8192
nat-nopat-address                         0           0    1048576
nat-pat-address                           0           0       2048
nat-port-ol-ipnumber                      0           0          4
nat-rule-referenced-prefix                0           0    1048576
nat-source-pool                           0           0       2048
nat-source-rule                           0           0       8192
nat-static-rule                           0           0      20480
policy                                    0           0      40000
policy-with-count                         0           0       1024
scheduler                                 0           0         64
zone                                      0           0        512
```

show system security-profile policy

```
user@host> show system security-profile policy
logical system name    security profile name    usage    reserved    maximum

ls-product-design      ls-design-profile            0          40         50
```

show system security-profile cpu

```
user@host> show system security-profile cpu
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)
drop rate(%)
root-logical-system    Default-Profile CP            0.00%          1.00%
    0.00%
```

```
                    root-logical-system   Default-Profile SPU            0.00%           1.00%
                         0.00%
```

show system security-profile cpu logical-system all

```
user@host> show system security-profile cpu logical-system all
CPU control: TRUE
CPU control target: 80.00%
logical system name    profile name    CPU name    usage(%)    reserved(%)
drop rate(%)
root-logical-system    Default-Profile CP            0.00%           1.00%
     0.00%
root-logical-system    Default-Profile SPU           0.00%           1.00%
     0.00%
ls-product-design      ls-design-profile CP          0.00%           0.00%
     0.00%
ls-product-design      ls-design-profile SPU         0.00%           0.00%
     0.00%
ls-marketing-dept      ls-acct-mrkt-profile CP       0.00%           0.00%
     0.00%
ls-marketing-dept      ls-acct-mrkt-profile SPU      0.00%           0.00%
     0.00%
```

Should the above output actually look as follows?

```
logical system name    security profile name    usage     reserved    maximum
root-logical-system    Default-Profile          67108864     0         134217728
lsys1                  profile1                  193536       6000      134217728
```

show system security-profile cpu summary

```
user@host> show system security-profile cpu summary
CPU control: TRUE
CPU control target: 80.00%

CPU type               :      CP
global used amount     :   0.00%
global maximum quota   : 80.00%
global available amount : 80.00%
total logical systems  :      3
total security profiles :      3
heaviest usage / user  :   0.00%       / root-logical-system
lightest usage / user  :   0.00%       / root-logical-system

CPU type               :      SPU
global used amount     :   0.00%
global maximum quota   : 80.00%
global available amount : 80.00%
total logical systems  :      3
total security profiles :      3
heaviest usage / user  :   0.00%       / root-logical-system
lightest usage / user  :   0.00%       / root-logical-system
```

show system security-profile nat-pat-portnum

```
user@host> show system security-profile cpu nat-pat-portnum
CPU control: TRUE
CPU control target: 80.00%
logical system name    security profile name        usage(%)     reserved(%)
maximum
root-logical-system    Default-Profile CP            67108864          0
     134217728
```

show system security-profile nat-pat-portnum summary

```
user@host> show system security-profile nat-pat-portnum summary
global used amount     :67302400
global maximum quota     :134217728
global available amount   :66915328
total logical systems    :2
total security profiles   :1
heaviest usage / user     :193536 / lsys1
```

show security softwires

Syntax   show security softwires <softwire-name *softwire-name*>
<logical-system (all | *logical-system-name*)>

Release Information   Command introduced in Junos OS Release 10.4. The **logical-system** option introduced in Junos OS Release 12.1.

Description   Display a summary of information of all the softwire concentrators and details on concentrators with specified name.

Options   **softwire-name** *softwire-name*—Display the details of the specified softwire concentrator.

logical-system (all | *logical-system-name*)—Display softwire information for all logical systems or for a specified logical system. This option is only available to the master administrator.

Required Privilege   view
Level

Related   • *Juniper Networks Devices Processing Overview*
Documentation

Sample Output

```
user@host> show security softwires
Softwire Name            SC Address          Status    Number of SI connected
  SC-CSSI-1                3001::1             Connected   2
  SC-CSSI-str00            3100::1             Active      0
  SC-CSSI-str01            3101::1             Inactive    0
  SC-CSSI-str02            3001::1             Connected   2520

user@host> show security softwires softwire-name SC-CSSI-1
Name of softwire: SC-CSSI-1
  SC status: Connected
  SC address: 3001::1
  Zone: trust
  VR ID: 0
 SI Address              SI Status              SPU
 3001::2                 Active                 spu-1
 3001::2                 Active                 spu-21
 SI number: 2

user@host> show security softwires logical-system ls-product-design
Softwire Name            SC Address          Status    Number of SI connected
  sc_1                    3000::1             Connected    1
```

show security zones

Syntax	show security zones <detail | terse> < *zone-name* >
Release Information	Command introduced in Junos OS Release 8.5. The **Description** output field added in Junos OS Release 12.1.
Description	Display information about security zones.
Options	• none—Display information about all zones. • **detail | terse**—(Optional) Display the specified level of output. • *zone-name* —(Optional) Display information about the specified zone.
Required Privilege Level	view
Related Documentation	• *Security Zones and Interfaces Overview* • *Supported System Services for Host Inbound Traffic* • *security-zone*
List of Sample Output	show security zones on page 463 show security zones abc on page 463 show security zones abc detail on page 463 show security zones terse on page 464
Output Fields	Table 44 on page 462 lists the output fields for the **show security zones** command. Output fields are listed in the approximate order in which they appear.

Table 44: show security zones Output Fields

Field Name	Field Description
Security zone	Name of the security zone.
Description	Description of the security zone.
Policy configurable	Whether the policy can be configured or not.
Interfaces bound	Number of interfaces in the zone.
Interfaces	List of the interfaces in the zone.
Zone	Name of the zone.
Type	Type of the zone.

Sample Output

show security zones

```
user@host> show security zones
Functional zone: management
   Description: This is the management zone.
   Policy configurable: No
   Interfaces bound: 1
   Interfaces:
     ge-0/0/0.0
Security zone: Host
   Description: This is the host zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     fxp0.0
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
Security zone: def
   Description: This is the def zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/2.0
```

Sample Output

show security zones abc

```
user@host> show security zones abc
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones abc detail

```
user@host> show security zones abc detail
Security zone: abc
   Description: This is the abc zone.
   Send reset for non-SYN session TCP packets: Off
   Policy configurable: Yes
   Interfaces bound: 1
   Interfaces:
     ge-0/0/1.0
```

Sample Output

show security zones terse

```
user@host> show security zones terse
Zone                    Type
my-internal             Security
my-external             Security
dmz                     Security
```