



**Concepts & Examples
ScreenOS Reference Guide**

**Volume 2:
Fundamentals**

Release 6.0.0, Rev. 04

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

Copyright Notice

Copyright © 2008 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

All specifications are subject to change without notice. Juniper Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Juniper Networks' installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR JUNIPER NETWORKS REPRESENTATIVE FOR A COPY.

Table of Contents

	About This Volume	ix
	Document Conventions.....	x
	Web User Interface Conventions	x
	Command Line Interface Conventions	x
	Naming Conventions and Character Types	xi
	Illustration Conventions	xii
	Technical Documentation and Support	xiii
Chapter 1	ScreenOS Architecture	1
	Security Zones	2
	Security Zone Interfaces.....	3
	Physical Interfaces.....	3
	Subinterfaces.....	3
	Virtual Routers	4
	Policies.....	5
	Virtual Private Networks	6
	Virtual Systems	9
	Packet-Flow Sequence.....	10
	Jumbo Frames.....	13
	ScreenOS Architecture Example.....	14
	Example: (Part 1) Enterprise with Six Zones.....	14
	Example: (Part 2) Interfaces for Six Zones.....	16
	Example: (Part 3) Two Routing Domains.....	18
	Example: (Part 4) Policies.....	20
Chapter 2	Zones	25
	Viewing Preconfigured Zones.....	26
	Security Zones	28
	Global Zone	28
	SCREEN Options.....	28
	Binding a Tunnel Interface to a Tunnel Zone.....	29
	Configuring Security Zones and Tunnel Zones	30
	Creating a Zone	30
	Modifying a Zone.....	31
	Deleting a Zone	32
	Function Zones	33
Chapter 3	Interfaces	35
	Interface Types	36
	Logical Interfaces.....	36
	Physical Interfaces	36
	Wireless Interfaces.....	36
	Bridge Group Interfaces.....	37
	Subinterfaces	37
	Aggregate Interfaces	37
	Redundant Interfaces.....	37
	Virtual Security Interfaces	38

Function Zone Interfaces	38
Management Interfaces.....	38
High Availability Interfaces.....	38
Tunnel Interfaces.....	39
Deleting Tunnel Interfaces	42
Viewing Interfaces	43
Configuring Security Zone Interfaces	44
Binding an Interface to a Security Zone	44
Unbinding an Interface from a Security Zone	45
Addressing an L3 Security Zone Interface.....	46
Public IP Addresses	47
Private IP Addresses.....	47
Addressing an Interface	48
Modifying Interface Settings	48
Creating a Subinterface in the Root System.....	49
Deleting a Subinterface.....	50
Creating a Secondary IP Address	50
Backup System Interfaces	51
Configuring a Backup Interface.....	52
Configuring an IP Tracking Backup Interface.....	52
Configuring a Tunnel-if Backup Interface	53
Configuring a Route Monitoring Backup Interface	57
Loopback Interfaces.....	58
Creating a Loopback Interface	59
Setting the Loopback Interface for Management.....	59
Setting BGP on a Loopback Interface	59
Setting VSIs on a Loopback Interface.....	60
Setting the Loopback Interface as a Source Interface.....	60
Interface State Changes.....	61
Physical Connection Monitoring	63
Tracking IP Addresses	63
Interface Monitoring.....	68
Monitoring Two Interfaces	69
Monitoring an Interface Loop.....	70
Security Zone Monitoring	73
Down Interfaces and Traffic Flow.....	74
Failure on the Egress Interface.....	75
Failure on the Ingress Interface.....	76
Chapter 4	Interface Modes 79
Transparent Mode.....	80
Zone Settings.....	81
VLAN Zone	81
Predefined Layer 2 Zones	81
Traffic Forwarding.....	81
Unknown Unicast Options.....	82
Flood Method.....	83
ARP/Trace-Route Method	84
Configuring VLAN1 Interface for Management.....	87
Configuring Transparent Mode.....	89
NAT Mode.....	92
Inbound and Outbound NAT Traffic	94
Interface Settings.....	95
Configuring NAT Mode	95

	Route Mode.....	98
	Interface Settings.....	99
	Configuring Route Mode.....	99
Chapter 5	Building Blocks for Policies	103
	Addresses	103
	Address Entries	104
	Adding an Address.....	104
	Modifying an Address	105
	Deleting an Address.....	105
	Address Groups	105
	Creating an Address Group	107
	Editing an Address Group Entry	108
	Removing a Member and a Group.....	108
	Services.....	109
	Predefined Services	109
	Internet Control Messaging Protocol	110
	Handling ICMP Unreachable Errors	113
	Internet-Related Predefined Services.....	114
	Microsoft Remote Procedure Call Services	115
	Dynamic Routing Protocols.....	117
	Streaming Video.....	117
	Sun Remote Procedure Call Services	118
	Security and Tunnel Services	118
	IP-Related Services.....	119
	Instant Messaging Services.....	119
	Management Services	119
	Mail Services	120
	UNIX Services	120
	Miscellaneous Services.....	121
	Custom Services	121
	Adding a Custom Service	122
	Modifying a Custom Service.....	123
	Removing a Custom Service.....	123
	Setting a Service Timeout	123
	Service Timeout Configuration and Lookup.....	123
	Contingencies	124
	Example.....	125
	Defining a Custom Internet Control Message Protocol Service.....	126
	Remote Shell ALG.....	127
	Sun Remote Procedure Call Application Layer Gateway.....	127
	Typical RPC Call Scenario.....	127
	Customizing Sun RPC Services.....	128
	Customizing Microsoft Remote Procedure Call ALG.....	128
	Real-Time Streaming Protocol Application Layer Gateway.....	129
	RTSP Request Methods	131
	RTSP Status Codes	133
	Configuring a Media Server in a Private Domain.....	134
	Configuring a Media Server in a Public Domain	136
	Service Groups.....	138
	Modifying a Service Group	139
	Removing a Service Group	140
	Dynamic IP Pools.....	140
	Port Address Translation	141

	Creating a DIP Pool with PAT	142
	Modifying a DIP Pool	143
	Sticky DIP Addresses	143
	Using DIP in a Different Subnet	144
	Using a DIP on a Loopback Interface	149
	Creating a DIP Group	153
	Setting a Recurring Schedule	156
Chapter 6	Policies	159
	Basic Elements	160
	Three Types of Policies	161
	Interzone Policies	161
	Intrazone Policies	161
	Global Policies	162
	Policy Set Lists	163
	Policies Defined	164
	Policies and Rules	164
	Anatomy of a Policy	165
	ID	166
	Zones	166
	Addresses	166
	Services	166
	Action	167
	Application	167
	Name	168
	VPN Tunneling	168
	L2TP Tunneling	168
	Deep Inspection	169
	Placement at the Top of the Policy List	169
	Source Address Translation	169
	Destination Address Translation	169
	User Authentication	170
	HA Session Backup	171
	Web Filtering	172
	Logging	172
	Counting	172
	Traffic Alarm Threshold	172
	Schedules	172
	Antivirus Scanning	173
	Traffic Shaping	173
	Policies Applied	174
	Viewing Policies	174
	Creating Policies	174
	Creating Interzone Policies Mail Service	175
	Creating an Interzone Policy Set	178
	Creating Intrazone Policies	182
	Creating a Global Policy	184
	Entering a Policy Context	185
	Multiple Items per Policy Component	185
	Setting Address Negation	186
	Modifying and Disabling Policies	189
	Policy Verification	189
	Reordering Policies	190
	Removing a Policy	191

Chapter 7	Traffic Shaping	193
	Managing Bandwidth at the Policy Level	194
	Setting Traffic Shaping	195
	Setting Service Priorities	198
	Setting Priority Queuing	199
	Ingress Policing	203
	Shaping Traffic on Virtual Interfaces	203
	Interface-Level Traffic Shaping	204
	Policy-Level Traffic Shaping	205
	Packet Flow	206
	Example: Route-Based VPN with Ingress Policing	206
	Example: Policy-Based VPN with Ingress Policing	210
	Traffic Shaping Using a Loopback Interface	214
	DSCP Marking and Shaping	214
Chapter 8	System Parameters	217
	Domain Name System Support	217
	DNS Lookup	218
	DNS Status Table	219
	Setting the DNS Server and Refresh Schedule	219
	Setting a DNS Refresh Interval	220
	Dynamic Domain Name System	220
	Setting Up DDNS for a DynDNS Server	221
	Setting Up DDNS for a DDO Server	222
	Proxy DNS Address Splitting	222
	Dynamic Host Configuration Protocol	225
	Configuring a DHCP Server	226
	Customizing DHCP Server Options	230
	Placing the DHCP Server in an NSRP Cluster	231
	DHCP Server Detection	231
	Enabling DHCP Server Detection	232
	Disabling DHCP Server Detection	232
	Assigning a Security Device as a DHCP Relay Agent	233
	Forwarding All DHCP Packets	237
	Configuring Next-Server-IP	237
	Using a Security Device as a DHCP Client	238
	Propagating TCP/IP Settings	240
	Configuring DHCP in Virtual Systems	242
	Setting DHCP Message Relay in Virtual Systems	242
	Point-to-Point Protocol over Ethernet	243
	Setting Up PPPoE	243
	Configuring PPPoE on Primary and Backup Untrust Interfaces	246
	Configuring Multiple PPPoE Sessions over a Single Interface	247
	PPPoE and High Availability	250
	License Keys	250
	Registration and Activation of Subscription Services	251
	Trial Service	252
	Updating Subscription Keys	252
	Adding Antivirus, Web Filtering, Anti-Spam, and Deep Inspection to an Existing or a New Device	253
	System Clock	253
	Date and Time	254
	Daylight Saving Time	254

Time Zone	254
Network Time Protocol.....	255
Configuring Multiple NTP Servers.....	255
Configuring a Backup NTP Server.....	255
Device as an NTP Server	256
Maximum Time Adjustment.....	256
NTP and NSRP	257
Setting a Maximum Time Adjustment Value to an NTP Server	257
Securing NTP Servers.....	257
Index.....	IX-I

About This Volume

Volume 2: Fundamentals describes the ScreenOS architecture and its elements, including examples for configuring various elements. This volume contains the following chapters:

- Chapter 1, “ScreenOS Architecture,” presents the fundamental elements of the architecture in ScreenOS and concludes with a four-part example illustrating an enterprise-based configuration incorporating most of those elements. In this and all subsequent chapters, each concept is accompanied by illustrative examples.
- Chapter 2, “Zones,” explains security, tunnel, and function zones.
- Chapter 3, “Interfaces,” describes the various physical, logical, and virtual interfaces.
- Chapter 4, “Interface Modes,” explains the concepts behind Transparent, Network Address Translation (NAT), and Route interface operational modes.
- Chapter 5, “Building Blocks for Policies,” discusses the elements used for creating policies and virtual private networks (VPNs): addresses (including VIP addresses), services, and DIP pools.
- Chapter 6, “Policies,” explores the components and functions of policies and offers guidance on their creation and application.
- Chapter 7, “Traffic Shaping,” explains how you can prioritize services and manage bandwidth at the interface and policy levels.
- Chapter 8, “System Parameters,” presents the concepts behind Domain Name System (DNS) addressing, using Dynamic Host Configuration Protocol (DHCP) to assign or relay TCP/IP settings, downloading and uploading system configurations and software, and setting the system clock.

Document Conventions

This document uses the conventions described in the following sections:

- “Web User Interface Conventions” on page x
- “Command Line Interface Conventions” on page x
- “Naming Conventions and Character Types” on page xi
- “Illustration Conventions” on page xii

Web User Interface Conventions

In the Web user interface (WebUI), the set of instructions for each task is divided into navigational path and configuration settings. To open a WebUI page where you can enter configuration settings, you navigate to it by clicking on a menu item in the navigation tree on the left side of the screen, then on subsequent items. As you proceed, your navigation path appears at the top of the screen, each page separated by angle brackets.

The following shows the WebUI path and parameters for defining an address:

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

```
Address Name: addr_1
IP Address/Domain Name:
  IP/Netmask: (select), 10.2.2.5/32
Zone: Untrust
```

To open Online Help for configuration settings, click the question mark (?) in the upper left of the screen.

The navigation tree also provides a Help > Config Guide configuration page to help you configure security policies and Internet Protocol Security (IPSec). Select an option from the dropdown menu and follow the instructions on the page. Click the ? character in the upper left for Online Help on the Config Guide.

Command Line Interface Conventions

The following conventions are used to present the syntax of command line interface (CLI) commands in examples and in text.

In examples:

- Anything inside square brackets [] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example:

```
set interface { ethernet1 | ethernet2 | ethernet3 } manage
```

- Variables are in *italic* type:

```
set admin user name1 password xyz
```

In text, commands are in **boldface** type and variables are in *italic* type.

NOTE: When entering a keyword, you only have to type enough letters to identify the word uniquely. Typing **set adm u whee j12fmt54** will enter the command **set admin user wheezer j12fmt54**. However, all the commands documented here are presented in their entirety.

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations:

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes; for example:


```
set address trust "local LAN" 10.1.1.0/24
```
- Any leading spaces or trailing text within a set of double quotes are trimmed; for example, " local LAN " becomes "local LAN".
- Multiple consecutive spaces are treated as a single space.
- Name strings are case-sensitive, although many CLI keywords are case-insensitive. For example, "local LAN" is different from "local lan".

ScreenOS supports the following character types:

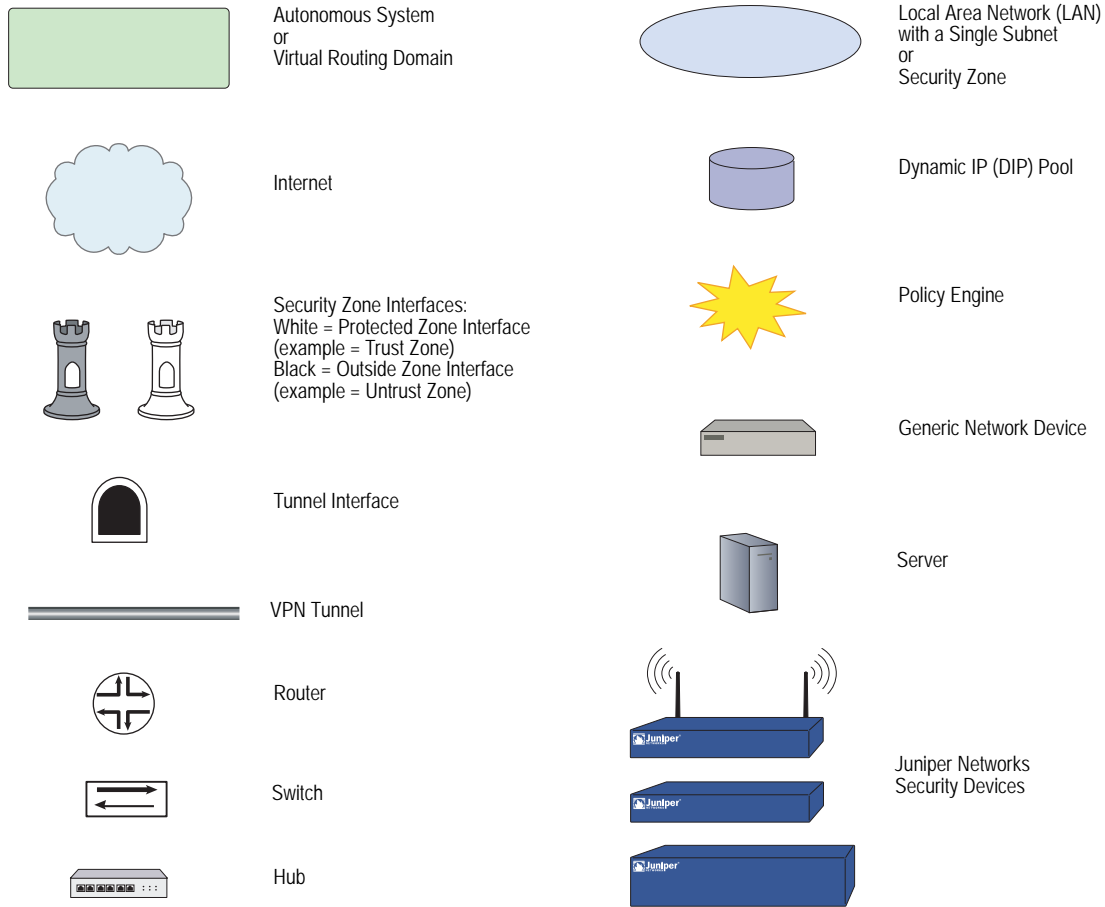
- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.
- ASCII characters from 32 (0x20 in hexadecimal) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NOTE: A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your browser supports.

Illustration Conventions

Figure 1 shows the basic set of images used in illustrations throughout this volume.

Figure 1: Images in Illustrations



Technical Documentation and Support

To obtain technical documentation for any Juniper Networks product, visit www.juniper.net/techpubs/.

For technical support, open a support case using the Case Management link at <http://www.juniper.net/customers/support/> or call 1-888-314-JTAC (from the United States, Canada, or Mexico) or 1-408-745-9500 (from elsewhere).

If you find any errors or omissions in this document, please contact Juniper Networks at techpubs-comments@juniper.net.

Chapter 1

ScreenOS Architecture

Juniper Networks ScreenOS architecture offers you flexibility in designing the layout of your network security. On Juniper Networks security devices with more than two interfaces, you can create numerous security zones and configure policies to regulate traffic between and within zones. You can bind one or more interfaces to each zone and enable different management and firewall options for each zone. ScreenOS allows you to create the number of zones required by your network environment, assign the number of interfaces required by each zone, and design each interface according to your needs.

This chapter presents an overview of ScreenOS. It contains the following sections:

- “Security Zones” on page 2
- “Security Zone Interfaces” on page 3
- “Virtual Routers” on page 4
- “Policies” on page 5
- “Virtual Private Networks” on page 6
- “Virtual Systems” on page 9
- “Packet-Flow Sequence” on page 10
- “Jumbo Frames” on page 13

The chapter concludes with a four-part example that illustrates a basic configuration for a security device using ScreenOS:

- “Example: (Part 1) Enterprise with Six Zones” on page 14
- “Example: (Part 2) Interfaces for Six Zones” on page 16
- “Example: (Part 3) Two Routing Domains” on page 18
- “Example: (Part 4) Policies” on page 20

Security Zones

A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic via policies (see “Policies” on page 5). Security zones are logical entities to which one or more interfaces are bound. With many types of Juniper Networks security devices, you can define multiple security zones, the exact number of which you determine based on your network needs. In addition to user-defined zones, you can also use the predefined zones: Trust, Untrust, and DMZ (for Layer 3 operation), or V1-Trust, V1-Untrust, and V1-DMZ (for Layer 2 operation). If you want, you can continue using just the predefined zones. You can also ignore the predefined zones and use user-defined zones exclusively. Optionally, you can use both kinds of zones—predefined and user-defined—side by side. This flexibility for zone configuration allows you to create a network design that best suits your specific needs. See Figure 2.


NOTE: The one security zone that requires no network segment is the global zone. (For more information, see “Global Zone” on page 28.) Additionally, any zone without an interface bound to it nor any address book entries can also be said not to contain any network segments.

If you upgrade from an earlier version of ScreenOS, all your configurations for these zones remain intact.

You cannot delete a predefined security zone. You can, however, delete a user-defined zone. When you delete a security zone, you also automatically delete all addresses configured for that zone.

Figure 2 shows a network configured with five security zones—three default zones (Trust, Untrust, DMZ) and two user-defined zones (Finance, Eng). Traffic passes from one security zone to another only if a policy permits it.

Figure 2: Predefined Security Zones



Security Zone Interfaces

An interface for a security zone can be thought of as a doorway through which TCP/IP traffic can pass between that zone and any other zone.

Through the policies you define, you can permit traffic between zones to flow in one direction or in both. With the routes that you define, you specify the interfaces that traffic from one zone to another must use. Because you can bind multiple interfaces to a zone, the routes you chart are important for directing traffic to the interfaces of your choice.

NOTE: For traffic to flow between interfaces bound to the same zone, no policy is required because both interfaces have security equivalency. ScreenOS requires policies for traffic between zones, not within a zone.

To permit traffic to flow from zone to zone, you bind an interface to the zone and—for an interface in Route or NAT mode (see “Interface Modes” on page 79)—assign an IP address to the interface. Two common interface types are physical interfaces and—for those devices with virtual-system support—subinterfaces (that is, a Layer 2 substantiation of a physical interface). For more information, see “Interfaces” on page 35.

Physical Interfaces

A physical interface relates to components that are physically present on the security device. The interface-naming convention differs from device to device.

NOTE: To see the naming conventions for a specific security device, refer to the installation and configuration guide for the device.

Subinterfaces

On devices that support virtual LANs (VLANs), you can logically divide a physical interface into several virtual subinterfaces, each of which borrows the bandwidth it needs from the physical interface from which it stems. A subinterface is an abstraction that functions identically to a physical interface and is distinguished by 802.1Q VLAN tagging. The security device directs traffic to and from a zone with a subinterface via its IP address and VLAN tag. For convenience, administrators usually use the same number for a VLAN tag as the subinterface number. For example, the interface ethernet1/2 using VLAN tag 3 is named *ethernet1/2.3*. This refers to the interface module in the first bay, the second port on that module, and subinterface number **3** (*ethernet1/2.3*).

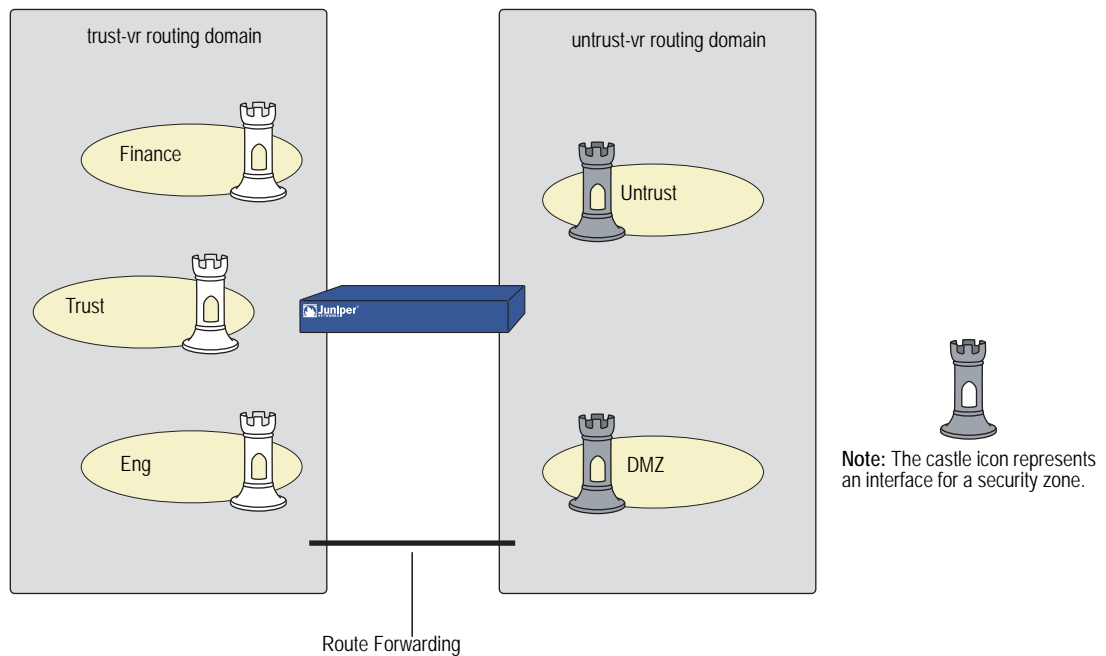
NOTE: 802.1Q is an IEEE standard that defines the mechanisms for the implementation of virtual bridged LANs and the ethernet frame formats used to indicate VLAN membership via VLAN tagging.

Note that although a subinterface shares part of its identity with a physical interface, the zone to which you bind it is not dependent on the zone to which you bind the physical interface. You can bind the subinterface *ethernet1/2.3* to a different zone than that to which you bind the physical interface *ethernet1/2*, or to which you bind *ethernet1/2.2*. Similarly, there are no restrictions in terms of IP-address assignments. The term *subinterface* does not imply that its address be in a subnet of the address space of the physical interface.

Virtual Routers

A virtual router (VR) functions as a router. It has its own interfaces and its own unicast and multicast routing tables. In ScreenOS, a security device supports two predefined virtual routers. This allows the security device to maintain two separate unicast and multicast routing tables and to conceal the routing information in one virtual router from the other. For example, the *untrust-vr* is typically used for communication with untrusted parties and does not contain any routing information for the protected zones. Routing information for the protected zones is maintained by the *trust-vr*. Thus, no internal network information can be gathered by the covert extraction of routes from the *untrust-vr*; see Figure 3.

Figure 3: Virtual Router Security Zones



When there are two virtual routers on a security device, traffic is *not* automatically forwarded between zones that reside in different VRs, even if there are policies that permit the traffic. If you want traffic to pass between virtual routers, you need to either export routes between the VRs or configure a static route in one VR that defines the other VR as the next hop. For more information about using two virtual routers, see *Volume 7: Routing*.

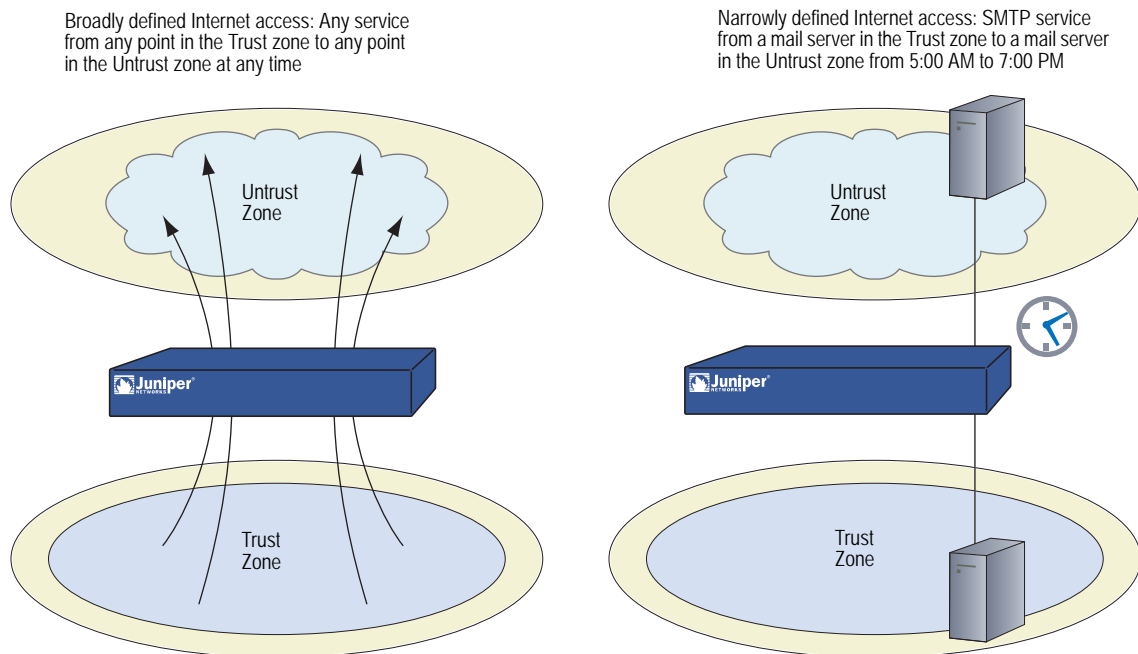
Policies

Juniper Networks security devices secure a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another.

By default, a security device denies all traffic in all directions. Through the creation of policies, you can control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from specified sources to specified destinations at scheduled times. At the broadest level, you can allow all kinds of traffic from any source in one zone to any destination in all other zones without any scheduling restrictions. At the narrowest level, you can create a policy that allows only one kind of traffic between a specified host in one zone and another specified host in another zone during a scheduled period, see Figure 4.

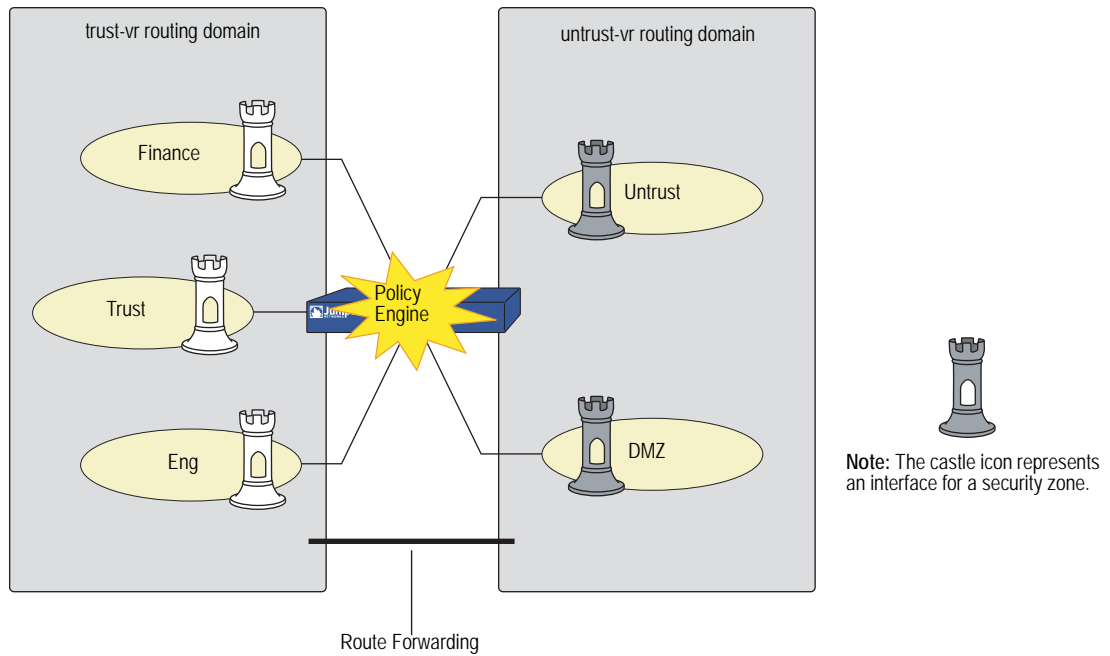
NOTE: Some security devices ship with a default policy that allows all outbound traffic from the Trust to the Untrust zone but denies all inbound traffic from the Untrust zone to the Trust zone.

Figure 4: Default Policy



Every time a packet attempts to pass from one zone to another or between two interfaces bound to the same zone, the security device checks its policy set lists for a policy that permits such traffic (see “Policy Set Lists” on page 163). To allow traffic to pass from one security zone to another—for example, from zone A to zone B—you must configure a policy that permits zone A to send traffic to zone B. To allow traffic to flow the other way, you must configure another policy permitting traffic from zone B to zone A. For any traffic to pass from one zone to another, there must be a policy that permits it. Also, if intrazone blocking is enabled, there must be a policy to permit traffic to pass from one interface to another within that zone. See Figure 5 on page 6.

Figure 5: Policy Architecture



NOTE: For more information, see “Policies” on page 159.

If you configure multicast routing on a security device, you might have to configure multicast policies. By default, a security device does not permit multicast control traffic between zones. Multicast control traffic refers to the messages transmitted by multicast protocols, such as Protocol Independent Multicast (PIM). Multicast policies control the flow of multicast control traffic only. To allow data traffic (both unicast and multicast) to pass between zones, you must configure firewall policies. (For more information, see “Multicast Policies” on page 7-153.)

Virtual Private Networks

ScreenOS supports several virtual private network (VPN) configuration options. The two main types are as follows:

- **Route-based VPN**—A route lookup determines which traffic the security device encapsulates. Policies either permit or deny traffic to the destination specified in the route. If the policy permits the traffic and the route references a tunnel interface bound to a VPN tunnel, then the security device also encapsulates it. This configuration separates the application of policies from the application of VPN tunnels. Once configured, such tunnels exist as available resources for securing traffic en route between one security zone and another.
- **Policy-based VPN**—A policy lookup determines which traffic the security device encapsulates when the policy references a particular VPN tunnel and specifies “tunnel” as the action.

A route-based VPN is a good choice for site-to-site VPN configurations because you can apply multiple policies to traffic passing through a single VPN tunnel. A policy-based VPN is a good choice for dialup VPN configurations because the dialup client might not have an internal IP address to which you can set a route. See Figure 6.


The following steps provide a sense of the main elements involved in a route-based VPN configuration:

1. While configuring the VPN tunnel (for example, *vpn-to-SF*, where *SF* is the destination or end entity), specify a physical interface or subinterface on the local device as the outgoing interface. (The IP address for this interface is what the remote peer must use when configuring its remote gateway.)
2. Create a tunnel interface (for example, *tunnel.1*), and bind it to a security zone.

NOTE: You do not have to bind the tunnel interface to the same zone for which VPN traffic is destined. Traffic to any zone can access a tunnel interface if a route points to that interface.

3. Bind the tunnel interface *tunnel.1* to the VPN tunnel *vpn-to-SF*.
4. To direct traffic through this tunnel, set up a route stating that traffic to *SF* must use *tunnel.1*.

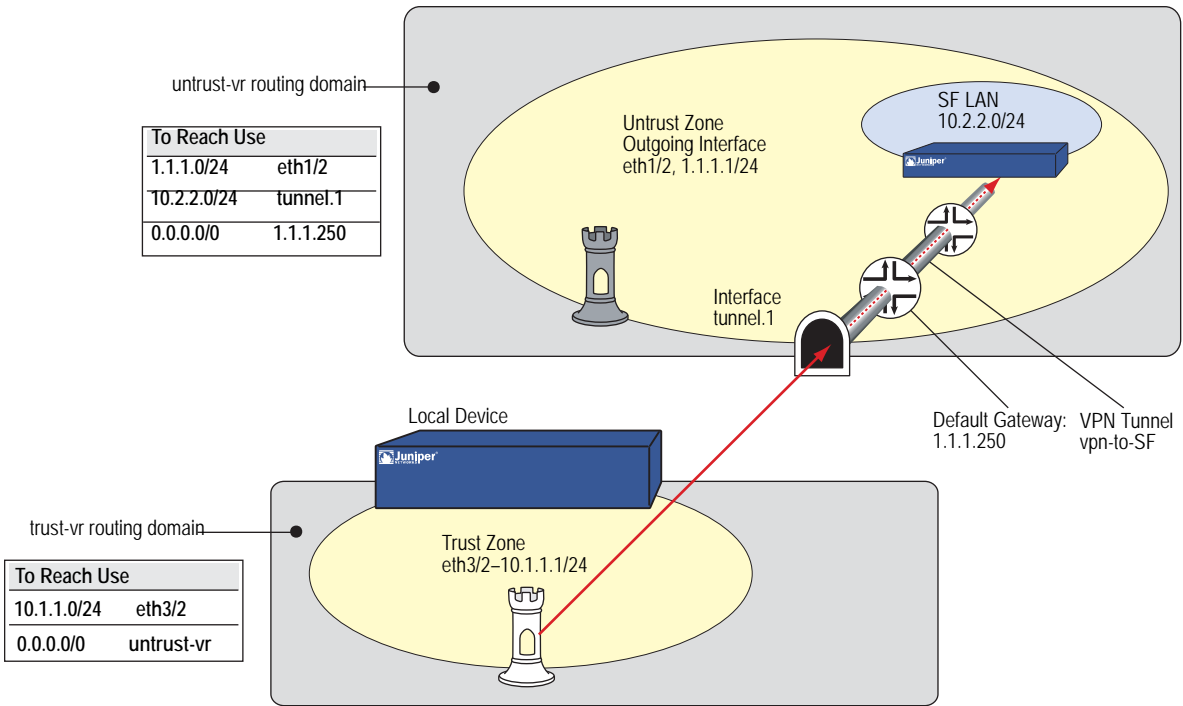
Figure 6: VPN Traffic



At this point, the tunnel is ready for traffic bound for *SF*. You can now create address-book entries, such as “Trust LAN” (10.1.1.0/24) and “SF LAN” (10.2.2.0/24) and set up policies to permit or block different types of traffic from a specified source, such as Trust LAN, to a specified destination, such as SF LAN. See Figure 7 on page 8.

Figure 7: VPN Traffic from Untrust Security Zone

The local security device routes traffic from the Trust zone to SF LAN in the Untrust zone through the tunnel.1 interface. Because tunnel.1 is bound to the VPN tunnel vpn-to-SF, the device encrypts the traffic and sends it through that tunnel to the remote peer.

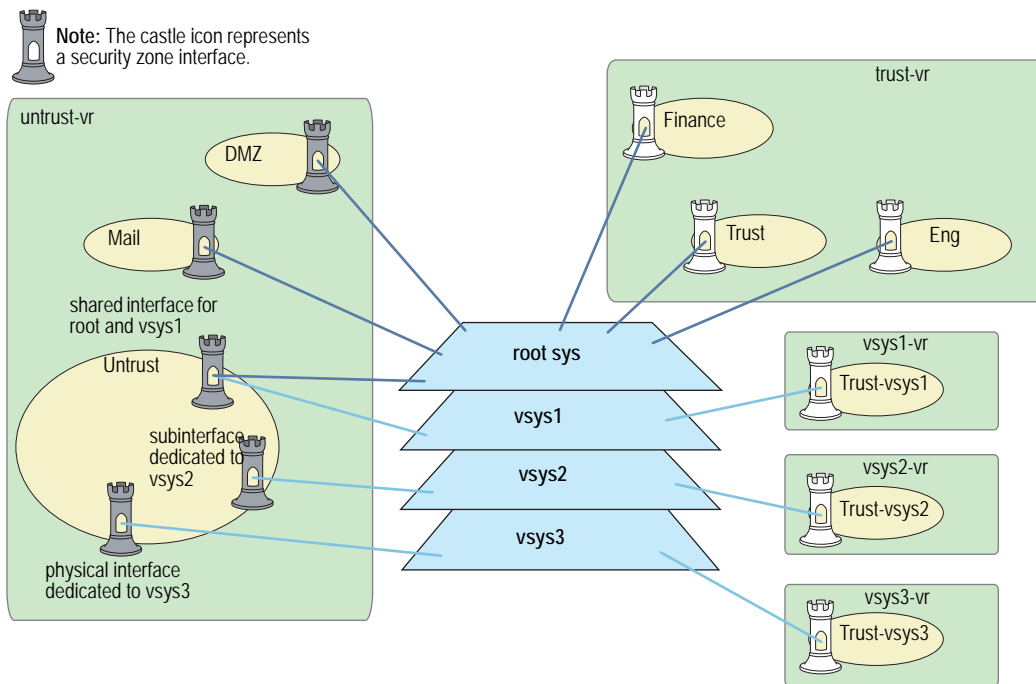


NOTE: For detailed information about VPNs, see *Volume 5: Virtual Private Networks*.

Virtual Systems

Some Juniper Networks security devices support virtual systems (vsys). A virtual system is a subdivision of the main system that appears to the user to be a standalone entity. Virtual systems reside separately from each other and from the root system within the same security device. The application of ScreenOS to virtual systems involves the coordination of three main components: zones, interfaces, and virtual routers. Figure 8 presents a conceptual overview of how ScreenOS integrates these components at both the root and vsys levels.

Figure 8: Vsys Architecture

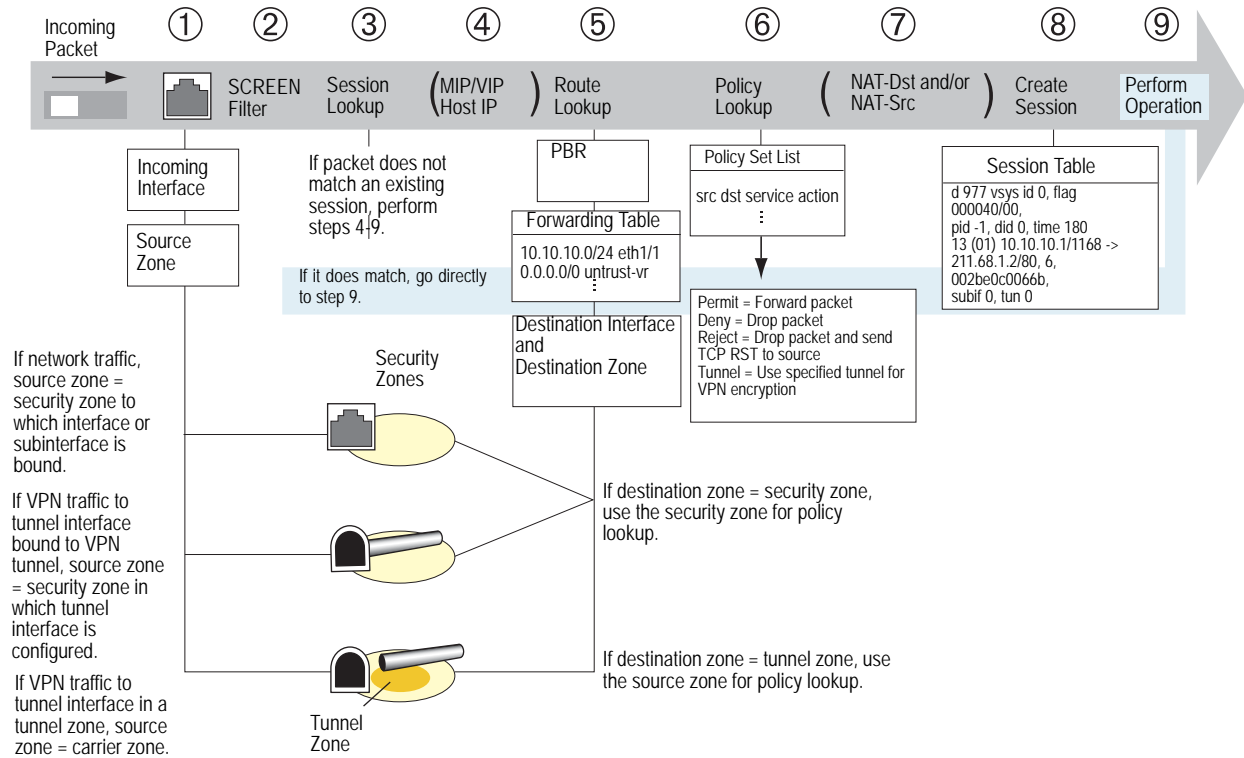


NOTE: For further information about virtual systems and the application of zones, interfaces, and virtual routers within the context of virtual systems, see *Volume 10: Virtual Systems*.

Packet-Flow Sequence

In ScreenOS, the flow sequence of an incoming packet progresses as presented in Figure 9.

Figure 9: Packet Flow Sequence Through Security Zones



1. The interface module identifies the incoming interface and, consequently, the source zone to which the interface is bound.

The interface module uses the following criteria to determine the source zone:

- If the packet is not encapsulated, the source zone is the security zone to which the incoming interface or subinterface is bound.
 - If the packet is encapsulated and the tunnel interface is bound to a VPN tunnel, the source zone is the security zone in which the tunnel interface is configured.
 - If the packet is encapsulated and the tunnel interface is in a tunnel zone, the source zone is the corresponding carrier zone (a security zone that carries a tunnel zone) for that tunnel zone.
2. If you have enabled SCREEN options for the source zone, the security device activates the SCREEN module at this point. SCREEN checking can produce one of the following three results:
 - If a SCREEN mechanism detects anomalous behavior for which it is configured to block the packet, the security device drops the packet and makes an entry in the event log.

- If a SCREEN mechanism detects anomalous behavior for which it is configured to record the event but not block the packet, the security device records the event in the SCREEN counters list for the ingress interface and proceeds to the next step.
 - If the SCREEN mechanisms detect no anomalous behavior, the security device proceeds to the next step.
3. The session module performs a session lookup, attempting to match the packet with an existing session.
 - If the packet does not match an existing session, the security device performs First Packet Processing, a procedure involving steps 4 through 9.
 - If the packet matches an existing session, the security device performs Fast Processing, using the information available from the existing session entry to process the packet. Fast Processing bypasses steps 4 through 8 because the information generated by those steps has already been obtained during the processing of the first packet in the session.
 4. If a mapped IP (MIP) or virtual IP (VIP) address is used, the address-mapping module resolves the MIP or VIP so that the routing table can search for the actual host address.
 5. Prior to route lookup, ScreenOS checks the packet for policy-based routing (PBR). If PBR is enabled on that in-interface, the following actions apply to the packet:
 - The PBR policy bound to that in-interface is applied to the packet.
 - If no PBR policy exists at the interface level, the PBR policy bound to the zone associated with the in-interface is applied to the packet.
 - If no PBR policy exists at the zone level, the PBR policy bound to the VR associated with the in-interface is applied to the packet.

NOTE: For more information about PBR, see *Volume 7: Routing*.

If PBR is not enabled, the route table lookup finds the interface that leads to the destination address. In so doing, the interface module identifies the destination zone to which that interface is bound.

The interface module uses the following criteria to determine the destination zone:

- If the destination zone is a security zone, that zone is used for the policy lookup.
 - If the destination zone is a tunnel zone, the corresponding carrier zone is used for the policy lookup.
 - If the destination zone is the same as the source zone and intrazone blocking is disabled for that zone, the security device bypasses steps 6 and 7 and creates a session (step 8). If intrazone blocking is enabled, then the security device drops the packet.
6. The policy engine searches the policy set lists for a policy between the addresses in the identified source and destination zones.

The action configured in the policy determines how the security device handles the packet:

- If the action is **permit**, the security device will forward the packet to its destination.
 - If the action is **deny**, the security device will drop the packet.
 - If the action is **reject**, the security device will drop the packet and—if the protocol is TCP—send a reset (RST) to the source IP address.
 - If the action is **tunnel**, the security device will forward the packet to the VPN module, which encapsulates the packet and transmits it using the specified VPN tunnel settings.
7. If destination address translation (NAT-dst) is specified in the policy, the NAT module translates the original destination address in the IP packet header to a different address.

If source address translation is specified (either interface-based NAT or policy-based NAT-src), the NAT module translates the source address in the IP packet header before forwarding it either to its destination or to the VPN module.

(If both NAT-dst and NAT-src are specified in the same policy, the security device first performs NAT-dst and then NAT-src.)

8. The session module creates a new entry in the session table containing the results of steps 1 through 7.

The security device then uses the information maintained in the session entry when processing subsequent packets of the same session.

9. The security device performs the operation specified in the session.

Some typical operations are source address translation, VPN tunnel selection and encryption, decryption, and packet forwarding.

Jumbo Frames

On some devices you can increase throughput by increasing the maximum packet size, or message transmission unit (MTU), the device can process. Refer to the installation and configuration guide for your device to find out if it supports jumbo frames.

Frame size ranges from 1514 through 9830 bytes. To put the device in jumbo frame mode, set the maximum frame size to a value from 1515 through 9830 inclusive, for example: **set envar max-frame-size= 9830**. Use the **unset envar max-frame-size** command to return the device to normal maximum frame size, which is 1514 bytes (alternatively, you can use the command: **set envar max-frame-size= 1514**). The maximum frame size does not include the 4-byte frame check sequence at the end of the frame. You must restart the system for changes to environmental variables to take effect.

In jumbo frame mode, the following apply:

- Deep inspection (DI) is not supported.
- Packets sent through aggregate interfaces might be out of order.
- NSRP forwarding is not supported.
- Maximum firewall or VPN throughput requires at least four sessions (for firewall) or tunnels (for VPN)

ScreenOS Architecture Example

The following sections comprise a four-part example that illustrates some of the concepts covered in the previous sections:

- “Example: (Part 1) Enterprise with Six Zones”
- “Example: (Part 2) Interfaces for Six Zones” on page 16
- “Example: (Part 3) Two Routing Domains” on page 18
- “Example: (Part 4) Policies” on page 20

Example: (Part 1) Enterprise with Six Zones


This is the first of a four-part example, the purpose of which is to illustrate some of the concepts covered in the previous sections. For the second part, in which the interfaces for each zone are set, see “Example: (Part 2) Interfaces for Six Zones” on page 16. Here you configure the following six zones for an enterprise:

- Finance
- Trust
- Eng
- Mail
- Untrust
- DMZ

The Trust, Untrust, and DMZ zones are pre-configured. You must define the Finance, Eng, and Mail zones. By default, a user-defined zone is placed in the trust-vr routing domain. Thus, you do not have to specify a virtual router for the Finance and Eng zones. However, in addition to configuring the Mail zone, you must also specify that it be in the untrust-vr routing domain. You must also shift virtual router bindings for the Untrust and DMZ zones from the trust-vr to the untrust-vr, see Figure 10 on page 15.

NOTE: For more information about virtual routers and their routing domains, see *Volume 7: Routing*.

Figure 10: Zone-to-Virtual Router Bindings

**WebUI**

Network > Zones > New: Enter the following, then click **OK**:

Zone Name: Finance
 Virtual Router Name: trust-vr
 Zone Type: Layer 3: (select)

Network > Zones > New: Enter the following, then click **OK**:

Zone Name: Eng
 Virtual Router Name: trust-vr
 Zone Type: Layer 3: (select)

Network > Zones > New: Enter the following, then click **OK**:

Zone Name: Mail
 Virtual Router Name: untrust-vr
 Zone Type: Layer 3: (select)

Network > Zones > Edit (for Untrust): Select **untrust-vr** in the Virtual Router Name drop-down list, then click **OK**.

Network > Zones > Edit (for DMZ): Select **untrust-vr** in the Virtual Router Name dropdown list, then click **OK**.


CLI

```
set zone name finance
set zone name eng
set zone name mail
set zone mail vrouter untrust-vr
set zone untrust vrouter untrust-vr
set zone dmz vrouter untrust-vr
save
```

Example: (Part 2) Interfaces for Six Zones

This is the second part of an ongoing example. For the first part, in which zones are configured, see “Example: (Part 1) Enterprise with Six Zones” on page 14. For the next part, in which virtual routers are configured, see “Example: (Part 3) Two Routing Domains” on page 18. This part of the example demonstrates how to bind interfaces to zones and configure them with an IP address and various management options, see Figure 11.

Figure 11: Interface-to-Zone Bindings



WebUI

1. Interface ethernet3/2

Network > Interfaces > Edit (for ethernet3/2): Enter the following, then click **OK**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24
 Manageable: (select)
 Management Services: WebUI, Telnet, SNMP, SSH (select)
 Other Services: Ping (select)

2. Interface ethernet3/2.1

Network > Interfaces > Sub-IF New: Enter the following, then click **OK**:

Interface Name: ethernet3/2.1
 Zone Name: Finance
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.2.1/24
 VLAN Tag: 1
 Other Services: Ping (select)

3. Interface ethernet3/1

Network > Interfaces > Edit (for ethernet3/1): Enter the following, then click **OK**:

Zone Name: Eng
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.3.1/24
 Other Services: Ping (select)

4. Interface ethernet1/1

Network > Interfaces > Edit (for ethernet1/1): Enter the following, then click **OK**:

Zone Name: Mail
 Static IP: (select this option when present)
 IP Address/Netmask: 1.3.3.1/24

5. Interface ethernet1/1.2

Network > Interfaces > Sub-IF New: Enter the following, then click **OK**:

Interface Name: ethernet1/1.2
 Zone Name: Mail
 Static IP: (select this option when present)
 IP Address/Netmask: 1.4.4.1/24
 VLAN Tag: 2

6. Interface ethernet1/2

Network > Interfaces > Edit (for ethernet1/2): Enter the following, then click **OK**:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24
 Manageable: (select)
 Management Services: SNMP (select)

7. Interface ethernet2/2

Network > Interfaces > Edit (for ethernet2/2): Enter the following, then click **OK**:

Zone Name: DMZ
 Static IP: (select)
 IP Address/Netmask: 1.2.2.1/24

CLI**1. Interface ethernet3/2**

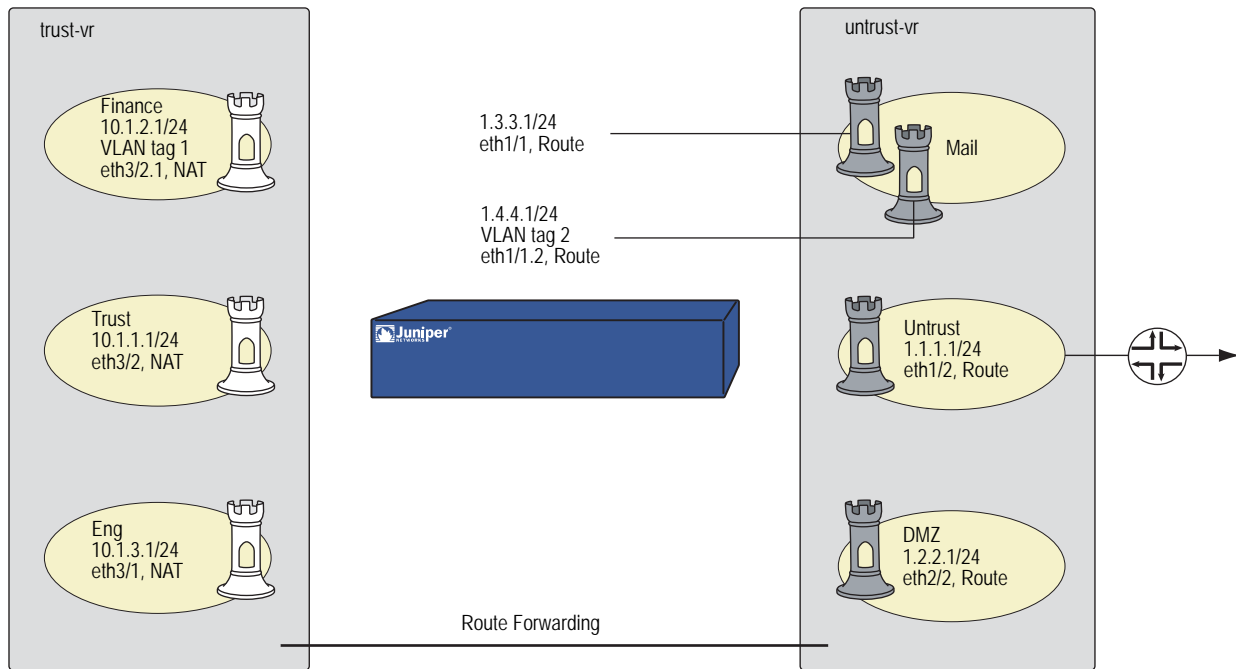
```
set interface ethernet3/2 zone trust
set interface ethernet3/2 ip 10.1.1.1/24
set interface ethernet3/2 manage ping
set interface ethernet3/2 manage webui
set interface ethernet3/2 manage telnet
set interface ethernet3/2 manage snmp
set interface ethernet3/2 manage ssh
```

2. **Interface ethernet3/2.1**
set interface ethernet3/2.1 tag 1 zone finance
set interface ethernet3/2.1 ip 10.1.2.1/24
set interface ethernet3/2.1 manage ping
3. **Interface ethernet3/1**
set interface ethernet3/1 zone eng
set interface ethernet3/1 ip 10.1.3.1/24
set interface ethernet3/1 manage ping
4. **Interface ethernet1/1**
set interface ethernet1/1 zone mail
set interface ethernet1/1 ip 1.3.3.1/24
5. **Interface ethernet1/1.2**
set interface ethernet1/1.2 tag 2 zone mail
set interface ethernet1/1.2 ip 1.4.4.1/24
6. **Interface ethernet1/2**
set interface ethernet1/2 zone untrust
set interface ethernet1/2 ip 1.1.1.1/24
set interface ethernet1/2 manage snmp
7. **Interface ethernet2/2**
set interface ethernet2/2 zone dmz
set interface ethernet2/2 ip 1.2.2.1/24
save

Example: (Part 3) Two Routing Domains

This is the third part of an ongoing example. For the previous part, in which interfaces for the various security zones are defined, see “Example: (Part 2) Interfaces for Six Zones” on page 16. For the next part, in which the policies are set, see “Example: (Part 4) Policies” on page 20. In this example, you only have to configure a route for the default gateway to the Internet. The other routes are automatically created by the security device when you create the interface IP addresses, see Figure 12 on page 19.

Figure 12: Routing Domains

**WebUI**

Network > Routing > Destination > (select trust-vr) New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Next Hop Virtual Router Name: (select); untrust-vr

Network > Routing > Destination > (select untrust-vr) New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet1/2
 Gateway IP Address: 1.1.1.254

CLI

```
set router trust-vr route 0.0.0.0/0 vrouter untrust-vr
set router untrust-vr route 0.0.0.0/0 interface eth1/2 gateway 1.1.1.254
save
```

The security device automatically creates the routes shown in Table 1 and Table 2 (except as indicated).

Table 1: Route Table for trust-vr

To Reach:	Use Interface:	Use Gateway/Vrouter:	Created by:
0.0.0.0/0	n/a	untrust-vr	User-configured
10.1.3.0/24	eth3/1	0.0.0.0	Security device
10.1.1.0/24	eth3/2	0.0.0.0	Security device
10.1.2.0/24	eth3/2.1	0.0.0.0	Security device

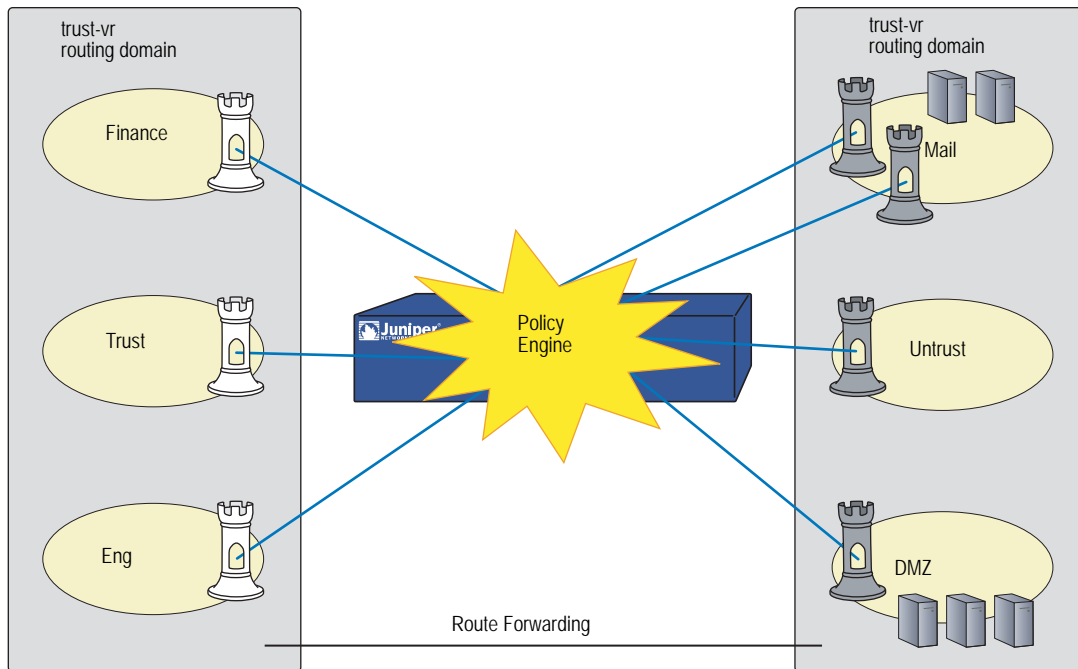
Table 2: Route Table for untrust-vr

To Reach:	Use Interface:	Use Gateway/Vrouter:	Created by:
1.2.2.0/24	eth2/2	0.0.0.0	Security device
1.1.1.0/24	eth1/2	0.0.0.0	Security device
1.4.4.0/24	eth1/1.2	0.0.0.0	Security device
1.3.3.0/24	eth1/1	0.0.0.0	Security device
0.0.0.0/0	eth1/2	1.1.1.254	User-configured

Example: (Part 4) Policies

This is the last part of an ongoing example. The previous part is “Example: (Part 3) Two Routing Domains” on page 18. This part of the example demonstrates how to configure new policies. See Figure 13.

Figure 13: Policies



For the purpose of this example, before you begin configuring new policies, you need to create new service groups.

NOTE: When you create a zone, the security device automatically creates the address **Any** for all hosts within that zone. This example makes use of the address **Any** for the hosts.

WebUI

1. Service Groups

Policy > Policy Elements > Services > Groups > New: Enter the following, then click **OK**:

Group Name: Mail-Pop3

Select **Mail**, then use the < < button to move that service from the Available Members column to the Group Members column.

Select **Pop3**, then use the < < button to move that service from the Available Members column to the Group Members column.

Policy > Policy Elements > Services > Groups > New: Enter the following, then click **OK**:

Group Name: HTTP-FTPGet

Select **HTTP**, then use the < < button to move that service from the Available Members column to the Group Members column.

Select **FTP-Get**, then use the < < button to move that service from the Available Members column to the Group Members column.

2. Policies

Policy > Policies > (From: Finance, To: Mail) New: Enter the following, then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Mail-Pop3

Action: Permit

Policy > Policies > (From: Trust, To: Mail) New: Enter the following, then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Mail-Pop3

Action: Permit

Policy > Policies > (From: Eng, To: Mail) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: Mail-Pop3
Action: Permit

Policy > Policies > (From: Untrust, To: Mail) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: Mail
Action: Permit

Policy > Policies > (From: Finance, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: HTTP-FTPGet
Action: Permit

Policy > Policies > (From: Finance, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: HTTP-FTPGet
Action: Permit

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: HTTP-FTPGet
Action: Permit

Policy > Policies > (From: Trust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: HTTP-FTPGet
 Action: Permit

Policy > Policies > (From: Eng, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: HTTP-FTPGet
 Action: Permit

Policy > Policies > (From: Eng, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: FTP-Put
 Action: Permit

Policy > Policies > (From: Untrust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: HTTP-FTPGet
 Action: Permit

CLI

1. Service Groups

```
set group service mail-pop3 add mail
set group service mail-pop3 add pop3
set group service http-ftpget add http
set group service http-ftpget add ftp-get
```

2. Policies

set policy from finance to mail any any mail-pop3 permit
set policy from trust to mail any any mail-pop3 permit
set policy from eng to mail any any mail-pop3 permit
set policy from untrust to mail any any mail permit
set policy from finance to untrust any any http-ftpget permit
set policy from finance to dmz any any http-ftpget permit
set policy from trust to untrust any any http-ftpget permit
set policy from trust to dmz any any http-ftpget permit
set policy from eng to untrust any any http-ftpget permit
set policy from eng to dmz any any http-ftpget permit
set policy from eng to dmz any any ftp-put permit
set policy from untrust to dmz any any http-ftpget permit
save

Chapter 2

Zones

A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or a logical entity that performs a specific function (a function zone).

This chapter examines each type of zone, with particular emphasis given to the security zone. It contains the following sections:

- “Viewing Preconfigured Zones” on page 26
- “Security Zones” on page 28
- “Binding a Tunnel Interface to a Tunnel Zone” on page 29
- “Configuring Security Zones and Tunnel Zones” on page 30
- “Function Zones” on page 33

Viewing Preconfigured Zones

When you first boot up a security device, you can see a number of preconfigured zones. To view these zones using the WebUI, click **Network > Zones** in the menu column on the left, see Figure 14.

To view these zones using the CLI, use the **get zone** command. See Figure 15 on page 27.

Figure 14: Network > Zones Page in the WebUI

ID	Name	Virtual Router	Vsys	Default IP	Type	Attribute	Configure
0	Null	untrust-vr	Root	hidden	Null	Shared	
2	Trust	trust-vr	Root	ethernet0/2	Security(L3)		Edit Screen Mal-URL
1	Untrust	trust-vr	Root	serial1/0	Security(L3)	Shared	Edit Screen Mal-URL
4	Self	trust-vr	Root	self	Function		
10	Global	trust-vr	Root	null	Security(L3)		
6	HA	trust-vr	Root	null	Function		
5	MGT	trust-vr	Root	null	Function		Edit Screen Mal-URL
16	Untrust-Tun	trust-vr	Root	hidden.1	Tunnel		
15	V1-Null	trust-vr	Root	i2v	Security(L2)	Shared	Edit Screen Mal-URL
12	V1-Trust	trust-vr	Root	v1-trust	Security(L2)	Shared	Edit Screen Mal-URL
11	V1-Untrust	trust-vr	Root	v1-untrust	Security(L2)	Shared	Edit Screen Mal-URL
3	DMZ	trust-vr	Root	ethernet0/1	Security(L3)		Edit Screen Mal-URL
13	V1-DMZ	trust-vr	Root	v1-dmz	Security(L2)	Shared	Edit Screen Mal-URL
14	VLAN	trust-vr	Root	vlan1	Function(vlan)	Shared	Edit

Figure 15 shows the output of the **get zone** CLI command.

Figure 15: Get Zone Output

device> get zone

Total of 13 zones in vsys root

ID	Name)	Type)	Attr)	VR)	Default-IF)	VSYS)
0)	Null)	Null)	Shared)	untrust-vr)	null)	Root)
1)	Untrust)	Sec(L3))	Shared)	trust-vr)	ethernet1/2)	Root)
2)	Trust)	Sec(L3)))	trust-vr)	ethernet3/2)	Root)
3)	DMZ)	Sec(L3)))	trust-vr)	ethernet2/2)	Root)
4)	Self)	Func))	trust-vr)	self)	Root)
5)	MGT)	Func))	trust-vr)	mgmt)	Root)
6)	HA)	Func))	trust-vr)	ha)	Root)
10)	Global)	Sec(L3)))	trust-vr)	null)	Root)
11)	V1-Untrust)	Sec(L2))	trust-vr)	v1-untrust)	Root)	Root)
12)	V1-Trust)	Sec(L2)))	trust-vr)	v1-trust)	Root)
13)	V1-DMZ)	Sec(L2)))	trust-vr)	v1-dmz)	Root)
14)	VLAN)	Func))	trust-vr)	vlan)	Root)
16)	Untrust-Tun)	Tun))	trust-vr)	null)	Root)	

Zone numbers 7-9 and 15 are reserved for future use.

The root and virtual systems share these zones.

These zones (ID 0 and 10) do not and cannot have an interface.

These zones (ID 1-3 and 11-14) provide backward compatibility when upgrading from a release prior to ScreenOS 3.1.0—the upper 3 for devices in NAT or Route mode, the lower 3 for devices in Transparent mode.

By default, VPN tunnel interfaces are bound to the Untrust-Tun zone, whose carrier zone is the Untrust zone. (When upgrading, existing tunnels are bound to the Untrust-Tun zone.)

The preconfigured zones shown in Figure 14 and Figure 15 can be grouped into four different zone types:

- Security: Untrust, Trust, DMZ, Global, V1-Untrust, V1-Trust, V1-DMZ, V1-Null
- Tunnel: Untrust-Tun
- Function: Null, Self, MGT, HA, VLAN
- Null: Null

Security Zones

On a single security device, you can configure multiple security zones, dividing the network into segments to which you can apply various security options to satisfy the needs of each segment. At a minimum, you must define two security zones, basically to protect one area of the network from the other. On some security platforms, you can define many security zones, bringing finer granularity to your network security design—and without deploying multiple security appliances to do so.

Global Zone

You can identify a security zone because it has an address book and can be referenced in policies. The Global zone satisfies these criteria. However, it does not have one element that all other security zones have—an interface. The Global zone serves as a storage area for mapped IP (MIP) and virtual IP (VIP) addresses. The predefined Global zone address “Any” applies to all MIPs, VIPs, and other user-defined addresses set in the Global zone. Because traffic going to these addresses is mapped to other addresses, the Global zone does not require an interface for traffic to flow through it.

The Global zone also contains addresses for use in global policies. For information about global policies, see “Global Policies” on page 162.

NOTE: Any policy that uses the Global zone as its destination cannot support NAT or traffic shaping.

SCREEN Options

A Juniper Networks firewall secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another. For every security zone, and the MGT zone, you can enable a set of predefined SCREEN options that detect and block various kinds of traffic that the security device determines as potentially harmful.

For more information about available SCREEN options, see *Volume 4: Attack Detection and Defense Mechanisms*.

Binding a Tunnel Interface to a Tunnel Zone

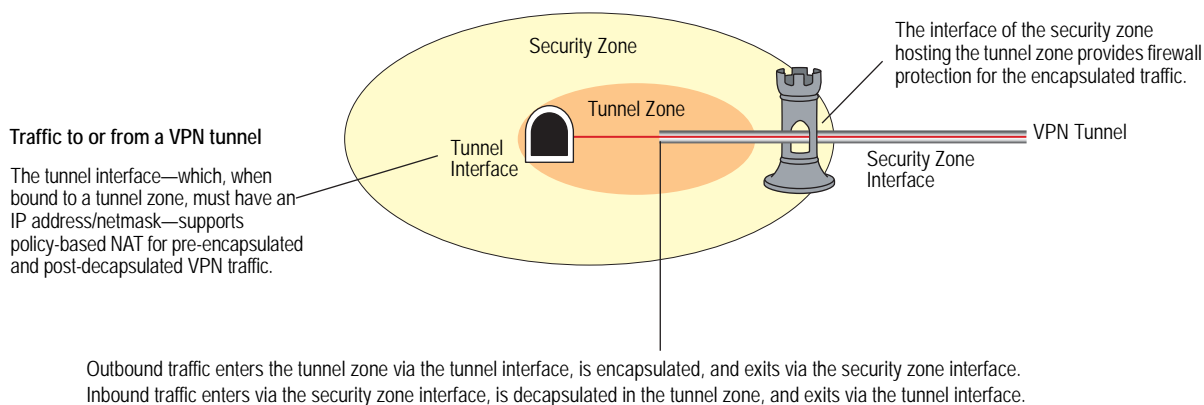
A tunnel zone is a logical segment that hosts one or more tunnel interfaces. A tunnel zone is conceptually affiliated with a security zone in a “child-parent” relationship. The security zone acting as the “parent,” which you can also conceive of as a carrier zone, provides the firewall protection to the encapsulated traffic. The tunnel zone provides packet encapsulation/decapsulation, and—by supporting tunnel interfaces with IP addresses and netmasks that can host mapped IP (MIP) addresses and dynamic IP (DIP) pools—can also provide policy-based NAT services.

The security device uses the routing information for the carrier zone to direct traffic to the tunnel endpoint. The default tunnel zone is Untrust-Tun, and it is associated with the Untrust zone. You can create other tunnel zones and bind them to other security zones, with a maximum of one tunnel zone per carrier zone per virtual system.

NOTE: The root system and all virtual systems can share the Untrust zone. However, each system has its own separate Untrust-Tun zone.

By default, a tunnel zone is in the trust-vr routing domain, but you can also move a tunnel zone into another routing domain. See Figure 16.

Figure 16: Tunnel Zone Routing Domain



When upgrading from a version of ScreenOS earlier than 3.1.0, existing tunnel interfaces are bound by default to the preconfigured Untrust-Tun tunnel zone, which is a “child” of the preconfigured Untrust security zone. You can bind multiple tunnel zones to the same security zone; however, you cannot bind a tunnel zone to another tunnel zone.

In this example, you create a tunnel interface and name it tunnel.3. You bind it to the Untrust-Tun zone, and assign it IP address 3.3.3.3/24. You then define a mapped IP (MIP) address on tunnel.3, translating 3.3.3.5 to 10.1.1.5, which is the address of a server in the Trust zone. Both the Untrust zone, which is the carrier zone for the Untrust-Tun zone, and the Trust zone are in the trust-vr routing domain.

WebUI

1. Tunnel Interface

Network > Interfaces > New Tunnel IF: Enter the following, then click **OK**:

Tunnel Interface Name: tunnel.3
 Zone (VR): Untrust-Tun (trust-vr)
 Fixed IP: (select)
 IP Address / Netmask: 3.3.3.3/24

2. MIP

Network > Interfaces > Edit (for tunnel.3) > MIP > New: Enter the following, then click **OK**:

Mapped IP: 3.3.3.5
 Host IP Address: 10.1.1.5
 Netmask: 255.255.255.255
 Host Virtual Router Name: trust-vr

CLI

1. Tunnel Interface

```
set interface tunnel.3 zone Untrust-Tun
set interface tunnel.3 ip 3.3.3.3/24
```

2. MIP

```
set interface tunnel.3 mip 3.3.3.5 host 10.1.1.5
save
```

Configuring Security Zones and Tunnel Zones

For best performance, always save your changes and reboot after creating a zone. The processes for creating, modifying, and deleting Layer 3 or Layer 2 security zones and tunnel zones are quite similar.

NOTE: You cannot delete predefined security zones or the predefined tunnel zone, although you can edit them.

Creating a Zone

To create a Layer 3 or Layer 2 security zone, or to create a tunnel zone, use either the WebUI or CLI.

WebUI

Network > Zones > New: Enter the following, then click **OK**:

Zone Name: Type a name for the zone.

Virtual Router Name: Select the virtual router in whose routing domain you want to place the zone.

Zone Type:

Select **Layer 3** to create a zone to which you can bind interfaces in NAT or Route mode.

Select **Layer 2** to create a zone to which you can bind interfaces in Transparent mode.

Select **Tunnel Out Zone** when creating a tunnel zone and binding it to a carrier zone, then select a specific carrier zone from the dropdown list.

Block Intra-Zone Traffic: Select this option to block traffic between hosts within the same security zone. By default, intra-zone blocking is disabled.

NOTE: The name of a Layer 2 security zone must begin with “L2-”; for example, “L2-Corp” or “L2-XNet.”

CLI

```
set zone name zone [ l2 vlan_id_num | tunnel sec_zone ]
set zone zone block
set zone zone vrouter name_str
save
```

NOTE: When creating a Layer 2 security zone, the VLAN ID number must be 1 (for VLAN1).

Modifying a Zone

To modify the name of a security zone or tunnel zone, or to change the carrier zone for a tunnel zone, you must first delete the zone and then create it again with the changes. You can change the intra-zone blocking option and the virtual router on an existing zone.

NOTE: Before you can remove a zone, you must first unbind all interfaces bound to it.

Before you can change the virtual router for a zone, you must first remove any interfaces bound to it.

WebUI**1. Modifying the Zone Name**

Network > Zones: Click **Remove** (for the security zone or tunnel zone whose name you want to change or for the tunnel zone whose carrier zone you want to change).

When the prompt appears, asking for confirmation of the removal, click **Yes**.

Network > Zones > New: Enter the zone settings with your changes, then click **OK**.

2. Changing the Intra-Zone Blocking Option or Virtual Router

Network > Zones > Edit (for the zone that you want to modify): Enter the following, then click **OK**:

Virtual Router Name: From the dropdown list, select the virtual router into whose routing domain you want to move the zone.

Block Intra-Zone Traffic: To enable, select the checkbox. To disable, clear it.

CLI

1. Modifying the Zone Name

```
unset zone zone
set zone name zone [ l2 vlan_id_num | tunnel sec_zone ]
```

2. Changing the Intra-Zone Blocking Option or Virtual Router

```
{ set | unset } zone zone block
set zone zone vrouter name_str
save
```

Deleting a Zone

For best performance, always save your changes and reboot after deleting a zone. To delete a security zone or tunnel zone, do either of the following:

WebUI

Network > Zones: Click **Remove** (for the zone you want to delete).

When the prompt appears, asking for confirmation of the removal, click **Yes**.

CLI

```
unset zone zone
save
```

NOTE: Before you can remove a zone, you must first unbind all interfaces bound to it. To unbind an interface from a zone, see “Binding an Interface to a Security Zone” on page 44.

Function Zones

The five function zones are Null, MGT, HA, Self, and VLAN. Each zone exists for a single purpose, as explained in Table 3.

Table 3: Function Zones

Zone	Description
Null Zone	This zone serves as temporary storage for any interfaces that are not bound to any other zone.
MGT Zone	This zone hosts the out-of-band management interface, MGT. You can set firewall options on this zone to protect the management interface from different types of attacks. For more information about firewall options, see <i>Volume 4: Attack Detection and Defense Mechanisms</i> .
HA Zone	This zone hosts the high availability interfaces, HA1 and HA2. Although you can set interfaces for the HA zone, the zone itself cannot be configured.
Self Zone	This zone hosts the interface for remote management connections. When you connect to the security device via HTTP, SCS, or Telnet, you connect to the Self zone.
VLAN Zone	This zone hosts the VLAN1 interface, which you use to manage the device and terminate VPN traffic when the device is in Transparent mode. You can also set firewall options on this zone to protect the VLAN1 interface from various attacks.

Chapter 3

Interfaces

Physical interfaces and subinterfaces allow traffic to enter and exit a security zone. To allow network traffic to flow in and out of a security zone, you must bind an interface to that zone and, if it is a Layer 3 zone, assign it an IP address. Then, you must configure policies to allow traffic to pass from interface to interface between zones. You can assign multiple interfaces to a zone, but you cannot assign a single interface to multiple zones.

This chapter contains the following sections:

- “Interface Types” on page 36
- “Viewing Interfaces” on page 43
- “Configuring Security Zone Interfaces” on page 44
- “Creating a Secondary IP Address” on page 50
- “Backup System Interfaces” on page 51
- “Loopback Interfaces” on page 58
- “Interface State Changes” on page 61

Interface Types

This section describes logical interfaces, function zone interfaces, and tunnel interfaces. For information about viewing a table of all these interfaces, see “Viewing Interfaces” on page 43.

Logical Interfaces

The purpose of logical interfaces is to provide an opening through which network traffic can pass between zones. ScreenOS supports the following types of logical Interfaces:

- Physical Interfaces
- Wireless Interfaces
- Bridge Group Interfaces
- Subinterfaces
- Aggregate Interfaces
- Redundant Interfaces
- Virtual Security Interfaces

Physical Interfaces

The name of a physical interface is composed of the media type, slot number (for some devices), and index number, for example, *ethernet3/2*, *ethernet0/2*, *wireless2*, *wireless0/2*, *bgroup2*, *serial0/0*, *serial0/2*, *bri0/0*, or *adsl0/2*. You can bind a physical interface to any security zone where it acts as a doorway through which traffic enters and exits the zone. Without an interface, no traffic can enter or leave the zone.

On security devices that support changes to interface-to-zone bindings, three of the physical ethernet interfaces are prebound to specific Layer 3 security zones—Trust, Untrust, and DMZ. Which interface is bound to which zone is specific to each platform. (For more information about security zones, see “Configuring Security Zone Interfaces” on page 44.)

Wireless Interfaces

A wireless interface, like a physical interface, acts as a doorway through which traffic enters and exits a security zone. Each wireless security device allows up to four wireless interfaces (*wireless0/0* — *wireless0/3*) to be active simultaneously.

A wireless interface cannot be bound to the Untrust security zone. (For more information, see “Wireless Local Area Network” on page 12-117.)

Bridge Group Interfaces

Some security devices support bridge groups (bgroups). Bgroups let you group multiple Ethernet and wireless interfaces together. Each bgroup constitutes its own broadcast domain and provides high-speed Ethernet switching between interfaces within the group. You can assign a single IP address to each bgroup interface. You can bind a bgroup interface to any zone.

For devices that are preconfigured with bridge groups, there are two different bridge group numbering systems. On some devices, the preconfigured bgroup interfaces are identified as bgroup0 through bgroup3. On other devices the preconfigured bgroup interfaces are identified as bgroup0/0 through bgroup0/2.

On some devices that support field-installable modules such as PIMs or mini-PIMs, you can create bgroups containing some or all of the Ethernet ports on the module. On these devices, you create bgroups identified as bgroupx/y, where x is the slot number for the module containing the grouped ports and y is a number you assign when creating the bridge group.

You can bind a bridge group interface to any zone. (For more information, see “Binding an Interface to a Security Zone” on page 44.)

Subinterfaces

A subinterface, like a physical interface, acts as a doorway through which traffic enters and exits a security zone. You can logically divide a physical interface into several virtual subinterfaces. Each virtual subinterface borrows the bandwidth it needs from the physical interface from which it stems, thus its name is an extension of the physical interface name, for example, *ethernet3/2.1* or *ethernet0/2.1*.

You can bind a subinterface to any Layer 3 zone. You can bind a subinterface to the same zone as its physical interface, or you can bind it to a different zone. (For more information, see “Binding an Interface to a Security Zone” on page 44.)

Aggregate Interfaces

Some security devices support aggregate interfaces. An aggregate interface is the accumulation of two or more physical interfaces that share the traffic load directed to the IP address of the aggregate interface equally among themselves. By using an aggregate interface, you can increase the amount of bandwidth available to a single IP address. Also, if one member of an aggregate interface fails, the other member or members can continue processing traffic—although with less bandwidth than previously available.

NOTE: For more information about aggregate interfaces, see “Interface Redundancy and Failover” on page 11-43.

Redundant Interfaces

You can bind two physical interfaces together to create one redundant interface, which you can then bind to a security zone. One of the two physical interfaces acts as the primary interface and handles all the traffic directed to the redundant interface. The other physical interface acts as the secondary interface and stands by

in case the active interface experiences a failure. If that occurs, traffic to the redundant interface fails over to the secondary interface, which becomes the new primary interface. The use of redundant interfaces provides a first line of redundancy before escalating a failover to the device level.

NOTE: For more information about redundant interfaces, see “Interface Redundancy and Failover” on page 11-43.

Virtual Security Interfaces

Virtual security interfaces (VSIs) are the virtual interfaces that two security devices forming a virtual security device (VSD) share when operating in HA mode. Network and VPN traffic use the IP address and virtual MAC address of a VSI. The VSD then maps the traffic to the physical interface, subinterface, or redundant interface to which you have previously bound the VSI. When two security devices are operating in HA mode, you must bind security zone interfaces that you want to provide uninterrupted service in the event of a device failover to one or more virtual security devices (VSDs). When you bind an interface to a VSD, the result is a virtual security interface (VSI).

NOTE: For more information about VSIs and how they function with VSDs in an HA cluster, see *Volume 11: High Availability*.

Function Zone Interfaces

Function zone interfaces, such as Management and HA, serve a special purpose.

Management Interfaces

On some security devices, you can manage the device through a separate physical interface—the Management (MGT) interface—moving administrative traffic outside the regular network user traffic. Separating administrative traffic from network user traffic greatly increases security and ensures constant management bandwidth.

NOTE: For information about configuring the device for administration, see “Administration” on page 3-1.

High Availability Interfaces

The high availability (HA) interface is a physical port used exclusively for HA functions. In a redundant group, one unit serves as the primary device—performing network firewall, VPN, and traffic-shaping functions—while the other unit serves as the backup device, waiting to take over the firewall functions when the primary unit fails. This is an active/passive configuration. You can also set up both members of the cluster to be primary and backup for each other. This is an active/active configuration. Both configurations are explained fully in *Volume 11: High Availability*.

Virtual HA Interfaces

On security devices without a dedicated HA interface, a virtual HA interface provides the same functionality. Because there is no separate physical port used exclusively for HA traffic, you must bind the virtual HA interface to one of the physical ethernet ports. You use the same procedure for binding a network interface to the HA zone as you do for binding a network interface to a security zone.

Tunnel Interfaces

A tunnel interface acts as a doorway to a VPN tunnel. Traffic enters and exits a VPN tunnel via a tunnel interface.

When you bind a tunnel interface to a VPN tunnel, you can reference that tunnel interface in a route to a specific destination and then reference that destination in one or more policies. With this approach, you can finely control the flow of traffic through the tunnel. It also provides dynamic routing support for VPN traffic. When there is no tunnel interface bound to a VPN tunnel, you must specify the tunnel in the policy itself and choose **tunnel** as the action. Because the action **tunnel** implies permission, you cannot specifically deny traffic from a VPN tunnel.

You can perform policy-based NAT on outgoing or incoming traffic using a pool of dynamic IP (DIP) addresses in the same subnet as the tunnel interface. A typical reason for using policy-based NAT on a tunnel interface is to avoid IP address conflicts between the two sites on either end of the VPN tunnel.

You must bind a route-based VPN tunnel to a tunnel interface so that the security device can route traffic to and from it. You can bind a route-based VPN tunnel to a tunnel interface that is either numbered (with IP address/netmask) or unnumbered (without IP address/netmask). If the tunnel interface is unnumbered, you must specify an interface from which the tunnel interface borrows an IP address. The security device only uses the borrowed IP address as a source address when the security device itself initiates traffic—such as OSPF messages—through the tunnel. The tunnel interface can borrow the IP address from an interface in the same security zone or from an interface in a different one as long as both zones are in the same routing domain.

You can achieve very secure control of VPN traffic routing by binding all the unnumbered tunnel interfaces to one zone, which is in its own virtual routing domain, and borrowing the IP address from a loopback interface bound to the same zone. For example, you can bind all the unnumbered tunnel interfaces to a user-defined zone named “VPN” and configure them to borrow an IP address from the loopback.1 interface, also bound to the VPN zone. The VPN zone is in a user-defined routing domain named “vpn-vr.” You put all destination addresses to which the tunnels lead in the VPN zone. Your routes to these addresses point to the tunnel interfaces, and your policies control VPN traffic between other zones and the VPN zone, see Figure 17.

Figure 17: Unnumbered Tunnel Interface Bindings

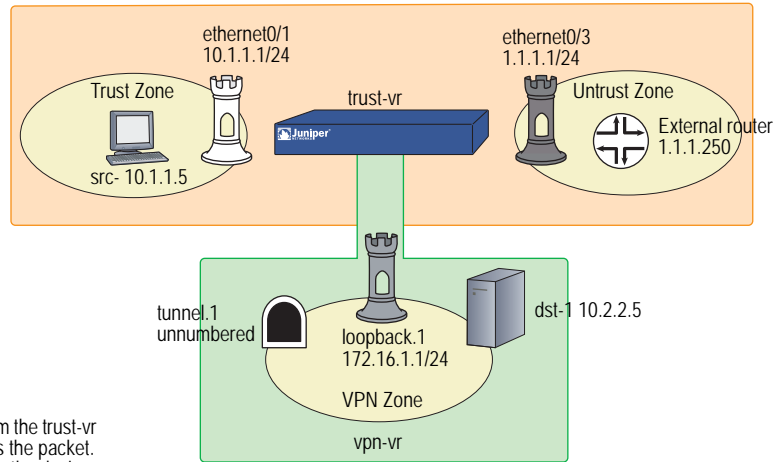
```
set router name vpn-vr
set zone name vpn vrouter vpn-vr
set interface loopback.1 zone vpn
set interface loopback.1 ip 172.16.1.1/24
set interface tunnel.1 zone vpn
set interface tunnel.1 ip unnumbered loopback.1
```

Configure addresses for src-1 and dst-1.
Configure a VPN tunnel and bind it to tunnel.1.

```
set router trust-vr route 10.2.2.5/32 vrouter vpn-vr
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3
gateway 1.1.1.250
set vrouter vpn-vr route 10.2.2.5 interface tunnel.1
```

```
set policy from trust to vpn scr-1 dst-1 any permit
```

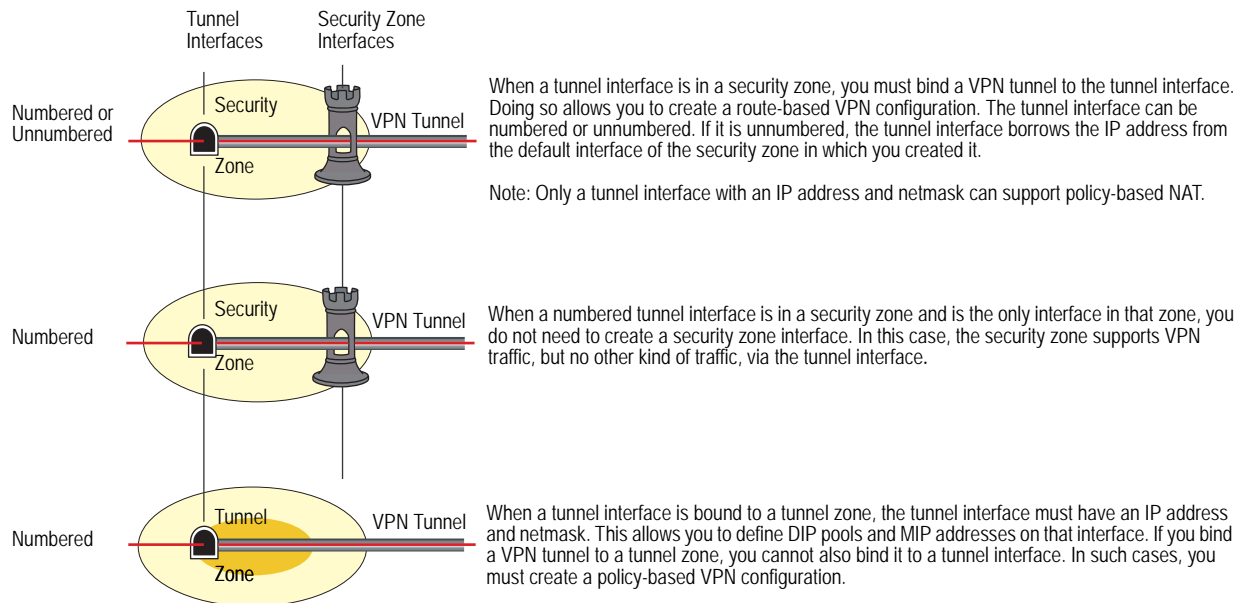
The security device sends traffic destined for 10.2.2.5/32 from the trust-vr to the vpn-vr. If tunnel.1 becomes disabled, the device drops the packet. Because the default route (to 0.0.0.0/0) is only in the trust-vr, the device does not attempt to send the packet in plain text out ethernet0/3.



Putting all the tunnel interfaces in such a zone is very secure because there is no chance for the failure of a VPN, which causes the route to the associated tunnel interface to become inactive, to redirect traffic intended for tunneling to use a non-tunneled route—such as the default route. (For several suggestions about how to avoid such a problem, see “Route-Based Virtual Private Network Security Considerations” on page 5-71.)

You can also bind a tunnel interface to a tunnel zone. When you do, it must have an IP address. The purpose of binding a tunnel interface to a tunnel zone is to make NAT services available for policy-based VPN tunnels. See Figure 18 on page 41.

NOTE: Network address translation (NAT) services include dynamic IP (DIP) pools and mapped IP (MIP) addresses defined in the same subnet as an interface.

Figure 18: Tunnel Interface to Zone Binding

Conceptually, you can view VPN tunnels as pipes that you have laid. They extend from the local device to remote gateways, and the tunnel interfaces are the openings to these pipes. The pipes are always there, available for use whenever the routing engine directs traffic to one of their interfaces.

Generally, assign an IP address to a tunnel interface if you want the interface to support one or more dynamic IP (DIP) pools for source address translation (NAT-src) and mapped IP (MIP) addresses for destination address translation (NAT-dst). For more information about VPNs and address translation, see “VPN Sites with Overlapping Addresses” on page 5-140. You can create a tunnel interface with an IP address and netmask in either a security or tunnel zone.

If the tunnel interface does not need to support address translation, and your configuration does not require the tunnel interface to be bound to a tunnel zone, you can specify the interface as unnumbered. You must bind an unnumbered tunnel interface to a security zone; you cannot bind it to a tunnel zone. You must also specify an interface with an IP address that is in the same virtual routing domain as the security zone to which the unnumbered interface is bound. The unnumbered tunnel interface borrows the IP address from that interface.

NOTE: For examples showing how to bind a tunnel interface to a tunnel, see the route-based VPN examples in “Site-to-Site Virtual Private Networks” on page 5-79 and “Dialup Virtual Private Networks” on page 5-159.

If you are transmitting multicast packets through a VPN tunnel, you can enable Generic Routing Encapsulation (GRE) on the tunnel interfaces to encapsulate multicast packets in unicast packets. Juniper Networks security devices support GREv1 for encapsulating IP packets in IPv4 unicast packets. For additional information about GRE, see “Configuring Generic Routing Encapsulation on Tunnel Interfaces” on page 7-151.

Deleting Tunnel Interfaces

You cannot immediately delete a tunnel interface that hosts mapped IP addresses (MIPs) or dynamic IP (DIP) address pools. Before you delete a tunnel interface hosting any of these features, you must first delete any policies that reference them. Then you must delete the MIPs and DIP pools on the tunnel interface. Also, if a route-based VPN configuration references a tunnel interface, you must first delete the VPN configuration before you can delete the tunnel interface.

In this example, tunnel interface `tunnel.2` is linked to DIP pool 8. DIP pool 8 is referenced in a policy (ID 10) for VPN traffic from the Trust zone to the Untrust zone through a VPN tunnel named `vpn1`. To remove the tunnel interface, you must first delete the policy (or remove the reference to DIP pool 8 from the policy) and then the DIP pool. Then, you must unbind `tunnel.2` from `vpn1`. After removing all the configurations that depend on the tunnel interface, you can then delete it.

WebUI

1. **Deleting Policy 10, Which References DIP Pool 8**
Policy > Policies (From: Trust, To: Untrust): Click **Remove** for Policy ID 10.
2. **Deleting DIP Pool 8, Which Is Linked to tunnel.2**
Network > Interfaces > Edit (for `tunnel.2`) > DIP: Click **Remove** for DIP ID 8.
3. **Unbinding tunnel.2 from vpn1**
VPNs > AutoKey IKE > Edit (for `vpn1`) > Advanced: Select **None** in the Bind to: Tunnel Interface drop-down list, click **Return**, then click **OK**.
4. **Deleting tunnel.2**
Network > Interfaces: Click **Remove** for `tunnel.2`.

CLI

1. **Deleting Policy 10, Which References DIP Pool 8**
`unset policy 10`
2. **Deleting DIP Pool 8, Which Is Linked to tunnel.2**
`unset interface tunnel.2 dip 8`
3. **Unbinding tunnel.2 from vpn1**
`unset vpn vpn1 bind interface`
4. **Deleting tunnel.2**
`unset interface tunnel.2`
`save`

Viewing Interfaces

You can view a table that lists all interfaces on your security device. Because they are predefined, physical interfaces are listed regardless of whether or not you configure them. Subinterfaces and tunnel interfaces are only listed once you create and configure them.

To view the interface table in the WebUI, click **Network > Interfaces**. You can specify the types of interfaces to display from the List Interfaces drop-down list.

To view the interface table in the CLI, use the **get interface** command.

The interface table displays the following information about each interface:

- **Name:** This field identifies the name of the interface.
- **IP/Netmask:** This field identifies the IP address and netmask address of the interface.
- **Zone:** This field identifies the zone to which the interface is bound.
- **Type:** This field indicates the interface type: Layer 2, Layer 3, tunnel, redundant, aggregate, VSI.
- **Link:** This field identifies whether the interface is active (up) or inactive (down).
- **Configure:** This field allows you modify or remove interfaces.

Figure 19: WebUI Interface Table

Name	Tag	IP/Netmask	Zone	Type	Link	PPPoE	Configure
bgroup5/0	-	0.0.0.0/0	V1-Untrust	Layer2	Down	-	Edit Remove
bgroup5/1	-	0.0.0.0/0	DMZ	Layer3	Down	-	Edit Remove
ethernet5/4	-				Down	-	Edit
ethernet5/5	-				Down	-	Edit
ethernet5/6	-				Down	-	Edit
ethernet5/13	-				Down	-	Edit
bgroup5/2	-	0.0.0.0/0	Untrust	Layer3	Down	-	Edit Remove
ethernet5/7	-				Down	-	Edit
ethernet5/8	-				Down	-	Edit
bgroup5/3	-	0.0.0.0/0	Null	Unused	Down	-	Edit Remove
bgroup5/4	-	0.0.0.0/0	Null	Unused	Down	-	Edit Remove
bgroup5/5	-	0.0.0.0/0	Null	Unused	Down	-	Edit Remove
bgroup5/6	-	0.0.0.0/0	Null	Unused	Down	-	Edit Remove
ethernet0/0	-	10.100.37.236/24	Trust	Layer3	Up	-	Edit
		-			Router		Edit

Figure 20: CLI Interface Table

```
ns500-> get interface

A - Active, I - Inactive, U - Up, D - Down, R - Ready

Interfaces in vsys Root:
Name          IP Address      Zone      MAC          VLAN State VSD Vsys
eth1/1        0.0.0.0/0       Null      0010.db0d.4ddc - D - Root
eth1/2        10.100.37.155/24 Untrust   0010.db0d.4dde - U - Root
eth2/1        0.0.0.0/0       Null      0010.db0d.4ddb - D - Root
eth2/2        1.1.2.5/24      Untrust   0010.db0d.4ddd - D - Root
eth3/1        2.2.2.0/24      Untrust   0010.db0d.4dd8 - D - Root
eth3/2        10.1.2.155/24   Trust     0010.db0d.4dda - U - Root
eth4/1        3.3.3.0/24      Untrust   0010.db0d.4dd7 - D - Root
eth4/2        0.0.0.0/0       Null      0010.db0d.4dd9 - D - Root
mgt           0.0.0.0/0       MGT       0010.db0d.4dd0 - D - Root
ha1           0.0.0.0/0       HA        0010.db0d.4dd5 - D - Root
ha2           0.0.0.0/0       HA        0010.db0d.4dd6 - D - Root
vlan1         0.0.0.0/0       VLAN      0010.db0d.4ddf 1 D - Root
null          0.0.0.0/0       Null      0010.dbff.0100 - U 0 Root
ns500->
```

Configuring Security Zone Interfaces

This section describes how to configure the following aspects of security zone interfaces:

- Binding and unbinding an interface to a security zone
- Assigning an address to a Layer 3 (L3) security zone interface
- Modifying physical interfaces and subinterfaces
- Creating subinterfaces
- Deleting subinterfaces

NOTE: For information about setting traffic bandwidth for an interface, see “Traffic Shaping” on page 193. For more information about the management and other services options available per interface, see “Controlling Administrative Traffic” on page 3-28.

Binding an Interface to a Security Zone

You can bind some physical interfaces to either an L2 or L3 security zone. WAN interfaces, except for ADSL, cannot be bound to L2 security zones. You can bind a subinterface only to an L3 security zone because a subinterface requires an IP address. You can only assign an IP address to an interface after you have bound it to an L3 security zone. Wireless interfaces cannot be bound to the Untrust security zone.

Some security devices allow you to group multiple interfaces. Before adding an interface to a group, the interface must be set to the Null security zone. After interfaces are added to a group, the group interface must be assigned to a security zone for connection to be established.

In this example, you bind ethernet0/5 to the Trust zone.

WebUI

Network > Interfaces > Edit (for ethernet0/5): Select **Trust** from the Zone Name drop-down list, then click **Apply**.

CLI

```
set interface ethernet0/5 zone trust
save
```

In this example, you set ethernet0/3 and ethernet0/4 to be in the Null security zone, group the interfaces in bgroup1, then bind the group to the DMZ security zone:

WebUI

Network > Interfaces > Edit (for ethernet0/3): Select **Null** from the Zone Name drop-down list, then click **Apply**.

Network > Interfaces > Edit (for ethernet0/4): Select **Null** from the Zone Name drop-down list, then click **Apply**.

Network > Interfaces > Edit (for bgroup1): Select **DMZ** from the Zone Name drop-down list, then click **Apply**.

Network > Interfaces > Edit (for bgroup1): Check **ethernet0/3** and **ethernet0/4** in the Bind to Current Bgroup column, then click **Apply**.

CLI

```
set interface ethernet0/3 zone null
set interface ethernet0/4 zone null
set interface bgroup1 port ethernet0/3
set interface bgroup1 port ethernet0/4
set interface bgroup1 zone DMZ
save
```

Unbinding an Interface from a Security Zone

If an interface is unnumbered, you can unbind it from one security zone and bind it to another. If an interface is numbered, you first must set its IP address and netmask to 0.0.0.0. Then, you can unbind it from one security zone and bind it to another one, and (optionally) reassign it an IP address/netmask.

In this example, ethernet0/3 has the IP address 210.1.1.1/24 and is bound to the Untrust zone. You set its IP address and netmask to 0.0.0.0/0 and bind it to the Null zone.

WebUI

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

```
Zone Name: Null
IP Address/Netmask: 0.0.0.0/0
```

CLI

```
set interface ethernet0/3 ip 0.0.0.0/0
set interface ethernet0/3 zone null
save
```

To unbind an interface from a group and reassign it to a different security zone, the interface must be released from the bgroup. Releasing the interface from the bgroup puts the interface in the Null security zone. Once in the Null security zone, the interface can be bound to any security zone then configured with an IP address.

WebUI

Network > Interfaces > Edit (for bgroup1) > Bind Port: Deselect **ethernet0/3** in the Bind to Current Bgroup column, then click **Apply**.

Network > Interfaces > Edit (for ethernet0/3): Select **Trust** from the Zone Name drop-down list, then click **Apply**.

CLI

```
unset interface bgroup1 port ethernet0/3
set interface ethernet0/3 zone trust
save
```

Addressing an L3 Security Zone Interface

When defining a Layer 3 (L3) security zone interface or subinterface, you must assign it an IP address and a netmask. If you bind the interface to a zone in the trust-vr, you can also specify the interface mode as NAT or Route. (If the zone to which you bind the interface is in the untrust-vr, the interface is always in Route mode.)

NOTE: For examples of NAT and Route mode configurations, see “Interface Modes” on page 79.

The two basic types of IP addresses to be considered when making interface address assignments are as follows:

- Public addresses, which Internet service providers (ISPs) supply for use on a public network like the Internet and which must be unique
- Private addresses, which a local network administrator assigns for use on a private network and which other administrators can assign for use on other private networks as well

NOTE: When you add an IP address to an interface, the security device checks via an ARP request to make sure that the IP address does not already exist on the local network. (The physical link must be up at the time.) If the IP address already exists, a warning is displayed.

Public IP Addresses

If an interface connects to a public network, it must have a public IP address. Also, if an L3 security zone in the untrust-vr connects to a public network and the interfaces of zones in the trust-vr are in Route mode, then all the addresses in the zones in the trust-vr—for interfaces and for hosts—must also be public addresses. Public IP addresses fall into three classes, A, B, and C, as shown in Table 4.

Table 4: Public Address Ranges

Address Class	Address Range	Excluded Address Range
A	0.0.0.0 – 127.255.255.255	10.0.0.0 – 10.255.255.255, 127.0.0.0 – 127.255.255.255
B	128.0.0.0 – 191.255.255.255	172.16.0.0 – 172.31.255.255
C	192.0.0.0 – 223.255.255.255	192.168.0.0 – 192.168.255.255

NOTE: There are also D and E class addresses, which are reserved for special purposes.

An IP address is composed of four octets, each octet being 8 bits long. In a class A address, the first 8 bits indicate the network ID, and the final 24 bits indicate the host ID (nnn.hhh.hhh.hhh). In a class B address, the first 16 bits indicate the network ID, and the final 16 bits indicate the host ID (nnn.nnn.hhh.hhh). In a class C address, the first 24 bits indicate the network ID, and the final 8 bits indicate the host ID (nnn.nnn.nnn.hhh).

Through the application of subnet masks (or netmasks), you can further divide networks. A netmask essentially masks part of the host ID so that the masked part becomes a subnet of the network ID. For example, the 24-bit mask in the address 10.2.3.4/24 indicates that the first 8 bits (that is, the first octet—010) identify the network portion of this private class A address, the next 16 bits (that is, the second and third octets—002.003) identify the subnetwork portion of the address, and the last 8 bits (the last octet—004) identify the host portion of the address. Using subnets to narrow large network address spaces into smaller subdivisions greatly increases the efficient delivery of IP data grams.

NOTE: The dotted-decimal equivalent of a 24-bit mask is 255.255.255.0.

Private IP Addresses

If an interface connects to a private network, a local network administrator can assign it any address, although it is conventional to use an address from the range of addresses reserved for private use—10.0.0.0/8, 172.16.0.0 – 172.31.255.255, 192.168.0.0/16— as defined in RFC 1918, *Address Allocation for Private Internets*.

If an L3 security zone in the untrust-vr connects to a public network and the interfaces bound to zones in the trust-vr are in NAT mode, then all the addresses in the zones in the trust-vr—for interfaces and for hosts—can be private addresses.

Addressing an Interface

In this example, you assign ethernet0/5 the IP address 210.1.1.1/24 and give it the Manage IP address 210.1.1.5. (Note that the Manage IP address must be in the same subnet as the security zone interface IP address.) Finally, you set the interface in NAT mode, which translates all internal IP addresses to the default interfaces bound to the other security zones.

NOTE: The default interface in a security zone is the first interface bound to the zone. To learn which interface is the default interface for a zone, see the Default IF column on Network > Zones in the WebUI, or the Default-If column in the output from the **get zone** command in the CLI.

WebUI

Network > Interfaces > Edit (for ethernet0/5): Enter the following, then click **OK**:

IP Address/Netmask: 210.1.1.1/24
Manage IP: 210.1.1.5

CLI

```
set interface ethernet0/5 ip 210.1.1.1/24
set interface ethernet0/5 manage-ip 210.1.1.5
save
```

Modifying Interface Settings

After you have configured a physical interface, a subinterface, a redundant interface, an aggregate interface, or a Virtual Security Interface (VSI), you can later change any of the following settings should the need arise:

- IP address and netmask.
- Manage IP address.
- (L3 zone interfaces) Management and network services.
- (Subinterface) Subinterface ID number and VLAN tag number.
- (Interfaces bound to L3 security zones in the trust-vr) Interface mode—NAT or Route.
- (Physical interface) Traffic bandwidth settings (see “Traffic Shaping” on page 193).
- (Physical, redundant, and aggregate interfaces) Maximum Transmission Unit (MTU) size.
- (L3 interfaces) Block traffic from coming in and going out the same interface, including traffic between a primary and secondary subnet or between secondary subnets (this is done using the CLI **set interface** command with the **route-deny** option).

For physical interfaces on some security devices, you can force the physical state of the link to be down or up. By forcing the physical state of the link to be down, you can simulate a disconnect of the cable from the interface port. (This is done with the CLI **set interface** command with the **phy link-down** option.)

In this example, you make some modifications to ethernet0/1, an interface bound to the Trust zone. You change the Manage IP address from 10.1.1.2 to 10.1.1.12. To enforce tighter security of administrative traffic, you also change the management services options, enabling SCS and SSL and disabling Telnet and WebUI.

WebUI

Network > Interfaces > Edit (for ethernet0/1): Make the following modifications, then click **OK**:

Manage IP: 10.1.1.12
Management Services: (select) SSH, SSL; (clear) Telnet, WebUI

CLI

```
set interface ethernet0/1 manage-ip 10.1.1.12
set interface ethernet0/1 manage ssh
set interface ethernet0/1 manage ssl
unset interface ethernet0/1 manage telnet
unset interface ethernet0/1 manage web
save
```

Creating a Subinterface in the Root System

You can create a subinterface on any physical interface in the root system or virtual system. A subinterface makes use of VLAN tagging to distinguish traffic bound for it from traffic bound for other interfaces. Note that although a subinterface stems from a physical interface, from which it borrows the bandwidth it needs, you can bind a subinterface to any zone, not necessarily that to which its “parent” interface is bound. Additionally, the IP address of a subinterface must be in a different subnet from the IP addresses of all other physical interfaces and subinterfaces.

NOTE: You can also configure subinterfaces on redundant interfaces and VSIs. For an example that includes the configuration of a subinterface on a redundant interface, see “Virtual System Failover” on page 11-96.

In this example, you create a subinterface for the Trust zone in the root system. You configure the subinterface on ethernet0/1, which is bound to the Trust zone. You bind the subinterface to a user-defined zone named “accounting,” which is in the trust-vr. You assign it subinterface ID 3, IP address 10.2.1.1/24, and VLAN tag ID 3. The interface mode is NAT.

WebUI

Network > Interfaces > New Sub-IF: Enter the following, then click **OK**:

```
Interface Name: ethernet0/1.3
Zone Name: accounting
IP Address/Netmask: 10.2.1.1/24
VLAN Tag: 3
```

CLI

```
set interface ethernet0/1.3 zone accounting
set interface ethernet0/1.3 ip 10.2.1.1/24 tag 3
save
```

Deleting a Subinterface

You cannot immediately delete a subinterface that hosts mapped IP addresses (MIPs), virtual IP addresses (VIPs), or dynamic IP (DIP) address pools. Before you delete a subinterface hosting any of these features, you must first delete any policies or IKE gateways that reference them. Then you must delete the MIPs, VIPs, and DIP pools on the subinterface.

In this example, you delete the subinterface ethernet0/1.1.

WebUI

Network > Interfaces: Click **Remove** for ethernet0/1.1.

A system message prompts you to confirm the removal.

Click **Yes** to delete the subinterface.

CLI

```
unset interface ethernet0/1.1
save
```

Creating a Secondary IP Address

Each ScreenOS interface has a single, unique *primary* IP address. However, some situations demand that an interface have multiple IP addresses. For example, an organization might have additional IP address assignments and might not wish to add a router to accommodate them. In addition, an organization might have more network devices than its subnet can handle, as when there are more than 254 hosts connected to a LAN. To solve such problems, you can add *secondary* IP addresses to an interface in the Trust, DMZ, or user-defined zone.

NOTE: You cannot set multiple secondary IP addresses for interfaces in the Untrust zone.

Secondary addresses have certain properties that affect how you can implement such addresses. These properties are as follows:

- There can be no subnet address overlap between any two secondary IP addresses. In addition, there can be no subnet address overlap between a secondary IP and any existing subnet on the security device.
- When you manage a security device through a secondary IP address, the address always has the same management properties as the primary IP address. Consequently, you cannot specify a separate management configuration for the secondary IP address.
- You cannot configure a gateway for a secondary IP address.
- Whenever you create a new secondary IP address, the security device automatically creates a corresponding routing table entry. When you delete a secondary IP address, the device automatically deletes its routing table entry.

Enabling or disabling routing between two secondary IP addresses causes no change in the routing table. For example, if you disable routing between two such addresses, the security device drops any packets directed from one interface to the other, but no change occurs in the routing table.

In this example, you set up a secondary IP address—192.168.2.1/24—for ethernet0/1, an interface that has IP address 10.1.1.1/24 and is bound to the Trust zone.

WebUI

Network > Interfaces > Edit (for ethernet0/1) > Secondary IP: Enter the following, then click **Add**:

IP Address/Netmask: 192.168.2.1/24

CLI

```
set interface ethernet0/1 ip 192.168.2.1/24 secondary
save
```

Backup System Interfaces

The interface backup feature allows you to configure a backup interface that can take over traffic from a configured primary interface. You can back up any type of interface with any other type of interface supported on the platform. The only requirement is that both interfaces must be in the untrust zone.

You set a backup interface so that the security device can switch traffic over to it in the event that the primary interface goes down (is unplugged, or fails), destinations on the primary interface become unreachable, or the tunnel bound to the primary interface becomes inactive. When the connection through the primary interface is restored, ScreenOS automatically switches traffic from the backup interface to the primary. The interface backup feature also provides a way for you manually to force the primary interface to switch over to the backup, and to force the backup to switch over to the primary. Each primary interface can have only one backup interface, and each backup interface can have only one primary.

You can configure the security device to switch over to the backup interface when any of the following conditions are met on the primary interface:

- Certain IP addresses become unreachable through the interface.
- Certain VPN tunnels on the interface become unreachable.
- A preconfigured route becomes unreachable through the interface.

For the security device to switch traffic over to a backup interface for any of these reasons, you must first configure the primary interface for that purpose, then configure the backup interface accordingly. You must also configure two default routes, one for the primary interface and one for the backup interface. You can configure the backup interface feature through the WebUI or at the CLI.

Configuring a Backup Interface

ScreenOS determines when to switch over to a backup interface by tracking or monitoring activity on the primary interface. You can configure the following types of backup interfaces:

- IP Tracking
- Tunnel-if tracking
- Route monitoring

Configuring an IP Tracking Backup Interface

In this example, you configure ScreenOS to track IP address 10.1.1.1 on the primary interface (ethernet0/0), and to switch over to the backup interface (dialer1), in the event that this address becomes unreachable. For a discussion of how IP tracking works, see “Interface Failover with IP Tracking” on page 11-50.

WebUI

1. Configure Interfaces

Network > Interfaces > (for ethernet0/0) Edit > Monitor > Add: Enter the following, then click **Apply**:

```
Track IP: 10.1.1.1
Weight: 200
Interval: 2
Threshold: 5
```

Network > Interfaces > Backup: Enter the following, then click **Apply**:

```
Primary interface (select): ethernet0/0
Backup Interface (select): dialer1
```

2. Configure Routes

Network > Routing > Destination > trust-tr New: Enter the following, then click **Apply**:

IP Address/Netmask: 0.0.0.0/0
Interface (select): ethernet0/0

Network > Routing > Destination > trust-tr New: Enter the following, then click **Apply**:

IP Address/Netmask: 0.0.0.0/0
Interface (select): dialer1

CLI

1. Configure Interfaces

```
set interface ethernet0/0 monitor track-ip
set interface ethernet0/0 monitor track-ip threshold 100
set interface ethernet0/0 monitor trackip ip 10.1.1.1 interval 2
set interface ethernet0/0 monitor trackip ip 10.1.1.1 threshold 5
set interface ethernet0/0 monitor trackip ip 10.1.1.1 weight 200
set interface ethernet0/0 backup interface dialer1 type track-ip
```

2. Configure Routes

```
set route 0.0.0.0/0 interface ether0/0
set route 0.0.0.0/0 interface dialer1
save
```

Configuring a Tunnel-if Backup Interface

In this example, you configure a pair of unidirectional VPN tunnels on ethernet0/0—one on Router-1 and one on Router-2—and you configure dialer1 as the backup interface on Route-1. The tunnels connect hosts in the Trust zone at a branch site to an SMTP server in the Trust zone at the corporate site. The zones at each site are in the trust-vr routing domain.

You configure both tunnels with the primary Untrust zone interface (ethernet0/0) as the outgoing interface, and the backup VPN tunnel with the backup Untrust zone interface (dialer1) as the outgoing interface. The security device monitors the primary VPN tunnels to determine when to switch over to the backup. It does this by comparing the backup weight with the VPN monitor threshold. You set the threshold to 100 (**set vpnmonitor threshold 100**) and the backup weight to 200 (**set vpn vpn backup-weight 200**). When the primary interface becomes inactive for any reason, ScreenOS compares the VPN monitor threshold with the backup weight, and if the backup weight is greater than the threshold, it switches the device to the backup.

You also enable the VPN monitor rekey feature. In the event of a failover, this feature enables the security device to revert traffic from the backup interface to the primary if the accumulated weight of the VPN tunnels on the primary interface becomes greater than the VPN monitor threshold.

The security device in the branch site receives its Untrust zone interfaces address, default gateway, and DNS server addresses dynamically from two different ISPs. Each ISP uses a different protocol. ISP-1 uses DHCP to assign an address to ethernet0/0, and ISP-2 uses PPP to assign an address to dialer1. The security device at the corporate site has a static IP address (2.2.2.2). The IP address of its default gateway is 2.2.2.250.

The destination address for VPN monitoring is not the default—the remote gateway IP address (2.2.2.2)—but the addresses of the server (10.2.2.10). If you use the remote gateway IP address and it becomes unreachable, the primary tunnel always switches over to the backup.

NOTE: Because this example is extensive, only the CLI configuration is included in its entirety. The WebUI section lists the navigational paths to the pages where you can set the various elements of the configuration. You can see what you need to set by referring to the CLI commands

WebUI (for Router-1)

1. Configure Interfaces

Network > Interfaces > Edit (for ethernet0/0)

Network > Interfaces > Edit (for bri1/0)

Network > Interfaces > New Tunnel IF

2. Configure VPN

VPNs > AutoKey IKE> New

VPNs > AutoKey Advanced > Gateway > New

3. Configure Asymmetric VPN

Network > Zones > Edit (for Trust)

4. Configure Backup Interface

Network > Interfaces > Backup

5. Configure Routes

Network > Routing > Destination > trust-vr

CLI (for Router-1)

1. Configure Interfaces

```
set interface ethernet0/0 zone untrust
set interface ethernet0/0 dhcp client
```

```
set interface bri2/0 isdn switch-type etsi
set interface dialer1 zone untrust
set dialer pool name pool-1
set interface bri2/0 dialer-pool-member pool-1
set interface dialer1 dialer-pool pool-1
```

```
set ppp profile isdn-ppp
set ppp profile isdn-ppp auth type chap
set ppp profile isdn-ppp auth local-name juniper
```

```

set ppp profile isdn-ppp auth secret juniper
set ppp profile isdn-ppp passive
set ppp dialer1 ppp profile isdn-ppp

set interface tunnel.1 untrust
set interface tunnel.1 ip unnumbered interface bgroup0
set interface tunnel.2 untrust
set interface tunnel.2 ip unnumbered interface bgroup0

```

2. Configure VPN

```

set ike gateway corp1 address 2.2.2.2 aggressive local-id ssg5ssg20-e0
  outgoing-interface ethernet0/0 preshare juniper1 sec-level basic
set ike gateway corp1 address 2.2.2.2 aggressive local-id ssg5ssg20-dialer
  outgoing-interface dialer1 preshare juniper2 sec-level basic

set vpn vpn1 gateway corp1 sec-level basic
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 smtp
set vpn vpn1 monitor source-interface bgroup0 destination-ip 10.2.2.10 rekey
set vpn vpn1 backup-weight 200

set vpn vpn2 gateway corp2 sec-level basic
set vpn vpn2 bind interface tunnel.2
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 smtp
set vpn vpn2 monitor source-interface bgroup0 destination-ip 10.2.2.10 rekey

```

3. Configure Asymmetric VPN

```
set zone trust asymmetric-vpn
```

4. Configure Backup Interface

```
set interface ethernet0/0 backup interface dialer1 type tunnel-if
set vpnmonitor threshold 100
```

5. Configure Routes

```
set route 10.2.2.10/32 interface tunnel.1
set route 10.2.2.10/32 interface tunnel.2
set route 0.0.0.0/0 interface ethernet0/0
save
```

WebUI (for Router-2)

1. Configure Interfaces

Network > Interfaces > Edit (for bgroup1)

Network > Interfaces > Edit (for ethernet0/0)

Network > Interfaces > New Tunnel IF

2. Configure Addresses

Policy > Policy Elements > Addresses > List > New

3. Configure VPN

VPNs > AutoKey IKE > New

VPNs > AutoKey Advanced > Gateway > New

4. Configure Asymmetric VPN

Network > Zones > Edit (for Trust)

Network > Interfaces > Edit (for ethernet0/0)

5. Configure Route

Network > Routing > Destination > trust-vr

6. Configure Policy

Policy > Policies (From Untrust to trust) > New

CLI (for Router-2)

1. Configure Interfaces

```
set interface bgroup zone trust
set interface bgroup ip 10.2.2.1/24
set interface bgroup nat
set interface ethernet0/0 zone untrust
set interface ethernet0/0 ip 2.2.2.2/24
set interface tunnel.1 zone trust
set interface tunnel.1 ip unnumbered interface bgroup0
set interface tunnel.2 zone untrust
set interface tunnel.2 ip unnumbered interface bgroup0
```

2. Configure Addresses

```
set address untrust branch 10.1.1.0/24
set address trust smtp-1 10.2.2.10/24
set address trust http-1 10.2.2.15/32
set group address trust servers add smtp-1
set group service vpn-srv add smtp
```

3. Configure VPN

```
set ike gateway branch1 dynamic ssg5ssg20-e0 aggressive outgoing-interface
ethernet0/0 preshare juniper1 sec-level basic
set ike gateway branch2 dynamic ssg5ssg20-dialer aggressive outgoing-interface
ethernet0/0 preshare juniper2 sec-level basic

set vpn vpn1 gateway branch1 sec-level basic
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 smtp
set vpn vpn2 gateway branch2 sec-level basic
set vpn vpn2 bind interface tunnel.2
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 smtp
```

4. Configure Asymmetric VPN

```
set zone trust asymmetric-vpn
```

5. Configure Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet 0/0 gateway 2.2.2.250
```

6. Configure Policy

```
set policy from untrust to trust branch servers vpn-srv permit
save
```

Configuring a Route Monitoring Backup Interface

In this example, on the primary interface (ethernet0/0) you configure a route to the network segment 5.5.5.0/24 using gateway 10.10.10.1; and you configure dialer1 as the backup interface, with a route to the same network segment. You then configure a default route for the primary interface to the same gateway (10.10.10.1), and a default route for the dialer1 interface.

WebUI

1. Configure Interfaces

Network > Routing > Destination > trust-vr > New: Enter the following, then click **Apply**:

IP Address/Netmask: 5.5.5.0/24
 Gateway: (select)
 Interface: (select) ethernet0/0
 Gateway IP Address: 10.10.10.1

Network > Interfaces > Backup: Enter the following, then click **Apply**:

Primary Interface: (select) ethernet0/0
 Backup Interface: (select) dialer1
 Type: (select) route vrouter: (select) trust-vr
 IP Address/Netmask: 5.5.5.0/24

2. Configure Route

Network > Routing > Source Interface (for ethernet0/0) > New: Enter the following, then click **OK**:

Gateway: 10.10.10.1

Network > Routing > Destination (for trust-vr) > New: Enter the following, then click **OK**:

Interface (select): Null

CLI

```
set vrouter trust-vr route 5.5.5.0/24 interface ethernet0/0 gateway 10.10.10.1
set interface ethernet0/0 backup interface dialer1 type route vrouter trust-vr
5.5.5.0/24
set route 0.0.0.0/0 interface ethernet0/0 gateway 10.10.10.1
set route 0.0.0.0/0 interface dialer1
save
```

Loopback Interfaces

A loopback interface is a logical interface that emulates a physical interface on the security device. However, unlike a physical interface, a loopback interface is always in the up state as long as the device on which it resides is up. Loopback interfaces are named `loopback.id_num`, where *id_num* is a number greater than or equal to 1 and denotes a unique loopback interface on the device. Like a physical interface, you must assign an IP address to a loopback interface and bind it to a security zone.

NOTE: The maximum *id_num* value you can specify is platform-specific.

After defining a loopback interface, you can then define other interfaces as members of its group. Traffic can reach a loopback interface if it arrives through one of the interfaces in its group. Any interface type can be a member of a loopback interface group—physical interface, subinterface, tunnel interface, redundant interface, or VSI.

After creating a loopback interface, you can use it in many of the same ways as a physical interface:

- Setting the Loopback Interface for Management
- Setting BGP on a Loopback Interface
- Setting VSIs on a Loopback Interface
- Setting the Loopback Interface as a Source Interface

NOTE: You cannot bind a loopback interface to an HA zone, nor can you configure a loopback interface for Layer 2 operation or as a redundant/aggregate interface. You cannot configure the following features on loopback interfaces: NTP, DNS, VIP, secondary IP, track IP, or WebAuth.

When a loopback interface is used as the outgoing Interface in a VPN configuration, then the loopback interface must be in the same zone as the outgoing physical interface.

You can define a MIP on a loopback interface. This allows the MIP to be accessed by a group of interfaces; this capability is unique to loopback interfaces. For information about using the loopback interface with MIPs, see “MIP and the Loopback Interface” on page 8-73.

You can manage the security device using either the IP address of a loopback interface or the manage IP address that you assign to a loopback interface.

Creating a Loopback Interface

In the following example, you create the loopback interface `loopback.1`, bind it to the `Untrust` zone, and assign the IP address `1.1.1.27/24` to it.

WebUI

Network > Interfaces > New Loopback IF: Enter the following, then click **OK**:

```
Interface Name: loopback.1
Zone: Untrust (select)
IP Address/Netmask: 1.1.1.27/24
```

CLI

```
set interface loopback.1 zone untrust
set interface loopback.1 ip 1.1.1.27
save
```

NOTE: The loopback interface is not directly accessible from networks or hosts that reside in other zones. You must define a policy to permit traffic to and from the interface.

Setting the Loopback Interface for Management

In the following example, you configure the previously defined `loopback.1` interface as a management interface for the device.

WebUI

Network > Interfaces > `loopback.1` > Edit: Select all the management options, then click **OK**.

CLI

```
set interface loopback.1 manage
save
```

Setting BGP on a Loopback Interface

The loopback interface can support the BGP dynamic routing protocol on the security device. In the following example, you enable BGP on the `loopback.1` interface.

NOTE: To enable BGP on the loopback interface, you must first create a BGP instance for the virtual router in which you plan to bind the interface. For information about configuring BGP on Juniper Networks security devices, See *Volume 7: Routing*.

WebUI

Network > Interfaces > `loopback.1` > Edit: Select **Protocol BGP**, then click **OK**.

CLI

```
set interface loopback.1 protocol bgp
save
```

Setting VSIs on a Loopback Interface

You can configure Virtual Security Interfaces (VSIs) for NSRP on a loopback interface. The physical state of the VSI on the loopback interface is always up. The interface can be active or not, depending upon the state of the VSD group to which the interface belongs.

WebUI

Network > Interfaces > New VSI IF: Enter the following, then click **OK**:

Interface Name: VSI Base: loopback.1
VSD Group: 1
IP Address/Netmask: 1.1.1.1/24

CLI

```
set interface loopback.1:1 ip 1.1.1.1/24
save
```

Setting the Loopback Interface as a Source Interface

You can use a loopback interface as a source interface for certain traffic that originates from the security device. (When you define a source interface for an application, the specified source interface address is used instead of the outbound interface address to communicate with an external device.) In the following example, you specify that the security device uses the previously defined loopback.1 interface for sending syslog packets.

WebUI

Configuration > Report Settings > Syslog: Enter the following, then click **Apply**:

Enable Syslog Messages: (select)
Source interface: loopback.1 (select)
Syslog Servers:
No.: 1 (select)
IP/Hostname: 10.1.1.1
Traffic Log: (select)
Event Log: (select)

CLI

```
set syslog config 10.1.1.1 log all
set syslog src-interface loopback.1
set syslog enable
save
```

Interface State Changes

An interface can be in one of the states described in Table 5:

Table 5: Interface States

State	Description
Physically Up	For physical ethernet interfaces operating at either Layer 2 (Transparent mode) or Layer 3 (Route Mode) in the Open Systems Interconnection (OSI) model. An interface is physically up when it is cabled to another network device and can establish a link to that device.
Logically Up	For both physical interfaces and logical interfaces (subinterfaces, redundant interfaces, and aggregate interfaces). An interface is logically up when traffic passing through that interface is able to reach specified devices (at tracked IP addresses) on a network.
Physically Down	An interface is physically down when it is not cabled to another network device or when it is cabled but cannot establish a link. You can also force an interface to be physically down with the following CLI command: set interface <i>interface</i> phy link-down .
Logically Down	An interface is logically down when traffic passing through that interface cannot reach specified devices (at tracked IP addresses) on a network.

The physical state of an interface takes precedence over its logical state. An interface can be physically up and—at the same time—be either logically up or logically down. If an interface is physically down, its logical state becomes irrelevant.

When the state of an interface is up, all routes that make use of that interface remain active and usable. When the state of an interface is down, the security device deactivates all routes using that interface—although, depending on whether the interface is physically or logically down, traffic might still flow through an interface whose state is down (see “Down Interfaces and Traffic Flow” on page 74). To compensate for the loss of routes caused by the loss of an interface, you can configure alternate routes using an alternate interface.

Depending on how you set up the action that an observed interface state change can cause, a state change from up to down in a monitored interface can cause the monitoring interface to change its state from down to up. To configure this behavior, you can use the following CLI command:

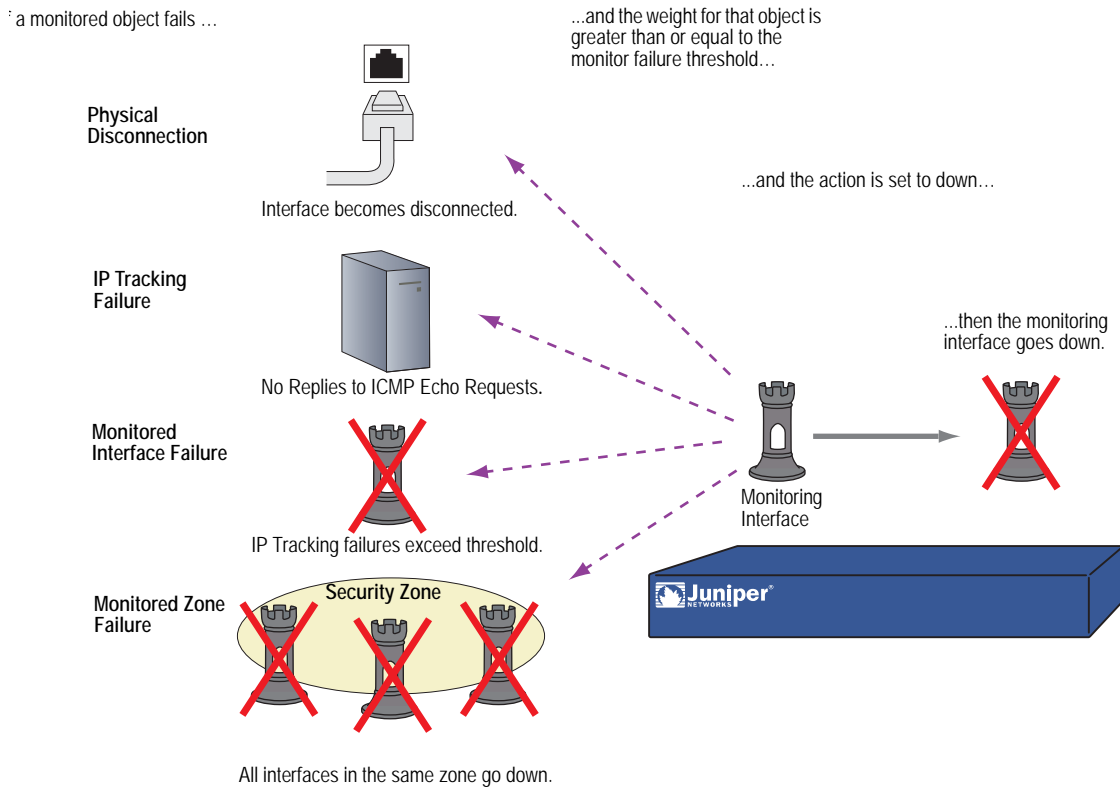
```
set interface interface monitor threshold number action up { logically | physically }
```

When you enter the above command, the security device automatically forces the monitoring interface into a down state. If the monitored object (tracked IP address, interface, zone) fails, then the state of the monitoring interface becomes up—either logically or physically, per your configuration.

An interface can monitor objects for one or more of the following events. See Figure 21 on page 62. Each of these events by itself or in combination can cause the state of the monitoring interface to change from up to down and from down to up:

- Physical disconnection/reconnection
- IP tracking failure/success
- Failure/success of a monitored interface
- Failure/success of a monitored security zone

Figure 21: Interface State Monitoring



If, after failing, a monitored object succeeds (the interface is reconnected or IP tracking again succeeds), then the monitoring interface comes back up. There is approximately a one-second delay between the monitored object succeeding and the monitoring interface reactivating.

Each of the above events is presented in the following sections.

Physical Connection Monitoring

All physical interfaces on a security device monitor the state of their physical connection to other network devices. When an interface is connected to and has established a link with another network device, its state is physically up, and all routes that use that interface are active.

You can see the state of an interface in the State column in the output of the **get interface** command and in the Link column on Network > Interfaces in the WebUI. It can be up or down.

You can see the state of a route in the status field of the **get route id number** command and on Network > Routing > Destination in the WebUI. If there is an asterisk, the route is active. If there is no asterisk, it is inactive.

Tracking IP Addresses

The security device can track specified IP addresses through an interface so that when one or more of them become unreachable, the security device can deactivate all routes associated with that interface, even if the physical link is still active. A deactivated route becomes active again after the security device regains contact with those IP addresses.

NOTE: For some ScreenOS appliances, this action also causes a failover to the backup interface that is bound to the same zone as the interface on which IP tracking is configured (see “Interface Failover with IP Tracking” on page 11-50).

ScreenOS uses Layer 3 path monitoring, or *IP tracking*, similar to that used for NSRP, to monitor the reachability of specified IP addresses through an interface. For example, if an interface connects directly to a router, you can track the next-hop address on the interface to determine if the router is still reachable. When you configure IP tracking on an interface, the security device sends ping requests on the interface to up to four target IP addresses at user-defined intervals. The security device monitors these targets to determine if it receives a response. If there is no response from a target for a specified number of times, that IP address is deemed to be unreachable. Failure to elicit a response from one or more targets can cause the security device to deactivate routes associated with that interface. If another route to the same destination is available, the security device then redirects traffic to use the new route.

You can define IP tracking on the following interfaces for which you have configured a manage IP address:

- Physical interface bound to a security zone (not the HA or MGT function zones)

NOTE: The interface can operate at Layer 2 (Transparent mode) or Layer 3 (Route mode).

- Subinterface
- Redundant interface
- Aggregate interface

NOTE: Although the interface can be a redundant interface or an aggregate interface, it cannot be a member of a redundant or aggregate interface.

On devices that support virtual systems, the interface on which you set IP tracking can belong to the root system or to a virtual system (vsys). However, to set IP tracking on a shared interface, you can only set it at the root level.

NOTE: From a vsys, you can set interface monitoring to monitor a shared interface from an interface that belongs to the vsys. However, from within a vsys, you cannot set interface monitoring from a shared interface. For more information, see “Interface Monitoring” on page 68.

For each interface, you can configure up to four IP addresses for the security device to track. On a single device, you can configure up to 64 track IP addresses. That total includes all track IP addresses whether they are for interface-based IP tracking, for NSRP-based IP tracking, at the root level, or at the vsys level.

The tracked IP addresses do not have to be in the same subnetwork as the interface. For each IP address to be tracked, you can specify the following:

- Interval, in seconds, at which the pings are sent to the specified IP address.
- Number of consecutive unsuccessful ping attempts before the connection to the IP address is considered failed.
- Weight of the failed IP connection (once the sum of the weights of all failed IP connections crosses a specified threshold, routes that are associated with the interface are deactivated).

You can also configure the security device to track the default gateway for an interface that is a PPPoE or DHCP client. To do that, use the “Dynamic” option: (CLI) **set interface *interface* monitor dynamic** or (WebUI) Network > Interfaces > Edit (for the DHCP or PPPoE client interface) > Monitor > Track IP > Add: Select **Dynamic**.

NOTE: When you configure an IP address for the security device to track, the security device does not add a host route for that IP address to the routing table.

There are two types of thresholds in configuring tracking IP addresses:


- Failure threshold for a specific tracked IP address — The number of consecutive failures to elicit a ping response from a specific IP address that constitutes a failure in reaching the IP address. Not exceeding the threshold indicates an acceptable level of connectivity with the address; exceeding the threshold indicates an unacceptable level. You set this threshold for each IP address at any value between 1 and 200. The default value is 3.

- Failure threshold for IP tracking on the interface — The total weight of the cumulative failed attempts to reach IP addresses on the interface that causes routes associated with the interface to be deactivated. You can set this threshold at any value between 1 and 255. The default value is 1, which means a failure to reach any configured tracked IP address causes routes associated with the interface to be deactivated.

By applying a *weight*, or a value, to a tracked IP address, you can adjust the importance of connectivity to that address in relation to reaching other tracked addresses. You can assign comparatively greater weights to relatively more important addresses, and less weight to relatively less important addresses. Note that the assigned weights only come into play when the failure threshold for a specific tracked IP address is reached. For example, if the failure threshold for IP tracking on an interface is 3, failure of a single tracked IP address with a weight of 3 meets the failure threshold for IP tracking on the interface, which causes routes associated with the interface to be deactivated. The failure of a single tracked IP address with a weight of 1 would not meet the failure threshold for IP tracking on the interface and routes associated with the interface would remain active.

In the following example, the interface ethernet0/1 is bound to the Trust zone and assigned the network address 10.1.1.1/24. The interfaces ethernet0/3 and ethernet0/4 are bound to the Untrust zone. The ethernet0/3 interface is assigned the network address 1.1.1.1/24 and is connected to the router at 1.1.1.250. The ethernet0/4 interface is assigned the network address 2.2.2.1/24 and is connected to the router at 2.2.2.250. See Figure 22.

Figure 22: Interface IP Tracking



There are two default routes configured: one uses ethernet0/3 as the outbound interface with the router address 1.1.1.250 as the gateway; the other uses ethernet0/4 as the outbound interface with the router address 2.2.2.250 as the gateway and is configured with a metric value of 10. The default route that uses ethernet0/3 is the preferred route since it has a lower metric (the default metric value for static routes is 1). The following output from the **get route** command shows four active routes for the trust-vr (active routes are denoted with an asterisk). The default route through ethernet0/3 is active, while the default route through ethernet0/4 is not active since it is less preferred.

Figure 23: Get Route Output

```

device-> get route
untrust-vr (0 entries)
-----
C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (4 entries)
-----
  ID  IP-Prefix  InterfaceGateway  P  Pref  Mtr  Vsys
-----
*  4  0.0.0.0/0  eth0/3  1.1.1.250  S  20  1  Root
*  2  1.1.1.0/24  eth0/3  0.0.0.0  C  0  0  Root
  3  0.0.0.0/0  eth0/4  2.2.2.250  S  20  10  Root
*  6  2.2.2.0/24  eth0/4  0.0.0.0  C  0  1  Root
*  5  10.1.1.0/24  eth0/1  0.0.0.0  C  20  1  Root

```

If the route through ethernet0/3 becomes unavailable, the default route through ethernet0/4 becomes active. You enable and configure IP tracking on the ethernet0/3 interface to monitor the router address 1.1.1.250. If IP tracking fails to reach 1.1.1.250, all routes associated with the ethernet0/3 interface become inactive on the security device. As a result, the default route through ethernet0/4 becomes active. When IP tracking is again able to reach 1.1.1.250, the default route through ethernet0/3 becomes active and, at the same time, the default route through ethernet0/4 becomes inactive, because it is less preferred than the default route through ethernet0/3.

The following enables IP tracking with an interface failure threshold of 5 and configures IP tracking on the ethernet0/3 interface to monitor the router IP address 1.1.1.250, which is assigned a weight of 10.

WebUI

Network > Interfaces > Edit (for ethernet0/3) > Monitor: Enter the following, then click **Apply**:

```

Enable Track IP: (select)
Threshold: 5
> Monitor Track IP ADD: Enter the following, then click Add:
  Static: (select)
  Track IP: 1.1.1.250
  Weight: 10

```

CLI

```

set interface ethernet0/3 monitor track-ip ip 1.1.1.250 weight 10
set interface ethernet0/3 monitor track-ip threshold 5
set interface ethernet0/3 monitor track-ip
save

```

In the example, the failure threshold for the target address is set to the default value of 3. That is, if the target does not return a response to three consecutive pings, a weight of 10 is applied toward the failure threshold for IP tracking on the interface. Because the failure threshold for IP tracking on the interface is 5, a weight of 10 causes routes associated with the interface to be deactivated on the security device.

You can verify the status of the IP tracking on the interface by issuing the CLI command **get interface ethernet0/3 track-ip**, as shown in Figure 24.

Figure 24: Get Interface Output

```
device-> get interface ethernet0/3 track-ip
ip address interval threshold wei gateway fail-count success-rate
1.1.1.250 1 1 10 0.0.0.0 343 46%
threshold: 5, failed: 1 ip(s) failed, weighted sum = 10
```

The **get route** command shows that the default route through ethernet0/4 is now active, while all routes through ethernet0/3 are no longer active.

Figure 25: Get Route Output With Activated Interfaces

```
device-> get route
untrust-vr (0 entries)
-----
C - Connected, S - Static, A - Auto-Exported, I - Imported, R - RIP
iB - IBGP, eB - EBGP, O - OSPF, E1 - OSPF external type 1
E2 - OSPF external type 2
trust-vr (4 entries)
-----
```

ID	IP-Prefix	Interface	Gateway	P	Pref	Mtr	Vsys
4	0.0.0.0/0	eth0/3	1.1.1.250	S	20	1	Root
2	1.1.1.0/24	eth0/3	0.0.0.0	C	0	0	Root
* 3	0.0.0.0/0	eth0/4	2.2.2.250	S	20	10	Root
* 6	2.2.2.0/24	eth0/4	0.0.0.0	C	0	1	Root
* 5	10.1.1.0/24	eth0/1	0.0.0.0	C	20	1	Root

Note that even though the routes through ethernet0/3 are no longer active, IP tracking uses the routes associated with ethernet0/3 to continue sending ping requests to the target IP address. When IP tracking is again able to reach 1.1.1.250, the default route through ethernet0/3 again becomes active on the security device. At the same time, the default route through ethernet0/4 becomes inactive, since it is less preferred than the default route through ethernet0/3.

Interface Monitoring

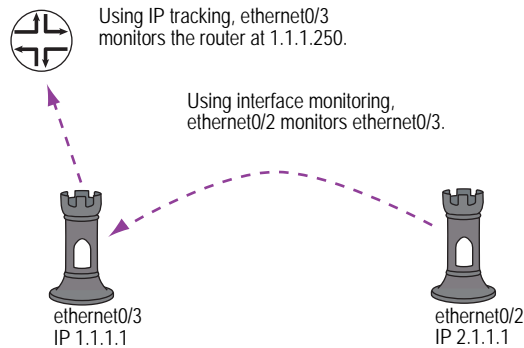
A security device can monitor the physical and logical state of interfaces and then take action based on observed changes, see Figure 26. For example, if the state of a monitored interface changes from up to down, the following can occur, as shown in Table 6.

Table 6: Monitored Interface

If:	Then:
The physical state of an interface changes from up to down	The state change might trigger another interface that is monitoring the one that just went down to also go down. You can specify whether you want the second interface to be physically or logically down. The state change of either interface going physically down, or the combined weight of both going physically down together, might trigger an NSRP failover. An NSRP device or a VSD group failover can only occur as a result of a change to the physical state of an interface.
The logical state of an interface changes from up to down as the result of an IP tracking failure	The state change might trigger another interface that is monitoring the one that just went down to also go down. Although the first interface is down logically, you can specify whether you want the down state of the second interface to be logical or physical.

Figure 26: Ethernet0/2 and Ethernet0/3 Interface Monitoring

One Interface Monitoring Another Interface



To set interface monitoring, do either of the following:

WebUI

Network > Interfaces > Edit (for the interface you want to do the monitoring)
> Monitor > Edit Interface: Enter the following, then click **Apply**:

Interface Name: Select the interface that you want to be monitored.
Weight: Enter a weight between 1 and 255.

CLI

```
set interface interface1 monitor interface interface2 [ weight number ]
save
```

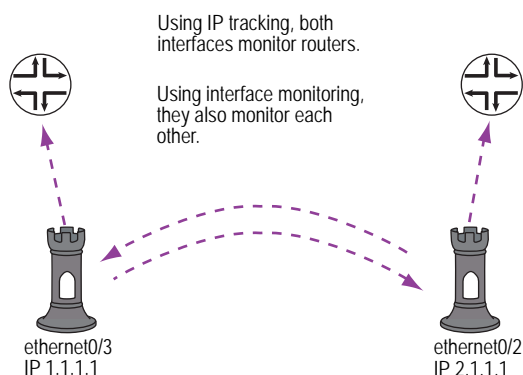
If you do not set a weight, the security device applies the default value, 255.

If two interfaces monitor each other, they form a loop. In that case, if either interface changes state, the other interface in the loop also changes state. See Figure 27 on page 69.

NOTE: An interface can only be in one loop at a time. We do not support configurations in which one interface belongs to multiple loops.

Figure 27: Loop Monitoring

Loop - Two Interfaces Monitoring Each Other



First State Change If:

- The number of unsuccessful ping attempts to either router exceeds the failure threshold for that tracked IP address.
- The weight of the failed track IP is greater than or equal to the track object failure threshold.
- The track object weight is greater than or equal to the monitor failure threshold.
- The failure action is a change from up to down.

Then:

The second interface also changes its state from up to down.

Second State Change If:

- The weight of the failure of the first interface is greater than or equal to the monitor failure threshold of the second interface.
- The failure action is a change from up to down.

Then:

The second interface also changes its state from up to down.


Monitoring Two Interfaces

In this example, you configure ethernet0/3 to monitor two interfaces—ethernet0/1 and ethernet0/2. Because the weight for each monitored interface (8 + 8) equals the monitor failure threshold (16), both ethernet0/1 and ethernet0/2 must fail (and change their state from up to down) concurrently to cause ethernet0/3 to fail (and change its state from up to down). See Figure 28.

NOTE: This example omits the configuration of IP tracking on the ethernet0/1 and ethernet0/2 interfaces (see “Tracking IP Addresses” on page 63). Without IP tracking, the only way that ethernet0/1 and ethernet0/2 might fail is if they become physically disconnected from other network devices or if they cannot maintain links with those devices.

If you set the monitor failure threshold to 8—or leave it at 16 and set the weight of each monitored interface to 16—the failure of either ethernet0/1 or ethernet0/2 can cause ethernet0/3 to fail.

Figure 28: Two-Loop Interface Monitoring



WebUI

Network > Interfaces > Edit (for ethernet0/3) > Monitor > Edit Interface:
Enter the following, then click **Apply**:

ethernet0/1: (select); Weight: 8
ethernet0/2: (select); Weight: 8

Network > Interfaces > Edit (for ethernet0/3) > Monitor: Enter **16** in the Monitor Threshold field, then click **Apply**.

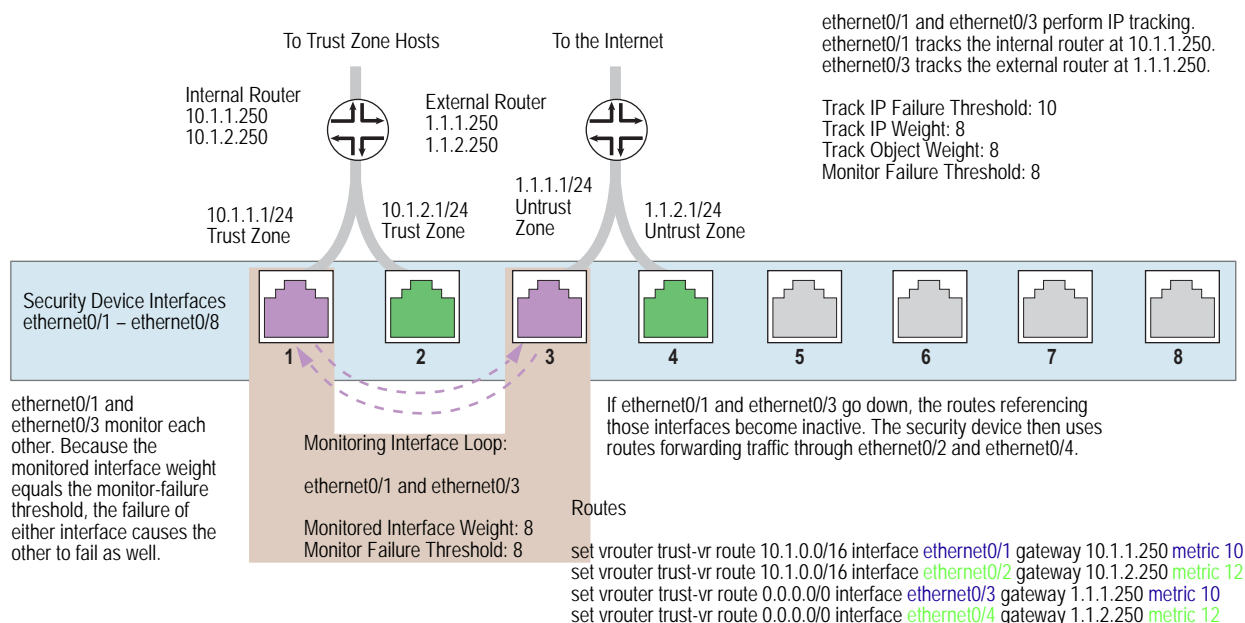
CLI

```
set interface ethernet0/3 monitor interface ethernet0/1 weight 8
set interface ethernet0/3 monitor interface ethernet0/2 weight 8
set interface ethernet0/3 monitor threshold 16
save
```

Monitoring an Interface Loop

In this example, you first configure IP tracking for two interfaces—ethernet0/1 and ethernet0/3. Then you configure these interfaces to monitor each other so that if one changes its state, the other does likewise. Finally, you define two sets of routes. The first set forwards traffic through ethernet0/1 and ethernet0/3. The second set has the same destination addresses, but these routes have lower ranked metrics and use different egress interfaces (ethernet0/2 and ethernet0/4) and different gateways from the first set. With this configuration, if the first set of interfaces fails, the security device can reroute all traffic through the second set. All zones are in the trust-vr routing domain.

Figure 29: Four-Interface Loop Monitoring



WebUI

1. IP Tracking

Network > Interfaces > Edit (for ethernet0/1) > Monitor: Enter the following, then click **Apply**.

Enable Track IP: (select)
Monitor Threshold: 8
Track IP Option: Threshold: 8
Weight: 8

NOTE: To control whether the state of an interface becomes logically or physically down (or up), you must use the CLI command **set interface interface monitor threshold number action { down | up } { logically | physically }**. Only physical interfaces bound to any security zone other than the Null zone can be physically up or down.

> Monitor Track IP ADD: Enter the following, then click **Add**:

Static: (select)
Track IP: 10.1.1.250
Weight: 8
Interval: 3 Seconds
Threshold: 10

Network > Interfaces > Edit (for ethernet0/3) > Monitor: Enter the following, then click **Apply**.

Enable Track IP: (select)
Monitor Threshold: 8
Track IP Option: Threshold: 8
Weight: 8

> Monitor Track IP ADD: Enter the following, then click **Add**:

Static: (select)
Track IP: 1.1.1.250
Weight: 8
Interval: 3 Seconds
Threshold: 10

2. Interface Monitoring

Network > Interfaces > Edit (for ethernet0/1) > Monitor > Edit Interface:
Enter the following, then click **Apply**:

ethernet0/3: (select); Weight: 8

Network > Interfaces > Edit (for ethernet0/3) > Monitor > Edit Interface:
Enter the following, then click **Apply**:

ethernet0/1: (select); Weight: 8

3. Routes

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 10.1.0.0/16
Gateway: (select)
Interface: ethernet0/1
Gateway IP Address: 10.1.1.250
Metric: 10

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 10.1.0.0/16
Gateway: (select)
Interface: ethernet0/2
Gateway IP Address: 10.1.2.250
Metric: 12

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet0/3
Gateway IP Address: 1.1.1.250
Metric: 10

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet0/4
Gateway IP Address: 1.1.2.250
Metric: 12

CLI**1. IP Tracking**

```

set interface ethernet0/1 track-ip ip 10.1.1.250 weight 8
set interface ethernet0/1 track-ip threshold 8
set interface ethernet0/1 track-ip weight 8
set interface ethernet0/1 track-ip
set interface ethernet0/3 track-ip ip 1.1.1.250 weight 8
set interface ethernet0/3 track-ip threshold 8
set interface ethernet0/3 track-ip weight 8
set interface ethernet0/3 track-ip

```

2. Interface Monitoring

```

set interface ethernet0/1 monitor interface ethernet0/3 weight 8
set interface ethernet0/1 monitor threshold 8 action down physically
set interface ethernet0/3 monitor interface ethernet0/1 weight 8
set interface ethernet0/3 monitor threshold 8 action down physically

```

3. Routes

```

set vrouter trust-vr route 10.1.0.0/16 interface ethernet0/1 gateway 10.1.1.250
metric 10
set vrouter trust-vr route 10.1.0.0/16 interface ethernet0/2 gateway 10.1.2.250
metric 12
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 1.1.1.250
metric 10
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/4 gateway 1.1.2.250
metric 12
save

```

Security Zone Monitoring

In addition to monitoring individual interfaces, an interface can monitor all the interfaces in a security zone—any security zone other than its own. For an entire security zone to fail, every interface bound to that zone must fail. As long as one interface bound to a monitored zone is up, the security device considers the entire zone to be up.

To configure an interface to monitor a security zone, do either of the following:

WebUI

Network > Interfaces > Edit (for the interface you want to do the monitoring)
> Monitor > Edit Zone: Enter the following, then click **Apply**:

Zone Name: Select the zone that you want to be monitored.
Weight: Enter a weight between 1 and 255.

CLI

```
set interface interface monitor zone zone [ weight number ]
```

If you do not set a weight, the security device applies the default value, 255.

Down Interfaces and Traffic Flow

Configuring IP tracking on an interface allows the security device to reroute outgoing traffic through a different interface if certain IP addresses become unreachable through the first interface. However, while the security device might deactivate routes associated with an interface because of an IP tracking failure, the interface can remain physically active and still send and receive traffic. For example, the security device continues to process incoming traffic for an existing session that might arrive on the original interface on which IP tracking failed. Also, the security device continues to use the interface to send ping requests to target IP addresses to determine if the targets again become reachable. In these situations, traffic still passes through an interface on which IP tracking has failed and for which the security device has deactivated routes.

How the security device handles session traffic on such an interface depends upon the following:

- If the interface on which you configure IP tracking functions as an egress interface for a session, session replies might continue to arrive at the interface and the security device still processes them.
- If the interface on which you configure IP tracking functions as an ingress interface for a session, applying the **set arp always-on-dest** command causes the security device to reroute session replies to another interface. If you do not set this command, the security device forwards session replies through the interface on which IP tracking failed even though the security device has deactivated routes using that interface. (By default, this command is unset.)

By default, a security device caches a session initiator's MAC address when it receives the initial packet for a new session. If you enter the CLI command **set arp always-on-dest**, the security device does not cache a session initiator's MAC address. Instead, the security device performs an ARP lookup when processing the reply to that initial packet. If the initiator's MAC address is in the ARP table, the security device uses that. If the MAC address is not in the ARP table, the security device sends an ARP request for the destination MAC address and then adds the MAC address it receives to its ARP table. The security device performs another ARP lookup whenever a route change occurs.

"Failure on the Egress Interface" describes separate scenarios in which IP tracking fails on the egress interface and on the ingress interface; and, in the case of the latter, what occurs when you use the command **set arp always-on-dest**.

NOTE: "Failure on the Egress Interface" describes how IP tracking triggers routing changes and how those changes can affect the packet flow through all Juniper Networks security devices.


Failure on the Egress Interface

In the following scenario, you configure IP tracking on ethernet0/2, which is the egress interface for sessions from Host A to Host B. Host A initiates the session by sending a packet to Host B, as shown in Figure 30.

NOTE: You must first create two routes to Host B, and both the egress interfaces must be in the same zone so that the same policy applies to traffic before and after the rerouting occurs.

Figure 30: Host A and Host B IP Tracking


Traffic Flow from Host A to Host B – Request (Session Initiation)



When Host B replies to Host A, the return traffic follows a similar path back through the security device, as shown in Figure 31.

Figure 31: Host B to Host A Egress Traffic Flow


Traffic Flow from Host A to Host B – Reply



If IP tracking on ethernet0/2 fails, the security device deactivates routes that use ethernet0/2 and uses ethernet0/3 for outbound traffic to Host B. However, replies from Host B to Host A can arrive through either ethernet0/2 or ethernet0/3 and the security device forwards them through ethernet0/1 to Host A. See Figure 32.

Figure 32: Egress IP Tracking Failure

Traffic Flow from Host B to Host A – IP Tracking Failure Triggers Rerouting




Failure on the Ingress Interface

In the following scenario, you again configure IP tracking on ethernet0/2, but this time ethernet0/2 is the ingress interface on the security device for sessions from Host B to Host A. Host B initiates the session by sending a packet to Host A, as shown in Figure 33.

Figure 33: Host B to Host A Ingress Traffic Flow


Traffic Flow from Host B to Host A – Request (Session Initiation)



When Host A replies to Host B, the return traffic follows a similar path back through the security device, as shown in Figure 34.

Figure 34: Ingress Host A to Host B Traffic Flow

Traffic Flow from Host B to Host A – Reply



1. Host A at 10.1.1.5 sends a reply packet destined for Host B (2.2.2.2) to ethernet0/1 at 10.1.1.1.
2. The security device performs a session lookup. Because this is a reply, the device matches it with an existing session and refreshes the session table entry.
3. By using the cached MAC address for the gateway at 1.1.1.254, or by doing an ARP lookup to discover its MAC address, the device forwards the packet through ethernet0/2 to the gateway.
4. When the gateway at 1.1.1.254 receives the reply, it forwards it to its next hop. Routing continues until Host B receives it.


If IP tracking on ethernet0/2 fails, the security device deactivates routes that use ethernet0/2 and uses ethernet0/3 for outbound traffic to Host B. However, requests from Host B to Host A can still arrive through ethernet0/2 and the security device still forwards them to Host A through ethernet0/1. The data flow for requests from Host B to Host A looks the same after an IP tracking failure as it did before. However, the replies from Host A can take one of two different paths, depending on the application of the **set arp always-on-dest** command.

If you set the command **set arp always-on-dest**, the security device sends an ARP request for the destination MAC address when processing the reply to the first packet in a session or when a route change occurs. (When this command is unset, the security device caches the session initiator's MAC address and uses that when processing replies. By default, this command is unset).

When IP tracking on ethernet0/2 fails, the security device first deactivates all routes using ethernet0/2 and then does a route lookup. It finds another route to reach Host B through ethernet0/3 and the gateway at 1.1.2.254. It then scans its session table and redirects all sessions to the new route. If you have the **set arp always-on-dest** command enabled, the security device does an ARP lookup when it receives the next packet from Host A because it is in a session affected by the route change. Despite the ingress interface on which packets from Host B arrive, the security device sends all further replies from Host A through ethernet0/3 to the gateway at 1.1.2.254. See Figure 35 on page 78.

Figure 35: Ingress IP Tracking Failure with Traffic Rerouting


Traffic Flow from Host B to Host A – IP Tracking Failure Triggers Rerouting



If you have set the command **unset arp always-on-dest** (which is the default configuration), the security device uses the MAC address for the gateway at 1.1.1.1 that it cached when Host B sent the initial session packet. The security device continues to send session replies through ethernet0/2. In this case, the IP tracking failure caused no change in the flow of data through the security device.

Figure 36: Ingress IP Tracking Failure with No Rerouting

Traffic Flow from Host B to Host A – IP Tracking Failure Triggers No Rerouting



Chapter 4

Interface Modes

Interfaces can operate in three different modes: Network Address Translation (NAT), Route, and Transparent. If an interface bound to a Layer 3 zone has an IP address, you can define the operational mode for that interface as either NAT or Route. An interface bound to a Layer 2 zone (such as the predefined v1-trust, v1-untrust, and v1-dmz zones, or a user-defined Layer 2 zone) must be in Transparent mode. You select an operational mode when you configure an interface.

NOTE: Although you can define the operational mode for an interface bound to any Layer 3 zone as NAT, the security device only performs NAT on traffic passing through that interface en route to the Untrust zone. ScreenOS does not perform NAT on traffic destined for any zone other than the Untrust zone. Also, note that ScreenOS allows you to set an Untrust zone interface in NAT mode, but doing so activates no NAT operations.


This chapter contains the following sections:

- “Transparent Mode” on page 80
- “NAT Mode” on page 92
- “Route Mode” on page 98

Transparent Mode

When an interface is in Transparent mode, the security device filters packets traversing the firewall without modifying any of the source or destination information in the IP packet header. All interfaces behave as though they are part of the same network, with the security device acting much like a Layer 2 switch or bridge. In Transparent mode, the IP addresses of interfaces are set at 0.0.0.0, making the presence of the security device transparent (invisible) to users. See Figure 37.

Figure 37: Transparent Mode



Transparent mode is a convenient means for protecting web servers or any other kind of server that mainly receives traffic from untrusted sources. Using Transparent mode offers the following benefits:

- No need to reconfigure the IP settings of routers or protected servers
- No need to create mapped IP (MIP) or virtual IP (VIP) addresses for incoming traffic to reach protected servers

Zone Settings

By default, ScreenOS creates one function zone, the VLAN zone, and three L2 security zones: V1-Trust, V1-Untrust, and V1-DMZ.

VLAN Zone

The VLAN zone hosts the VLAN1 interface, which has the same configuration and management abilities as a physical interface. When the security device is in Transparent mode, you use the VLAN1 interface for managing the device and terminating VPN traffic. You can configure the VLAN1 interface to permit hosts in the L2 security zones to manage the device. To do that, you must set the VLAN1 interface IP address in the same subnet as the hosts in the L2 security zones.

For management traffic, the VLAN1 Manage IP takes precedence over the VLAN1 interface IP. You can set the VLAN1 Manage IP for management traffic and dedicate the VLAN1 interface IP solely for VPN tunnel termination.

Predefined Layer 2 Zones

ScreenOS provides three L2 security zones by default: V1-Trust, V1-Untrust, and V1-DMZ. These three zones share the same L2 domain. When you configure an interface in one of the zones, it gets added to the L2 domain shared by all interfaces in all the L2 zones. All hosts in the L2 zones must be on the same subnet to communicate.

As stated in the previous section, when the device is in transparent mode, you use the VLAN1 interface to manage the device. For management traffic to reach the VLAN1 interface, you must enable the management options on the VLAN1 interface and on the zone(s) through which the management traffic passes. By default, all management options are enabled in the V1-Trust zone. To enable hosts in other zones to manage the device, you must set those options on the zones to which they belong.

NOTE: To see which physical interfaces are prebound to the L2 zones for each Juniper Networks security device, refer to the installation and configuration guide for that device.

Traffic Forwarding

A security device operating at Layer 2 (L2) does not permit any inter-zone traffic unless there is a policy configured on the device. For more information about setting policies, see “Policies” on page 2-159. After you configure a policy on the security device, it does the following:

- Allows or denies the traffic specified in the policy.
- Allows ARP and L2 non-IP multicast and broadcast traffic. The security device can then receive and pass L2 broadcast traffic for the spanning tree protocol.
- Continues to block all non-IP and non-ARP unicast traffic and IPSec traffic.

You can change the forwarding behavior of the device as follows:

- To block all L2 non-IP and non-ARP traffic, including multicast and broadcast traffic, enter the **unset interface vlan1 bypass-non-ip-all** command.
- To allow all L2 non-IP traffic to pass through the device, enter the **set interface vlan1 bypass-non-ip** command.
- To revert to the default behavior of the device, which is to block all non-IP and non-ARP unicast traffic, enter the **unset interface vlan1 bypass-non-ip** command.

Note that the **unset interface vlan1 bypass-non-ip-all** command always overwrites the **unset interface vlan1 bypass-non-ip** command when both commands are in the configuration file. Therefore, if you had previously entered the **unset interface vlan1 bypass-non-ip-all** command, and you now want the device to revert to its default behavior of blocking only the non-IP and non-ARP unicast traffic, you should first enter the **set interface vlan1 bypass-non-ip** command to allow all non-IP and non-ARP traffic, including multicast, unicast, and broadcast traffic to pass through the device. Then you must enter the **unset interface vlan1 bypass-non-ip** command to block only the non-IP, non-ARP unicast traffic.

- To allow a security device to pass IPSec traffic without attempting to terminate it, use the **set interface vlan1 bypass-others-ipsec** command. The security device then allows the IPSec traffic to pass through to other VPN termination points.

NOTE: A security device with interfaces in Transparent mode requires routes for two purposes: to direct self-initiated traffic, such as SNMP traps, and to forward VPN traffic after encapsulating or decapsulating it.

Unknown Unicast Options

When a host or any kind of network device does not know the MAC address associated with the IP address of another device, it uses the Address Resolution Protocol (ARP) to obtain it. The requestor broadcasts an ARP query (arp-q) to all the other devices on the same subnet. The arp-q requests the device at the specified destination IP address to send back an ARP reply (arp-r), which provides the requestor with the MAC address of the replier. When all the other devices on the subnet receive the arp-q, they check the destination IP address and, because it is not their IP address, drop the packet. Only the device with the specified IP address returns an arp-r. After a device matches an IP address with a MAC address, it stores the information in its ARP cache.

As ARP traffic passes through a security device in Transparent mode, the device notes the source MAC address in each packet and learns which interface leads to that MAC address. In fact, the security device learns which interface leads to which MAC address by noting the source MAC addresses in all packets it receives. It then stores this information in its forwarding table.

NOTE: A security device in Transparent mode does not permit any traffic between zones unless there is a policy configured on the device. For more information about how the device forwards traffic when it is in Transparent mode, see “Traffic Forwarding” on page 81.

The situation can arise when a device sends a unicast packet with a destination MAC address, which it has in its ARP cache, but which the security device does not have in its forwarding table. For example, the security device clears its forwarding table every time it reboots. (You can also clear the forwarding table with the CLI command **clear arp**.) When a security device in Transparent mode receives a unicast packet for which it has no entry in its forwarding table, it can follow one of two courses:

- After doing a policy lookup to determine the zones to which traffic from the source address is permitted, flood the initial packet out the interfaces bound to those zones, and then continue using whichever interface receives a reply. This is the Flood option, which is enabled by default.
- Drop the initial packet, flood ARP queries (and, optionally, trace-route packets, which are ICMP echo requests with the time-to-live value set to 1) out all interfaces (except the interface at which the packet arrived), and then send subsequent packets through whichever interface receives an ARP (or trace-route) reply from the router or host whose MAC address matches the destination MAC address in the initial packet. The trace-route option allows the security device to discover the destination MAC address when the destination IP address is in a nonadjacent subnet.

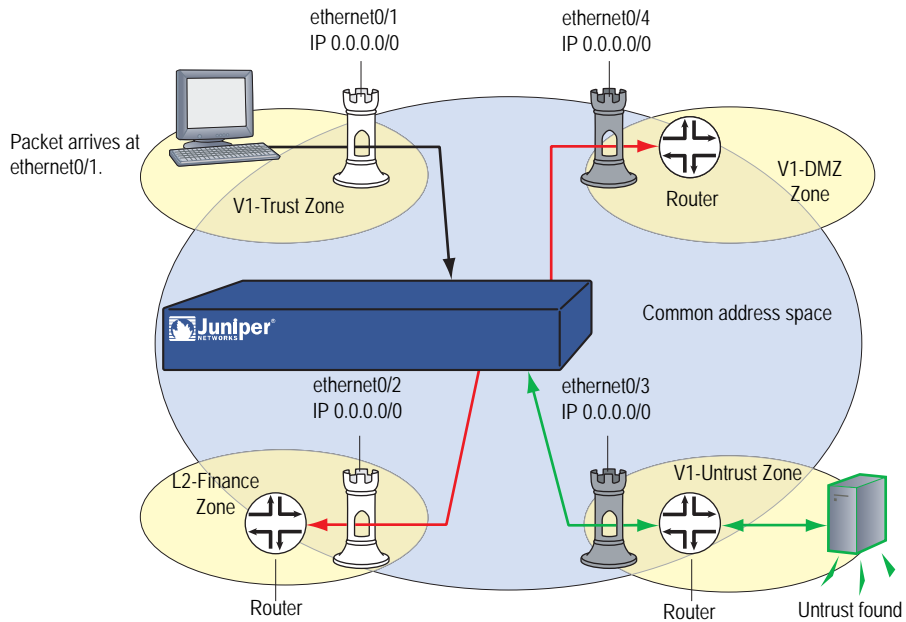
NOTE: Of the two methods—flood and ARP/trace-route—ARP/trace-route is more secure because the security device floods ARP queries and trace-route packets—not the initial packet—out all interfaces.

Flood Method

The flood method forwards packets in the same manner as most Layer 2 switches. A switch maintains a forwarding table that contains MAC addresses and associated ports for each Layer 2 domain. The table also contains the corresponding interface through which the switch can forward traffic to each device. Every time a packet arrives with a new source MAC address in its frame header, the switch adds the MAC address to its forwarding table. It also tracks the interface at which the packet arrived. If the destination MAC address is unknown to the switch, the switch duplicates the packet and floods it out all interfaces (other than the interface at which the packet arrived). It learns the previously unknown MAC address and its corresponding interface when a reply with that MAC address arrives at one of its interfaces.

When you enable the flood method and the security device receives an Ethernet frame with a destination MAC address that is not listed in the security device MAC table, it floods the packet out all interfaces.

Figure 38: Flood Method



The security device floods the packet out ethernet0/2 but receives no reply.

To enable the flood method for handling unknown unicast packets, do either of the following:

WebUI

Network > Interfaces > Edit (for VLAN1): For the broadcast options, select **Flood**, then click **OK**.

CLI

```
set interface vlan1 broadcast flood
save
```

ARP/Trace-Route Method

When you enable the ARP method with the trace-route option and the security device receives an Ethernet frame with a destination MAC address that is not listed in its MAC table, the security device performs the following series of actions:

NOTE: When you enable the ARP method, the trace-route option is enabled by default. You can also enable the ARP method without the trace-route option. However, this method only allows the security device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnet as the ingress IP address. (For more information about the ingress IP address, see the following Note.)

1. The security device notes the destination MAC address in the initial packet (and, if it is not already there, adds the source MAC address and its corresponding interface to its forwarding table).
2. The security device drops the initial packet.

3. The security device generates two packets—ARP query (arp-q) and a trace-route (an ICMP echo request, or PING) with a time-to-live (TTL) field of 1—and floods those packets out all interfaces except the interface at which the initial packet arrived. For the arp-q packets and ICMP echo requests, the security device uses the source and destination IP addresses from the initial packet. For arp-q packets, the security device replaces the source MAC address from the initial packet with the MAC address for VLAN1, and it replaces the destination MAC address from the initial packet with ffff.ffff.ffff. For the trace-route option, the security device uses the source and destination MAC addresses from the initial packet in the ICMP echo requests that it broadcasts.

If the destination IP address belongs to a device in the same subnet as the ingress IP address, the host returns an ARP reply (arp-r) with its MAC address, thus indicating the interface through which the security device must forward traffic destined for that address. (See Figure 39, “ARP Method,” on page 86.)

NOTE: The ingress IP address refers to the IP address of the last device to send the packet to the security device. This device might be the source that sent the packet or a router forwarding the packet.

If the destination IP address belongs to a device in a subnet beyond that of the ingress IP address, the trace-route returns the IP and MAC addresses of the router leading to the destination, and more significantly, indicates the interface through which the security device must forward traffic destined for that MAC address. (See Figure 40 on page 87.)

NOTE: Actually, the trace-route returns the IP and MAC addresses of all the routers in the subnet. The security device then matches the destination MAC address from the initial packet with the source MAC address on the arp-r packets to determine which router to target, and consequently, which interface to use to reach that target.

4. Combining the destination MAC address gleaned from the initial packet with the interface leading to that MAC address, the security device adds a new entry to its forwarding table.
5. The security device forwards all subsequent packets it receives out the correct interface to the destination.

To enable the ARP/trace-route method for handling unknown unicast packets, do either of the following:

WebUI

Network > Interfaces > Edit (for VLAN1): For the broadcast options, select **ARP**, then click **OK**.

CLI

```
set interface vlan1 broadcast arp
save
```

NOTE: The trace-route option is enabled by default. If you want to use ARP without the trace-route option, enter the following command: **unset interface vlan1 broadcast arp trace-route**. This command unsets the trace-route option but does not unset ARP as the method for handling unknown unicast packets.

Figure 39 shows how the ARP method can locate the destination MAC when the destination IP address is in an adjacent subnet.

Figure 39: ARP Method

Note: Only the relevant elements of the packet header and the last four digits in the MAC addresses are shown below.

If the following packet:

Ethernet Frame			IP Datagram	
dst	src	type	src	dst
11bb	11aa	0800	210.1.1.5	210.1.1.75

arrives at ethernet0/1 and the forwarding table does not have an entry for MAC address 00bb.11bb.11bb, the security device floods the following arp-q packet:

Ethernet Frame			ARM Message	
dst	src	type	src	dst
ffff	39ce	0806	210.1.1.5	210.1.1.75

When the device receives the following arp-r at ethernet0/2:

Ethernet Frame			ARM Message	
dst	src	type	src	dst
39ce	11bb	0806	210.1.1.75	210.1.1.5

It can now associate the MAC address with the interface leading to it.

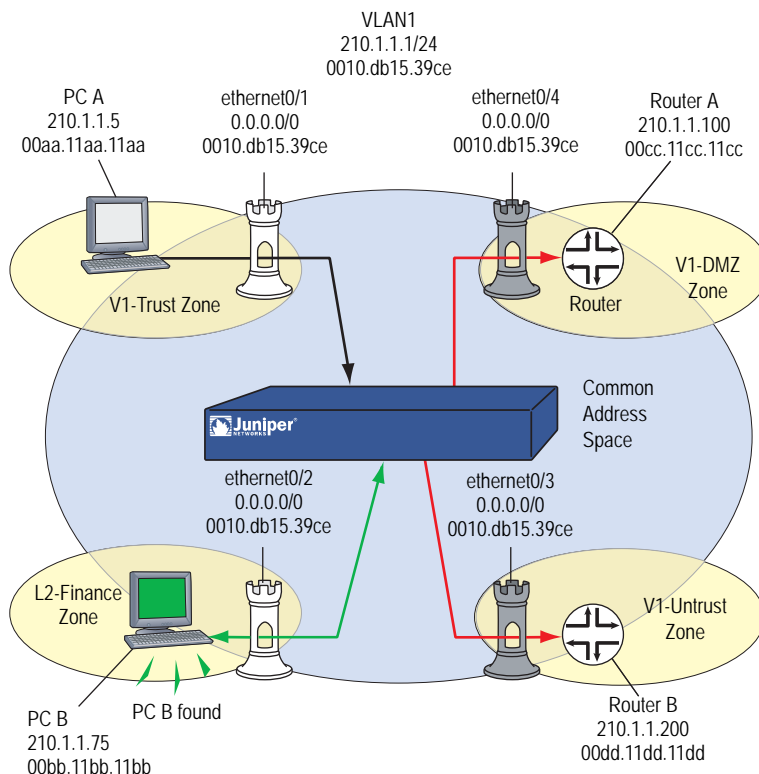


Figure 40 shows how the trace-route option can locate the destination MAC when the destination IP address is in a nonadjacent subnet.

Figure 40: Trace-Route

Note: Only the relevant elements of the packet header and the last four digits in the MAC addresses are shown below.

If the following packet:

Ethernet Frame			IP Datagram	
dst	src	type	src	dst
11dd	11aa	0800	210.1.1.5	195.1.1.5

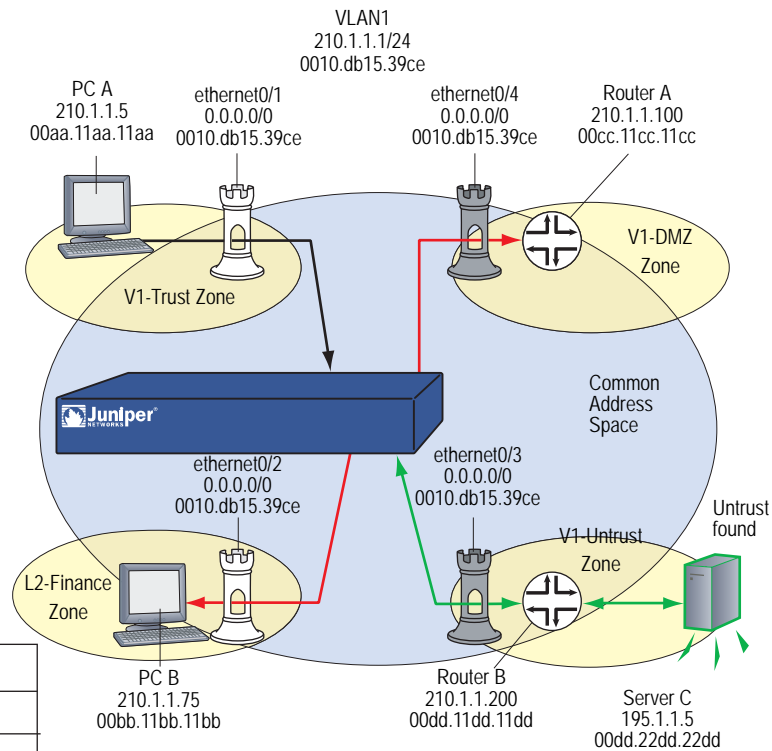
arrives at ethernet0/1 and the forwarding table does not have an entry for MAC address 00dd.11dd.11dd, the security device floods the following trace-route packet out ethernet0/2, ethernet0/3, and ethernet0/4:

Ethernet Frame			ICMP Message		
dst	src	type	src	dst	TTL
11dd	11aa	0800	210.1.1.5	195.1.1.5	1

When the device receives the following response at ethernet0/3:

Ethernet Frame			ICMP Message		
dst	src	type	src	dst	msg
11aa	11dd	0800	210.1.1.200	210.1.1.5	Time Exceeded

It can now associate the MAC address with the interface leading to it.



Configuring VLAN1 Interface for Management

In this example, you configure the security device for management to its VLAN1 interface as follows:

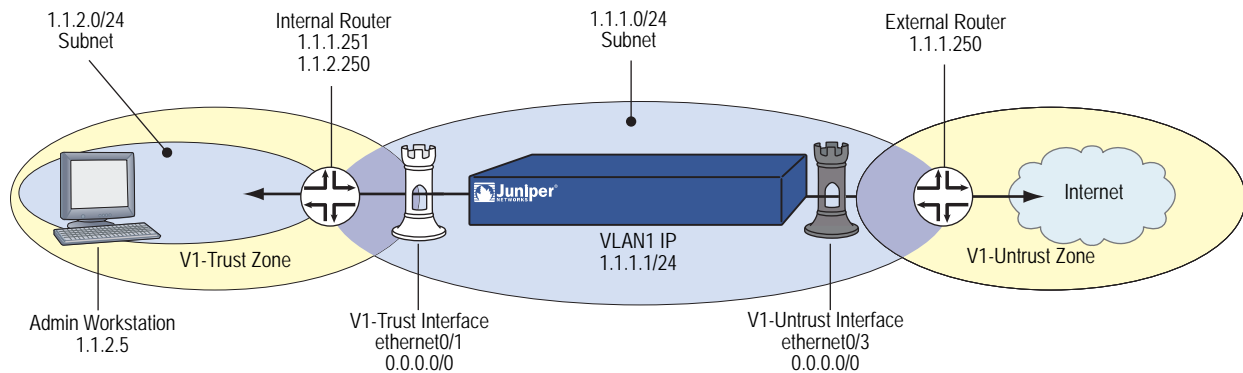
- Assign the VLAN1 interface an IP address of 1.1.1.1/24.
- Enable Web, Telnet, SSH, and Ping on both the VLAN1 interface and the V1-Trust security zone.

NOTE: By default, ScreenOS enables the management options for the VLAN1 interface and V1-Trust security zone. Enabling these options is included in this example for illustrative purposes only. Unless you have previously disabled them, you really do not need to enable them manually.

To manage the device from a Layer 2 security zone, you must set the same management options for both the VLAN1 interface and the Layer 2 security zone.

- Add a route in the trust virtual router (all Layer 2 security zones are in the trust-vr routing domain) to enable management traffic to flow between the security device and an administrative workstation beyond the immediate subnet of the security device. All security zones are in the trust-vr routing domain.

Figure 41: Transparent VLAN



WebUI

1. VLAN1 Interface

Network > Interfaces > Edit (for VLAN1): Enter the following, then click **OK**:

IP Address/Netmask: 1.1.1.1/24
 Management Services: WebUI, Telnet, SSH (select)
 Other Services: Ping (select)

2. V1-Trust Zone

Network > Zones > Edit (for V1-Trust): Select the following, then click **OK**:

Management Services: WebUI, Telnet, SSH
 Other Services: Ping

3. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 1.1.2.0/24
 Gateway: (select)
 Interface: vlan1(trust-vr)
 Gateway IP Address: 1.1.1.251
 Metric: 1

CLI

1. VLAN1 Interface

```
set interface vlan1 ip 1.1.1.1/24
set interface vlan1 manage web
set interface vlan1 manage telnet
set interface vlan1 manage ssh
set interface vlan1 manage ping
```

2. V1-Trust Zone

```
set zone v1-trust manage web
set zone v1-trust manage telnet
set zone v1-trust manage ssh
set zone v1-trust manage ping
```

3. Route


```
set vrouter trust-vr route 1.1.2.0/24 interface vln1 gateway 1.1.1.251 metric 1
save
```

Configuring Transparent Mode

The following example illustrates a basic configuration for a single LAN protected by a security device in Transparent mode. Policies permit outgoing traffic for all hosts in the V1-Trust zone, incoming SMTP services for the mail server, and incoming FTP-GET services for the FTP server.

To increase the security of management traffic, you change the HTTP port number for WebUI management from 80 to 5555, and the Telnet port number for CLI management from 23 to 4646. You use the VLAN1 IP address—1.1.1.1/24—to manage the security device from the V1-Trust security zone. You define addresses for the FTP and mail servers. You also configure a default route to the external router at 1.1.1.250, so that the security device can send outbound VPN traffic to it. (The default gateway on all hosts in the V1-Trust zone is also 1.1.1.250.)

NOTE: For an example of configuring a VPN tunnel for a security device with interfaces in Transparent mode, see “Transparent Mode VPN” on page 5-151.

Figure 42: Basic Transparent Mode

WebUI

1. VLAN1 Interface

Network > Interfaces > Edit (for the VLAN1 interface): Enter the following, then click **OK**:

IP Address/Netmask: 1.1.1.1/24
 Management Services: WebUI, Telnet (select)
 Other Services: Ping (select)

2. HTTP Port

Configuration > Admin > Management: In the HTTP Port field, type **5555** and then click **Apply**.

NOTE: The default port number is 80. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access to the configuration. When logging in to manage the device later, enter the following in the URL field of your browser: http://1.1.1.1:5555.

3. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Zone Name: V1-Trust
 IP Address/Netmask: 0.0.0.0/0

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: V1-Untrust
 IP Address/Netmask: 0.0.0.0/0

4. V1-Trust Zone

Network > Zones > Edit (for v1-trust): Select the following, then click **OK**:

Management Services: WebUI, Telnet
 Other Services: Ping

5. Addresses

Policy > Policy Elements > Addresses > List > New: Enter the following and then click **OK**:

Address Name: FTP_Server
 IP Address/Domain Name:
 IP/Netmask: (select), 1.1.1.5/32
 Zone: V1-Trust

Policy > Policy Elements > Addresses > List > New: Enter the following and then click **OK**:

Address Name: Mail_Server
 IP Address/Domain Name:
 IP/Netmask: (select), 1.1.1.10/32
 Zone: V1-Trust

6. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: vlan1(trust-vr)
 Gateway IP Address: 1.1.1.250
 Metric: 1

7. Policies

Policy > Policies > (From: V1-Trust, To: V1-Untrust) New: Enter the following and then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit

Policy > Policies > (From: V1-Untrust, To: V1-Trust) New: Enter the following and then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Mail_Server
 Service: Mail
 Action: Permit

Policy > Policies > (From: V1-Untrust, To: V1-Trust) New: Enter the following and then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), FTP_Server
 Service: FTP-GET
 Action: Permit

CLI**1. VLAN1**

```
set interface vlan1 ip 1.1.1.1/24
set interface vlan1 manage web
set interface vlan1 manage telnet
set interface vlan1 manage ping
```

2. Telnet

```
set admin telnet port 4646
```

NOTE: The default port number for Telnet is 23. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access to the configuration. When logging in to manage the device later via Telnet, enter the following address: 1.1.1.1 4646.

3. Interfaces

```
set interface ethernet0/1 ip 0.0.0.0/0
set interface ethernet0/1 zone v1-trust
set interface ethernet0/3 ip 0.0.0.0/0
set interface ethernet0/3 zone v1-untrust
```

4. V1-Trust Zone

```
set zone v1-trust manage web
set zone v1-trust manage telnet
set zone v1-trust manage ping
```

5. Addresses

```
set address v1-trust FTP_Server 1.1.1.5/32
set address v1-trust Mail_Server 1.1.1.10/32
```

6. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface vlan1 gateway 1.1.1.250 metric 1
```


7. Policies

```
set policy from v1-trust to v1-untrust any any permit
set policy from v1-untrust to v1-trust any Mail_Server mail permit
set policy from v1-untrust to v1-trust any FTP_Server ftp-get permit
save
```

NAT Mode

When an ingress interface is in Network Address Translation (NAT) mode, the security device, acting like a Layer 3 switch (or router), translates two components in the header of an outgoing IP packet destined for the Untrust zone: its source IP address and source port number. The security device replaces the source IP address of the originating host with the IP address of the Untrust zone interface. Also, it replaces the source port number with another random port number generated by the security device.

Figure 43: NAT Topology



When the reply packet arrives at the security device, the device translates two components in the IP header of the incoming packet: the destination address and port number, which are translated back to the original numbers. The security device then forwards the packet to its destination.

NAT adds a level of security not provided in Transparent mode: The addresses of hosts sending traffic through an ingress interface in NAT mode (such as a Trust zone interface) are never exposed to hosts in the egress zone (such as the Untrust zone) unless the two zones are in the same virtual routing domain and the security device is advertising routes to peers through a dynamic routing protocol (DRP). Even then, the Trust zone addresses are only reachable if you have a policy permitting inbound traffic to them. (If you want to keep the Trust zone addresses hidden while using a DRP, then put the Untrust zone in the untrust-vr and the Trust zone in the trust-vr, and do not export routes for internal addresses in the trust-vr to the untrust-vr.)

If the security device uses static routing and just one virtual router, the internal addresses remain hidden when traffic is outbound, due to interface-based NAT. The policies you configure control inbound traffic. If you use only mapped IP (MIP) and virtual IP (VIP) addresses as the destinations in your inbound policies, the internal addresses still remain hidden.

Also, NAT preserves the use of public IP addresses. In many environments, resources are not available to provide public IP addresses for all devices on the network. NAT services allow many private IP addresses to have access to Internet resources through one or a few public IP addresses. The following IP address ranges are reserved for private IP networks and must not get routed on the Internet:

- 10.0.0.0 – 10.255.255.255
- 172.16.0.0 – 172.31.255.255
- 192.168.0.0 – 192.168.255.255

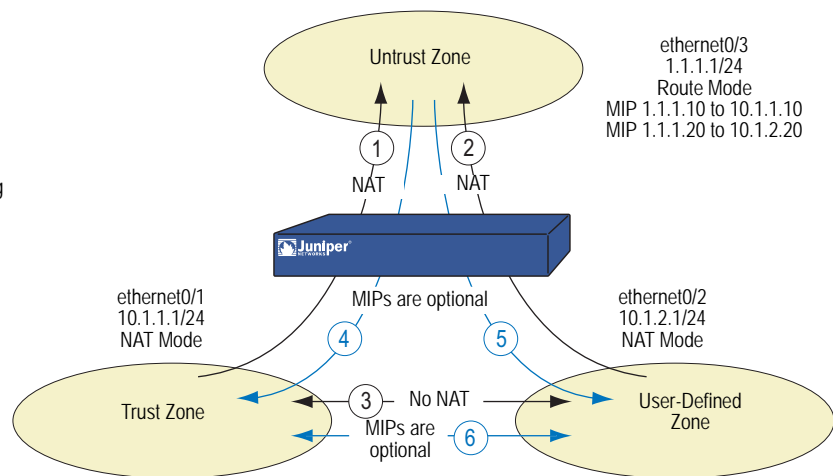
Inbound and Outbound NAT Traffic

A host in a zone sending traffic through an interface in NAT mode can initiate traffic to the Untrust zone—assuming that a policy permits it. In releases prior to ScreenOS 5.0.0, a host behind an interface in NAT mode was unable to receive traffic from the Untrust zone unless a mapped IP (MIP), virtual IP (VIP), or VPN tunnel was set up for it. However, in ScreenOS 5.0.0, traffic to a zone with a NAT-enabled interface from any zone—including the Untrust zone—does not need to use a MIP, VIP, or VPN. If you want to preserve the privacy of addresses or if you are using private addresses that do not occur on a public network such as the Internet, you can still define a MIP, VIP, or VPN for traffic to reach them. However, if issues of privacy and private IP addresses are not a concern, traffic from the Untrust zone can reach hosts behind an interface in NAT mode directly, without the use of a MIP, VIP, or VPN.

NOTE: You can define a virtual IP (VIP) address only on an interface bound to the Untrust zone.

Figure 44: NAT Traffic Flow

1. Interface-based NAT on traffic from the Trust zone to the Untrust zone.
 2. Interface-based NAT on traffic from the User-Defined zone to the Untrust zone.
- (Note: This is possible only if the User-Defined and Untrust zones are in different virtual routing domains.)
3. No interface-based NAT on traffic between the Trust and User-Defined zones.
 - 4 and 5. You can use MIPs, VIPs, or VPNs for traffic from the Untrust zone to reach the Trust zone or the User-Defined zone, but they are not required.



NOTE: For more information about MIPs, see “Mapped IP Addresses” on page 8-63. For more about VIPs, see “Virtual IP Addresses” on page 8-80.

Interface Settings

For NAT mode, define the following interface settings, where *ip_addr1* and *ip_addr2* represent numbers in an IP address, *mask* represents the numbers in a netmask, *vlan_id_num* represents the number of a VLAN tag, *zone* represents the name of a zone, and *number* represents the bandwidth size in kbps:

Zone Interfaces	Settings	Zone Subinterfaces
Trust, DMZ, and user-defined zones using NAT	IP: <i>ip_addr1</i> Netmask: <i>mask</i> Manage IP ¹ : <i>ip_addr2</i> Traffic Bandwidth ² : <i>number</i> NAT ³ : (select)	IP: <i>ip_addr1</i> Netmask: <i>mask</i> VLAN Tag: <i>vlan_id_num</i> Zone Name: <i>zone</i> NAT ¹ : (select)
Untrust ⁴	IP: <i>ip_addr1</i> Netmask: <i>mask</i> Manage IP*: <i>ip_addr2</i> Traffic Bandwidth ¹ : <i>number</i>	IP: <i>ip_addr1</i> Netmask: <i>mask</i> VLAN Tag: <i>vlan_id_num</i> Zone Name: <i>zone</i>

1. You can set the manage IP address on a per interface basis. Its primary purpose is to provide an IP address for administrative traffic separate from network traffic. You can also use the manage IP address for accessing a specific device when it is in a high availability configuration.
2. Optional setting for traffic shaping.
3. Selecting NAT defines the interface mode as NAT. Selecting Route defines the interface mode as Route.
4. Although you are able to select NAT as the interface mode on an interface bound to the Untrust zone, the security device does not perform any NAT operations on that interface.


NOTE: In NAT mode, you can manage a security device from any interface—and from multiple interfaces—using the system IP address, interface IP addresses, manage IP addresses, or the MGT IP address.

Configuring NAT Mode

The following example illustrates a simple configuration for a LAN with a single subnet in the Trust zone. The LAN is protected by a security device in NAT mode. Policies permit outgoing traffic for all hosts in the Trust zone and incoming mail for the mail server. The incoming mail is routed to the mail server through a Virtual IP address. Both the Trust and Untrust zones are in the trust-vr routing domain.

NOTE: Compare Figure 45 with that for Route mode in Figure 47, “Device in Route Mode,” on page 99.

Figure 45: Device in NAT Mode



WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24
 Enter the following, then click **OK**:
 Interface Mode: NAT

NOTE: By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24
 Interface Mode: Route

NOTE: If the IP address in the Untrust zone on the security device is dynamically assigned by an ISP, leave the IP address and netmask fields empty and select **Obtain IP using DHCP**. If the ISP uses Point-to-Point Protocol over Ethernet, select **Obtain IP using PPPoE**, click the **Create new PPPoE settings** link, and enter the name and password.

2. VIP

Network > Interfaces > Edit (for ethernet0/3) > VIP: Enter the following, then click **Add**:

Virtual IP Address: 1.1.1.5

Network > Interfaces > Edit (for ethernet0/3) > VIP > New VIP Service: Enter the following, then click **OK**:

Virtual Port: 25
 Map to Service: Mail
 Map to IP: 10.1.1.5

NOTE: For information about virtual IP (VIP) addresses, see “Virtual IP Addresses” on page 8-80.

3. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/3
 Gateway IP Address: 1.1.1.250

4. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: ANY
 Action: Permit

Policy > Policies > (From: Untrust, To: Global) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), VIP(1.1.1.5)
 Service: MAIL
 Action: Permit

CLI**1. Interfaces**

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/1 nat
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 1.1.1.1/24
set interface ethernet0/3 route
```

NOTE: The **set interface ethernet*n* nat** command determines that the security device operates in NAT mode.

If the IP address in the Untrust zone on the security device is dynamically assigned by an ISP, use the following command: **set interface untrust dhcp**. If the ISP uses Point-to-Point Protocol over Ethernet, use the **set pppoe** and **exec pppoe** commands. For more information, refer to the *ScreenOS CLI Reference Guide: IPv4 Command Descriptions*.

2. VIP

```
set interface ethernet0/3 vip 1.1.1.5 25 mail 10.1.1.5
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 1.1.1.250
```


4. Policies

```
set policy from trust to untrust any any any permit
set policy from untrust to global any vip(1.1.1.5) mail permit
save
```

Route Mode

When an interface is in Route mode, the security device routes traffic between different zones without performing source NAT (NAT-src); that is, the source address and port number in the IP packet header remain unchanged as it traverses the security device. Unlike NAT-src, you do not need to establish mapped IP (MIP) and virtual IP (VIP) addresses to allow inbound traffic to reach hosts when the destination zone interface is in Route mode. Unlike Transparent mode, the interfaces in each zone are on different subnets.

Figure 46: NAT Topology



You do not have to apply Source Network Address Translation (NAT-src) at the interface level so that all source addresses initiating outgoing traffic get translated to the IP address of the destination zone interface. Instead, you can perform NAT-src selectively at the policy level. You can determine which traffic to route and on which traffic to perform NAT-src by creating policies that enable NAT-src for specified source addresses on either incoming or outgoing traffic. For network traffic, NAT can use the IP address or addresses of the destination zone interface from a dynamic IP (DIP) pool, which is in the same subnet as the destination zone interface. For VPN traffic, NAT can use a tunnel interface IP address or an address from its associated DIP pool.

NOTE: For more information about configuring policy-based NAT-src, see “Source Network Address Translation” on page 8-13.

Interface Settings

For Route mode, define the following interface settings, where *ip_addr1* and *ip_addr2* represent numbers in an IP address, *mask* represents the numbers in a netmask, *vlan_id_num* represents the number of a VLAN tag, *zone* represents the name of a zone, and *number* represents the bandwidth size in kbps:

Table 7: Interface Settings

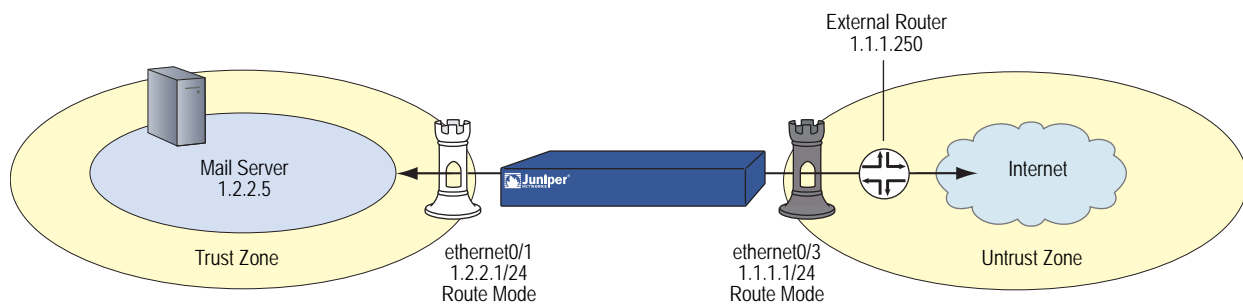
Zone Interfaces	Settings	Zone Subinterfaces
Trust, Untrust, DMZ, and user-defined zones	IP: <i>ip_addr1</i> Netmask: <i>mask</i> Manage IP ¹ : <i>ip_addr2</i> Traffic Bandwidth ² : <i>number</i> Route ³ : (select)	IP: <i>ip_addr1</i> Netmask: <i>mask</i> VLAN Tag: <i>vlan_id_num</i> Zone Name: <i>zone</i> Route [†] : (select)

1. You can set the manage IP address on a per interface basis. Its primary purpose is to provide an IP address for administrative traffic separate from network traffic. You can also use the manage IP address for accessing a specific device when it is in a high availability configuration.
2. Optional setting for traffic shaping.
3. Selecting Route defines the interface mode as Route. Selecting NAT defines the interface mode as NAT.

Configuring Route Mode

In “Configuring NAT Mode” on page 95, the hosts in the Trust zone LAN have private IP addresses and a mapped IP for the mail server. In the following example of the same network protected by a security device operating in Route mode, note that the hosts have public IP addresses and that a MIP is unnecessary for the mail server. Both security zones are in the trust-vr routing domain.

Figure 47: Device in Route Mode



WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.2.2.1/24
 Enter the following, then click **OK**:
 Interface Mode: Route

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24

NOTE: Selecting **Route** determines that the security device operates in Route mode, without performing NAT on traffic entering or exiting the Trust zone.

If the IP address in the Untrust zone on the security device is dynamically assigned by an ISP, leave the IP address and netmask fields empty and select **Obtain IP using DHCP**. If the ISP uses Point-to-Point Protocol over Ethernet, select **Obtain IP using PPPoE**, click the **Create new PPPoE settings** link, and then enter the name and password.

2. Address

Policy > Policy Elements > Addresses > List > New: Enter the following and then click **OK**:

Address Name: Mail Server
 IP Address/Domain Name:
 IP/Netmask: (select), 1.2.2.5/32
 Zone: Trust

3. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/3
 Gateway IP Address: 1.1.1.250

4. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: ANY
 Action: Permit

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Mail Server
 Service: MAIL
 Action: Permit

CLI**1. Interfaces**

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 1.2.2.1/24
set interface ethernet0/1 route
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 1.1.1.1/24
set interface ethernet0/3 route
```

NOTE: The **set interface ethernet *number* route** command determines that the security device operates in Route mode.

2. Address

```
set address trust mail_server 1.2.2.5/24
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 1.1.1.250
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy from untrust to trust any mail_server mail permit
save
```


Chapter 5

Building Blocks for Policies

This chapter discusses the components, or building blocks, that you can reference in policies. It contains the following sections:

- “Addresses” on page 103
- “Services” on page 109
- “Dynamic IP Pools” on page 140
- “Setting a Recurring Schedule” on page 156

NOTE: For information about user authentication, see *Volume 9: User Authentication*.

Addresses

ScreenOS classifies the addresses of all other devices by location and netmask. Each zone possesses its own list of addresses and address groups.

Individual hosts have only a single IP address defined and therefore, must have a netmask setting of 255.255.255.255 (which masks out all but this host).

Subnets have an IP address and a netmask (for example, 255.255.255.0 or 255.255.0.0).

Before you can configure policies to permit, deny, or tunnel traffic to and from individual hosts and subnets, you must make entries for them in ScreenOS address lists, which are organized by zones.

NOTE: You do not have to make address entries for “Any.” This term automatically applies to all devices physically located within their respective zones.

Address Entries

Before you can set up many of the Juniper Networks firewall, VPN, and traffic shaping features, you need to define addresses in one or more address lists. The address list for a security zone contains the IP addresses or domain names of hosts or subnets whose traffic is either allowed, blocked, encrypted, or user-authenticated.

NOTE: Before you can use domain names for address entries, you must configure the security device for Domain Name System (DNS) services. For information about DNS configuration, see “Domain Name System Support” on page 217.

For information regarding ScreenOS naming conventions—which apply to the names you create for addresses—see “Command Line Interface Conventions” on page x.

Adding an Address

In this example, you add the subnet “Sunnyvale_Eng” with the IP address 10.1.10.0/24 as an address in the Trust zone, and the address www.juniper.net as an address in the Untrust zone.

WebUI

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: Sunnyvale_Eng
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.10.0/24
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: Juniper
 IP Address/Domain Name:
 Domain Name: (select), www.juniper.net
 Zone: Untrust

CLI

```
set address trust Sunnyvale_Eng 10.1.10.0/24
set address untrust Juniper www.juniper.net
save
```

Modifying an Address

In this example, you change the address entry for the address “Sunnyvale_Eng” to reflect that this department is specifically for software engineering and has a different IP address—10.1.40.0/24.

WebUI

Policy > Policy Elements > Addresses > List > Edit (for Sunnyvale_Eng):
Change the name and IP address to the following, then click **OK**:

Address Name: Sunnyvale_SW_Eng
IP Address/Domain Name:
IP/Netmask: (select), 10.1.40.0/24
Zone: Trust

CLI

```
unset address trust Sunnyvale_Eng
set address trust Sunnyvale_SW_Eng 10.1.40.0/24
save
```

NOTE: After you define an address—or an address group—and associate it with a policy, you cannot change the address location to another zone (such as from Trust to Untrust). To change its location, you must first disassociate it from the underlying policy.

Deleting an Address

In this example, you remove the address entry for the address “Sunnyvale_SW_Eng”.

WebUI

Policy > Policy Elements > Addresses > List: Click **Remove** in the Configure column for Sunnyvale_SW_Eng.

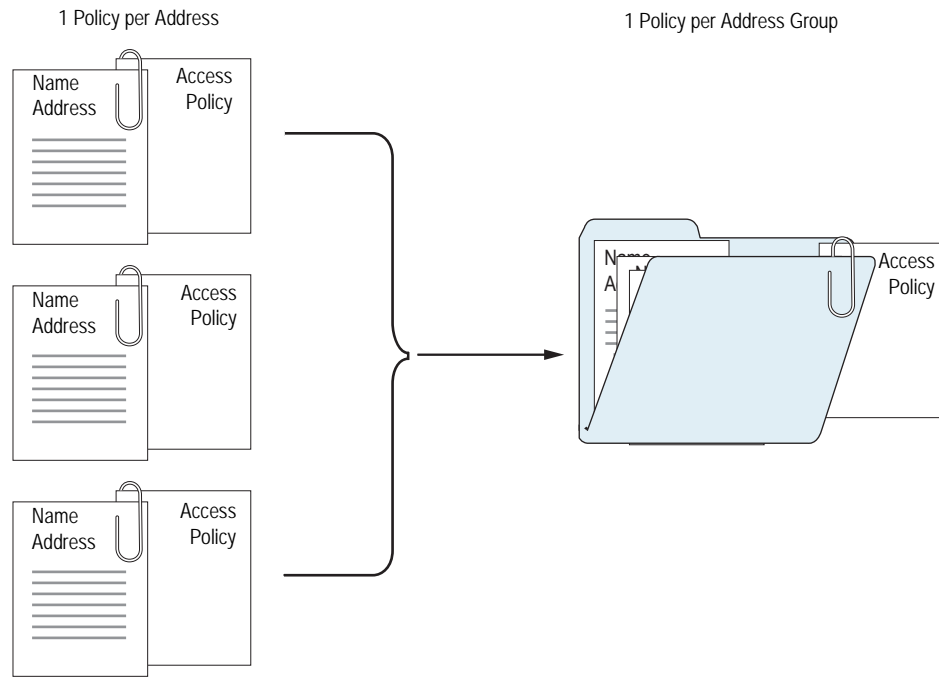
CLI

```
unset address trust "Sunnyvale_SW_Eng"
save
```

Address Groups

“Address Entries” on page 104 explained how you create, modify, and delete address book entries for individual hosts and subnets. As you add addresses to an address list, it becomes difficult to manage how policies affect each address entry. ScreenOS allows you to create groups of addresses. Rather than manage a large number of address entries, you can manage a small number of groups. Changes you make to the group are applied to each address entry in the group.

Figure 48: Address Groups



The address group option has the following features:

- You can create address groups in any zone.
- You can create address groups with existing users, or you can create empty address groups and later fill them with users.
- An address group can be a member of another address group.

NOTE: To ensure that a group does not accidentally contain itself as a member, the security device performs a sanity check when you add one group to another. For example, if you add group A as a member to group B, the security device automatically checks that A does not already contain B as its member.

- You can reference an address group entry in a policy like an individual address book entry.
 - ScreenOS applies policies to each member of the group by internally creating individual policies for each group member. While you only have to create one policy for a group, ScreenOS actually creates an internal policy for each member in the group (as well as for each service configured for each user).
-

NOTE: The automatic nature by which the security device applies policies to each address group member, saves you from having to create them one by one for each address. Furthermore, ScreenOS writes these policies to ASIC which makes lookups run very fast.

- When you delete an individual address book entry from the address book, the security device automatically removes it from all groups to which it belonged.

The following constraints apply to address groups:

- Address groups can only contain addresses that belong to the same zone.
- Address names cannot be the same as group names. If the name “Paris” is used for an individual address entry, it cannot be used for a group name.
- If an address group is referenced in a policy, the group cannot be removed. It can, however, be edited.
- When a single policy is assigned to an address group, it is applied to each group member individually, and the security device makes an entry for each member in the access control list (ACL). If you are not vigilant, it is possible to exceed the number of available policy resources, especially if both the source and destination addresses are address groups and the specified service is a service group.
- You cannot add the predefined addresses: “Any,” “All Virtual IPs,” and “Dial-Up VPN” to groups.

Creating an Address Group

In the following example, you create a group named “HQ 2nd Floor” that includes “Santa Clara Eng” and “Tech Pubs,” two addresses that you have already entered in the address book for the Trust zone.

WebUI

Policy > Policy Elements > Addresses > Groups > (for Zone: Trust) New:
Enter the following group name, move the following addresses, then click **OK**:

Group Name: HQ 2nd Floor

Select **Santa Clara Eng** and use the < < button to move the address from the Available Members column to the Group Members column.

Select **Tech Pubs** and use the < < button to move the address from the Available Members column to the Group Members column.

CLI

```
set group address trust "HQ 2nd Floor" add "Santa Clara Eng"
set group address trust "HQ 2nd Floor" add "Tech Pubs"
save
```

Editing an Address Group Entry

In this example, you add “Support” (an address that you have already entered in the address book) to the “HQ 2nd Floor” address group.

WebUI

Policy > Policy Elements > Addresses > Groups > (for Zone: Trust) Edit (for HQ 2nd Floor): Move the following address, then click **OK**:

Select **Support** and use the < < button to move the address from the Available Members column to the Group Members column.

CLI

```
set group address trust "HQ 2nd Floor" add Support
save
```

Removing a Member and a Group

In this example, you remove the member “Support” from the HQ 2nd Floor address group, and delete “Sales,” an address group that you had previously created.

WebUI

Policy > Policy Elements > Addresses > Groups > (for Zone: Trust) Edit (HQ 2nd Floor): Move the following address, then click **OK**:

Select **Support** and use the > > button to move the address from the Group Members column to the Available Members column.

Policy > Policy Elements > Addresses > Groups > (Zone: Trust): Click **Remove** in the Configure column for Sales.

CLI

```
unset group address trust "HQ 2nd Floor" remove Support
unset group address trust Sales
save
```

NOTE: The security device does not automatically delete a group from which you have removed all names.

Services

Services are types of traffic for which protocol standards exist. Each service has a transport protocol and destination port number(s) associated with it, such as TCP/port 21 for FTP and TCP/port 23 for Telnet. When you create a policy, you must specify a service for it. You can select one of the predefined services from the service book, or a custom service or service group that you created. You can see which service you can use in a policy by viewing the Service drop-down list on the Policy Configuration page (WebUI) or by using the **get service** command (CLI).

Predefined Services

You can view the list of predefined or custom services or service groups on the security device using the WebUI or the CLI.

- Using the WebUI:
 - Policy > Policy Elements > Services > Predefined
 - Policy > Policy Elements > Services > Custom
 - Policy > Policy Elements > Services > Groups
- Using the CLI:


```
get service [ group | predefined | user ]
```

NOTE: Each predefined service has a source port range of 1-65535, which includes the entire set of valid port numbers. This prevents potential attackers from gaining access by using a source port outside of the range. If you need to use a different source port range for any predefined service, create a custom service. For information, see “Custom Services” on page 121.

This section contains information about the following predefined services:

- “Internet Control Messaging Protocol” on page 110
- “Internet-Related Predefined Services” on page 114
- “Microsoft Remote Procedure Call Services” on page 115
- “Dynamic Routing Protocols” on page 117
- “Streaming Video” on page 117
- “Sun Remote Procedure Call Services” on page 118
- “Security and Tunnel Services” on page 118
- “IP-Related Services” on page 119
- “Instant Messaging Services” on page 119
- “Management Services” on page 119

- “Mail Services” on page 120
- “UNIX Services” on page 120
- “Miscellaneous Services” on page 121

You can find more detailed information about some of these listed on the following pages:

- “Defining a Custom Internet Control Message Protocol Service” on page 126
- “Remote Shell ALG” on page 127
- “Sun Remote Procedure Call Application Layer Gateway” on page 127
- “Customizing Microsoft Remote Procedure Call ALG” on page 128
- “Real-Time Streaming Protocol Application Layer Gateway” on page 129

Internet Control Messaging Protocol

Internet Control Messaging Protocol (ICMP) is a part of IP and provides a way to query a network (ICMP Query Messages) and to receive feedback from the network for error patterns (ICMP Error Messages). ICMP does not, however, guarantee error message delivery or report all lost data grams; and it is not a reliable protocol. ICMP code and type codes describe ICMP Query Messages and ICMP Error Messages.

You can choose to permit or deny any or specific types of ICMP messages to improve network security. Some types of ICMP messages can be exploited to gain information about your network that might compromise security. For example, ICMP, TCP, or UDP packets can be constructed to return ICMP error messages that contain information about a network, such as its topology, and access list filtering characteristics. Table 8 on page 111 lists ICMP message names, the corresponding code, type, and description.

Table 8: ICMP Information (page 1 of 3)

ICMP Message Name	Type	Code	Description
ICMP-ANY	all	all	ICMP-ANY affects any protocol using ICMP. Denying ICMP-ANY impairs any attempt to ping or monitor a network using ICMP. Permitting ICMP-ANY allows all ICMP messages.
ICMP-ADDRESS-MASK			ICMP Address Mask Query is used for systems that need the local subnet mask from a bootstrap server.
■ Request	17	0	
■ Reply	18	0	Denying ICMP Address Mask request messages can adversely affect diskless systems. Permitting ICMP Address Mask request messages might allow others to fingerprint the operating system of a host in your network.
ICMP-DEST-UNREACH	3	0	ICMP Destination Unreachable Error message indicates that the destination host is configured to reject the packets. Codes 0, 1, 4, or 5 can be from a gateway. Codes 2 or 3 from a host (RFC 792). Denying ICMP Destination Unreachable Error messages can remove the assumption that a host is up and running behind a security device. Permitting ICMP Destination Unreachable Error messages can allow some assumptions, such as security filtering, to be made about the network.
ICMP Fragment Needed	3	4	ICMP Fragmentation Error message indicates that fragmentation is needed but the don't fragment flag is set. Denying these messages from the Internet (untrust) to the trusted network is recommended.
ICMP Fragment Reassembly	11	1	ICMP Fragment Reassembly Time Exceeded Error indicates that a host reassembling a fragmented message ran out of time and dropped the packet. This message is sometimes sent. Denying these messages from the Internet (untrust) to the trusted network is recommended.
ICMP-HOST-UNREACH	3	1	ICMP Host Unreachable Error messages indicate that routing table entries do not list or list as infinity a particular host. Sometimes this error is sent by gateways that cannot fragment when a packet requiring fragmentation is received. Denying these messages from the Internet (untrust) to the trusted network is recommended. Permitting these messages allows others to be able to determine your internal hosts IP addresses by a process of elimination or make assumptions about gateways and fragmentation.
ICMP-INFO			ICMP-INFO Query messages allow diskless host systems to query the network and self-configure.
■ Request	15	0	
■ Reply	16	0	Denying ICMP Address Mask request messages can adversely affect diskless systems. Permitting ICMP Address Mask request messages might allow others to broadcast information queries to a network segment to determine computer type.
ICMP-PARAMETER-PROBLEM	12	0	ICMP Parameter Problem Error messages notify you when incorrect header parameters are present and caused a packet to be discarded Denying these messages from the Internet (untrust) to a trusted network is recommended. Permitting ICMP Parameter Problem error messages allows others to make assumptions about your network.

Table 8: ICMP Information (page 2 of 3)

ICMP Message Name	Type	Code	Description
ICMP-PORT-UNREACH	3	3	ICMP Port Unreachable Error messages indicate that gateways processing data grams requesting certain ports are unavailable or unsupported in the network. Denying these messages from the Internet (untrust) to the trusted network is recommended. Permitting ICMP Port Unreachable Error messages can allow others to determine which ports you use for certain protocols.
ICMP-PROTOCOL-UNREACH	3	2	ICMP Protocol Unreachable Error messages indicate that gateways processing data grams requesting certain protocols are unavailable or unsupported in the network. Denying these messages from the Internet (untrust) to the trusted network is recommended. Permitting ICMP Protocol Unreachable Error messages can allow others to determine what protocols your network is running.
ICMP-REDIRECT	5	0	ICMP Redirect Network Error messages are sent by routers. Denying these messages from the Internet (untrust) to the trusted network is recommended.
ICMP-REDIRECT-HOST	5	1	ICMP redirect messages indicate data grams destined for the specified host to be sent along another path.
ICMP-REDIRECT-TOS-HOST	5	3	ICMP Redirect Type of Service (TOS) and Host Error is a type of message.
ICMP-REDIRECT-TOS-NET	5	2	ICMP Redirect TOS and Network Error is a type of message.
ICMP-SOURCE-QUENCH	4	0	ICMP Source Quench Error message indicates that a router does not have the buffer space available to accept, queue, and send the packets on to the next hop. Denying these messages will not help or impair internal network performance. Permitting these messages can allow others to know that a router is congested making it a viable attack target.
ICMP-SOURCE-ROUTE-FAIL	3	5	ICMP Source Route Failed Error message Denying these messages from the Internet (untrust) is recommended.
ICMP-TIME-EXCEEDED	11	0	ICMP Time to Live (TTL) Exceeded Error message indicates that a packet's TTL setting reached zero before the packet reached its destination. This ensures that older packets are discarded before resent ones are processed. Denying these messages from a trusted network out to the Internet is recommended.
ICMP-TIMESTAMP			ICMP-TIMESTAMP Query messages provide the mechanism to synchronize time and coordinate time distribution in a large, diverse network.
■ Request	13	0	
■ Reply	14	0	

Table 8: ICMP Information (page 3 of 3)

ICMP Message Name	Type	Code	Description
Ping (ICMP ECHO)	8	0	Packet Internet Groper is a utility to determine whether a specific host is accessible by its IP address.0/0 echo reply. Denying ping functionality removes your ability to check to see if a host is active. Permitting ping can allow others to execute a denial of service (DoS) or Smurf attack.
ICMP-ECHO-FRAGMENT-ASSEMBLY-EXPIRE	11	1	ICMP Fragment Echo Reassembly Time Expired error message indicates that the reassembly time was exceeded. Denying these messages is recommended.
Traceroute			Traceroute is a utility to indicate the path to access a specific host.
■ Forward	30	0	Denying this utility from the Internet (untrust) to your internal network (trust) is recommended.
■ Discard	30	1	

Handling ICMP Unreachable Errors

For different levels of security, the default behavior for ICMP unreachable errors from downstream routers is handled as follows:

- Sessions do not close for ICMP type 3 code 4 messages.
ICMP messages pass through without dropping sessions. Packets are however, dropped per session.
- Sessions do not close on receiving any kind of ICMP unreachable messages.
- Sessions store ICMP unreachable message, thereby restricting the number of messages flowing through to 1.

One ICMP unreachable message is generated globally per device. The remaining ICMP unreachable errors are dropped.

Internet-Related Predefined Services

Table 9 lists Internet-related predefined services. Depending on your network requirements, you can choose to permit or deny any or all of these services. Each entry lists the service name, default receiving port, and service description.

Table 9: Predefined Services

Service Name	Port(s)	Service Description
AOL	5190-5193	America Online Internet Service Provider (ISP) provides Internet, chat, and instant messaging services.
DHCP-Relay	67 (default)	Dynamic Host Configuration.
DHCP	68 (client) 67 (server)	Dynamic Host Configuration Protocol allocates network addresses and delivers configuration parameters from server to hosts.
DNS	53	Domain Name System translates domain names into IP addresses.
FTP ■ FTP-Get ■ FTP-Put	20 (data) 21 (control)	File Transfer Protocol (FTP) allows the sending and receiving of files between machines. You can choose to deny or permit ANY (GET or PUT) or to selectively permit or deny either GET or PUT. GET receives files from another machine and PUT sends files to another machine. We recommend denying FTP services from untrusted sources (Internet).
Gopher	70	Gopher organizes and displays Internet servers' contents as a hierarchically structured list of files. We recommend denying Gopher access to avoid exposing your network structure.
HTTP	8080	HyperText Transfer Protocol is the underlying protocol used by the World Wide Web (WWW). Denying HTTP service disables your users from viewing the Internet. Permitting HTTP service allows your trusted hosts to view the Internet.
HTTP-EXT	—	Hypertext Transfer Protocol with extended non-standard ports
HTTPS	443	Hypertext Transfer Protocol with Secure Socket Layer (SSL) is a protocol for transmitting private documents through the Internet. Denying HTTPS disables your users from shopping on the Internet and from accessing certain online resources that require secure password exchange. Permitting HTTPS allows your trusted hosts to participate in password exchange, shop online and visit various protected online resources that require user login.
Internet Locator Service	—	Internet Locator Service includes LDAP, User Locator Service, and LDAP over TSL/SSL.
IRC	6665-6669	Internet Relay Chat (IRC) allows people connected to the Internet to join live discussions.
LDAP	389	Lightweight Directory Access Protocol is a set of protocols used to access information directories.
PC-Anywhere	—	PC-Anywhere is a remote control and file transfer software.
TFTP	69	TFTP is a protocol for simple file transfer.
WAIS	—	Wide Area Information Server is a program that finds documents on the Internet.

Microsoft Remote Procedure Call Services

Table 10 lists predefined Microsoft services, parameters associated with each service, and a brief description of each service. Parameters include Universal Unique Identifiers (UUIDs) and TCP/UDP source and destination ports. A UUID is a 128-bit unique number generated from a hardware address, a timestamp, and seed values.

Table 10: Microsoft Services (page 1 of 3)

Service	Parameter/UUID	Description
MS-RPC-EPM	135 e1af8308-5d1f-11c9-91a4-08002b14a0fa	Microsoft Remote Procedure Call (RPC) Endpoint Mapper (EPM) Protocol
MS-RPC-ANY	—	Any Microsoft Remote Procedure Call (RPC) Services
MS-AD	4 members	Microsoft Active Directory Service Group includes: <ul style="list-style-type: none"> ■ MS-AD-BR ■ MS-AD-DRSUAPI ■ MS-AD-DSROLE ■ MS-AD-DSSETUP
MS-EXCHANGE	6 members	Microsoft Exchange Service Group includes: <ul style="list-style-type: none"> ■ MS-EXCHANGE-DATABASE ■ MS-EXCHANGE-DIRECTORY ■ MS-EXCHANGE-INFO-STORE ■ MS-EXCHANGE-MTA ■ MS-EXCHANGE-STORE ■ MS-EXCHANGE-SYSATD
MS-IIS	6 members	Microsoft IIS Server Service Group includes: <ul style="list-style-type: none"> ■ MS-IIS-COM ■ MS-IIS-IMAP4 ■ MS-IIS-INETINFO ■ MS-IIS-NNTP ■ MS-IIS-POP3 ■ MS-IIS-SMTP
MS-AD-BR	ecec0d70-a603-11d0-96b1-00a0c91ece30 16e0cf3a-a604-11d0-96b1-00a0c91ece30	Microsoft Active Directory Backup and Restore Services
MS-AD-DRSUAPI	e3514235-4b06-11d1-ab04-00c04fc2dcd2	Microsoft Active Directory Replication Service
MS-AD-DSROLE	1cbcad78-df0b-4934-b558-87839ea501c9	Microsoft Active Directory DSROLE Service
MS-AD-DSSETUP	3919286a-b10c-11d0-9ba8-00c04fd92ef5	Microsoft Active Directory Setup Service
MS-DTC	906b0ce0-c70b-1067-b317-00dd010662da	Microsoft Distributed Transaction Coordinator Service
MS-EXCHANGE-DATABASE	1a190310-bb9c-11cd-90f8-00aa00466520	Microsoft Exchange Database Service
MS-EXCHANGE-DIRECTORY	f5cc5a18-4264-101a-8c59-08002b2f8426 f5cc5a7c-4264-101a-8c59-08002b2f8426 f5cc59b4-4264-101a-8c59-08002b2f8426	Microsoft Exchange Directory Service

Table 10: Microsoft Services (page 2 of 3)

Service	Parameter/UUID	Description
MS-EXCHANGE-INFO-STORE	0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde 1453c42c-0fa6-11d2-a910-00c04f990f3b 10f24e8e-0fa6-11d2-a910-00c04f990f3b 1544f5e0-613c-11d1-93df-00c04fd7bd09	Microsoft Exchange Information Store Service
MS-EXCHANGE-MTA	9e8ee830-4459-11ce-979b-00aa005ffebe 38a94e72-a9bc-11d2-8faf-00c04fa378ff	Microsoft Exchange MTA Service
MS-EXCHANGE-STORE	99e66040-b032-11d0-97a4-00c04fd6551d 89742ace-a9ed-11cf-9c0c-08002be7ae86 a4f1db00-ca47-1067-b31e-00dd010662da a4f1db00-ca47-1067-b31f-00dd010662da	Microsoft Exchange Store Service
MS-EXCHANGE-SYSATD	67df7c70-0f04-11ce-b13f-00aa003bac6c f930c514-1215-11d3-99a5-00a0c9b61b04 83d72bf0-0d89-11ce-b13f-00aa003bac6c 469d6ec0-0d87-11ce-b13f-00aa003bac6c 06ed1d30-d3d3-11cd-b80e-00aa004b9c30	Microsoft Exchange System Attendant Service
MS-FRS	f5cc59b4-4264-101a-8c59-08002b2f8426 d049b186-814f-11d1-9a3c-00c04fc9b232 a00c021c-2be2-11d2-b678-0000f87a8f8e	Microsoft File Replication Service
MS-IIS-COM	70b51430-b6ca-11d0-b9b9-00a0c922e750 a9e69612-b80d-11d0-b9b9-00a0c922e70	Microsoft Internet Information Server COM GUID/UUID Service
MS-IIS-IMAP4	2465e9e0-a873-11d0-930b-00a0c90ab17c	Microsoft Internet Information Server IMAP4 Service
MS-IIS-INETINFO	82ad4280-036b-11cf-972c-00aa006887b0	Microsoft Internet Information Server Administration Service
MS-IIS-NNTP	4f82f460-0e21-11cf-909e-00805f48a135	Microsoft Internet Information Server NNTP Service
MS-IIS-POP3	1be617c0-31a5-11cf-a7d8-00805f48a135	Microsoft Internet Information Server POP3 Service
MS-IIS-SMTP	8cfb5d70-31a4-11cf-a7d8-00805f48a135	Microsoft Internet Information Server SMTP Service
MS-ISMSERV	68dcd486-669e-11d1-ab0c-00c04fc2dcd2 130ceefb-e466-11d1-b78b-00c04fa32883	Microsoft Inter-site Messaging Service
MS-MESSENGER	17fdd703-1827-4e34-79d4-24a55c53bb37 5a7b91f8-ff00-11d0-a9b2-00c04fb6e6fc	Microsoft Messenger Service
MS-MQQM	fdb3a030-065f-11d1-bb9b-00a024ea5525 76d12b80-3467-11d3-91ff-0090272f9ea3 1088a980-eae5-11d0-8d9b-00a02453c33 5b5b3580-b0e0-11d1-b92d-0060081e87f0 41208ee0-e970-11d1-9b9e-00e02c064c39	Microsoft Windows Message Queue Management Service
MS-NETLOGON	12345678-1234-abcd-ef00-01234567cffb	Microsoft Netlogon Service

Table 10: Microsoft Services (page 3 of 3)

Service	Parameter/UUID	Description
MS-SCHEDULER	1ff70682-0a51-30e8-076d-740be8cee98b 378e52b0-c0a9-11cf-822d-00aa0051e40f 0a74ef1c-41a4-4e06-83ae-dc74fb1cdd53	Microsoft Scheduler Service
MS-WIN-DNS	50abc2a4-574d-40b3-9d66-ee4fd5fba076	Microsoft Windows DNS Server
MS-WINS	45f52c28-7f9f-101a-b52b-08002b2efabe 811109bf-a4e1-11d1-ab54-00a0c91e9b45	Microsoft WINS Service

Dynamic Routing Protocols

Depending on your network requirements, you can choose to permit or deny messages generated from and packets of these dynamic routing protocols. Table 11 lists each supported dynamic routing protocol by name, port, and description.

Table 11: Dynamic Routing Protocols

Dynamic Routing Protocol	Port	Description
RIP	520	RIP is a common distance-vector routing protocol.
OSPF	89	OSPF is a common link-state routing protocol.
BGP	179	BGP is an exterior/interdomain routing protocol.

Streaming Video

Table 12 lists each supported streaming video service by name and includes the default port and description. Depending on your network requirements, you can choose to permit or deny any or all of these services.

Table 12: Streaming Video Services

Service	Port	Description
H.323	TCP source 1-65535; TCP destination 1720, 1503, 389, 522, 1731 UDP source 1-65535; UDP source 1719	H.323 is a standard approved by the International Telecommunication Union (ITU) that defines how audiovisual conference data is transmitted across networks.
NetMeeting	TCP source 1-65535; TCP destination 1720, 1503, 389, 522 UDP source 1719	Microsoft NetMeeting uses TCP to provide teleconferencing (video and audio) services over the Internet.
Real media	TCP source 1-65535; TCP destination 7070	Real Media is streaming video and audio technology.
RTSP	554	Real Time Streaming Protocol (RTSP) for streaming media applications
SIP	5056	Session Initiation Protocol (SIP) is an Application Layer control protocol for creating, modifying, and terminating sessions.
VDO Live	TCP source 1-65535; TCP destination 7000-7010	VDOLive is a scalable, video streaming technology.

Sun Remote Procedure Call Services

Table 13 lists each Sun Remote Procedure Call Application Layer Gateway (RPC ALG) service name, parameters, and full name.

Table 13: RPC ALG Services

Service	Program Numbers	Full Name
SUN-RPC-PORTMAPPER	111 100000	Sun RPC Portmapper Protocol
SUN-RPC-ANY	ANY	Any Sun RPC services
SUN-RPC-PROGRAM-MOUNTD	100005	Sun RPC Mount Daemon
SUN-RPC-PROGRAM-NFS	100003 100227	Sun RPC Network File System
SUN-RPC-PROGRAM-NLOCKMGR	100021	Sun RPC Network Lock Manager
SUN-RPC-PROGRAM-RQUOTAD	100011	Sun RPC Remote Quota Daemon
SUN-RPC-PROGRAM-RSTATD	100001	Sun RPC Remote Status Daemon
SUN-RPC-PROGRAM-RUSERD	100002	Sun RPC Remote User Daemon
SUN-RPC-PROGRAM-SADMIND	100232	Sun RPC System Administration Daemon
SUN-RPC-PROGRAM-SPRAYD	100012	Sun RPC SPRAY Daemon
SUN-RPC-PROGRAM-STATUS	100024	Sun RPC STATUS
SUN-RPC-PROGRAM-WALLD	100008	Sun RPC WALL Daemon
SUN-RPC-PROGRAM-YPBIND	100007	SUN RPC Yellow Page Bind Service

Security and Tunnel Services

Table 14 lists each supported service and gives the default port(s) and a description of each entry.

Table 14: Supported Protocol Services

Service	Port	Description
IKE	UDP source 1-65535; UDP destination 500 4500 (used for NAT traversal)	IKE is a protocol to obtain authenticated keying material for use with ISAKMP. When configuring auto IKE, you can choose from three predefined Phase 1 or Phase 2 proposals: <ul style="list-style-type: none"> ■ Standard: AES and 3DES ■ Basic: DES and two different types of authentication algorithms ■ Compatible: four commonly used authentication and encryption algorithms
L2TP	1723	L2TP combines PPTP with Layer 2 Forwarding (L2F) for remote access.
PPTP	—	Point-to-Point Tunneling Protocol allows corporations to extend their own private network through private <i>tunnels</i> over the public Internet.

IP-Related Services

Table 15 lists the predefined IP-related services. Each entry includes the default port and a description of the service.

Table 15: IP-Related Services

Service	Port	Description
Any	—	Any service
TCP-ANY	1-65535	Any protocol using the Transport Control Protocol TCPMUX port 1
UDP-ANY	137	Any protocol using the User Datagram Protocol

Instant Messaging Services

Table 16 lists predefined Internet-messaging services. Each entry includes the name of the service, the default or assigned port, and a description of the service.

Table 16: Internet-Messaging Services

Service	Port	Description
Gnutella	6346 (default)	Gnutella File Sharing Protocol is a public domain file sharing protocol that operates over a distributed network. You can assign any port, but the default is 6346.
MSN	1863	Microsoft Network Messenger is a utility that allows you to send instant messages and talk online.
NNTP	119	Network News Transport Protocol is a protocol used to post, distribute, and retrieve USENET messages.
SMB	445	Server Message Block (SMB) Protocol over IP allows you to read and write files to a server on a network.
YMSG	5010	Yahoo! Messenger is a utility that allows you to check when others are online, send instant messages, and talk online.

Management Services

Table 17 lists the predefined management services. Each entry includes the name of the service, the default or assigned port, and a description of the service.

Table 17: Management Services (page 1 of 2)

Service	Port	Description
NBNAME	137	NetBIOS Name Service displays all NetBIOS name packets sent on UDP port 137.
NDBDS	138	NetBIOS Datagram Service, published by IBM, provides connectionless (datagram) services to PCs connected with a broadcast medium to locate resources, initiate sessions and terminate sessions. It is unreliable and the packets are not sequenced.
NFS	—	Network File System uses UDP to allow network users to access shared files stored on computers of different types. SUN RPC is a building block of NFS.
NS Global	—	NS-Global is the central management protocol for Juniper Networks Firewall/VPN devices.
NS Global PRO	—	NS Global-PRO is the scalable monitoring system for the Juniper Networks Firewall/VPN device family.
NSM	—	NetScreen-Security Manager
NTP	123	Network Time Protocol provides a way for computers to synchronize to a time referent.
RLOGIN	513	RLOGIN starts a terminal session on a remote host.

Table 17: Management Services (page 2 of 2)

Service	Port	Description
RSH	514	RSH executes a shell command on a remote host.
SNMP	161	Simple Network Management Protocol is a set of protocols for managing complex networks.
SQL*Net V1	66	SQL*Net Version 1 is a database language that allows for the creation, access, modification, and protection of data.
SQL*Net V2	66	SQL*Net Version 2 is a database language that allows for the creation, access, modification, and protection of data.
MSSQL	1433 (default instance)	Microsoft SQL is a proprietary database server tool that allows for the creation, access, modification, and protection of data.
SSH	22	Secure Shell is a program to log into another computer over a network through strong authentication and secure communications on an unsecure channel.
SYSLOG	514	Syslog is a UNIX program that sends messages to the system logger.
Talk	517-518	Talk is a visual communication program that copies lines from your terminal to that of another user.
Telnet	23	Telnet is a UNIX program that provides a standard method of interfacing terminal devices and terminal-oriented processes to each other.
WinFrame	—	WinFrame is a technology that allows users on non-Windows machines to run Windows applications.
X-Windows	—	X-Windows is the windowing and graphics system that Motif and OpenLook are based on.

Mail Services

Table 18 lists the predefined mail services. Each includes the name of the service, the default or assigned port number, and a description of the service.

Table 18: Mail Services

Service	Port	Description
IMAP	143	Internet Message Access Protocol is a protocol used for retrieving messages.
Mail (SMTP)	25	Simple Mail Transfer Protocol is a protocol for sending messages between servers.
POP3	110	Post office protocol is a protocol used for retrieving email.

UNIX Services

Table 19 lists the predefined UNIX services. Each entry includes the name of the service, the default or assigned port, and a description of the service.

Table 19: UNIX Services

Service	Port	Description
FINGER	79	Finger is a UNIX program that provides information about the users.
UUCP	117	Unix-to-Unix Copy Protocol (UUCP) is a UNIX utility that enables file transfers between two computers over a direct serial or modem connection.

Miscellaneous Services

Table 20 lists predefined miscellaneous services. Each entry includes the service name, default or assigned port, and a description of the service.

Table 20: Miscellaneous Services

Service	Port	Description
CHARGEN	19	Character Generator Protocol is a UDP or TCP-based debugging and measurement tool.
DISCARD	9	Discard Protocol is an application layer protocol that describes a process for discarding TCP or UDP data sent to port 9.
IDENT	113	Identification Protocol is a TCP/IP application layer protocol used for TCP client authentication.
LPR	515 listen; 721-731 source range (inclusive)	Line Printer Daemon Protocol is a TCP-based protocol used for printing services.
RADIUS	1812	Remote Authentication Dial-In User Service is a server program used for authentication and accounting purposes.
SQLMON	1434 (SQL Monitor Port)	SQL monitor (Microsoft)
VNC	5800	Virtual Network Computing facilitates viewing and interacting with another computer or mobile device connected to the Internet.
WHOIS	43	Network Directory Service Protocol is a way to look up domain names.
IPSEC-NAT	—	IPSEC-NAT allows Network Address Translation for ISAKMP and ESP packets.
SCCP	2000	Cisco Station Call Control Protocol uses the signaling connection control port (SCCP) to provide high availability and flow control.
VOIP	—	Voice over IP Service Group provides voice services over the Internet and includes H.323 and Session Initiation Protocol (SIP).

Custom Services

Instead of using predefined services, you can easily create custom services. You can assign each custom service the following attributes:

- Name
- Transport protocol
- Source and destination port numbers for services using TCP or UDP
- Type and code values for services using ICMP
- Timeout value

If you create a custom service in a virtual system (vsys) that has the same name as a previously defined custom service in the root system, the service in the vsys takes the default timeout for the specified transport protocol (TCP, UDP, or ICMP). To define a custom timeout for a service in a vsys that is different from the default when a custom service with the same name in the root system has its own timeout, create the custom service in the vsys and root system in the following order:

1. First, create the custom service with a custom timeout in the vsys.
2. Then create another custom service with the same name but a different timeout in the root system.

The following examples describe how to add, modify, and remove a custom service.

NOTE: For information regarding ScreenOS naming conventions—which apply to the names you create for custom services—see “Command Line Interface Conventions” on page x.

Adding a Custom Service

To add a custom service to the service book, you need the following information:

- A name for the service: in this example “cust-telnet.”
- A range of source port numbers: 1 – 65535.
- A range of destination port numbers to receive the service request: for example: 23000 – 23000.
- Whether the service uses TCP or UDP protocol, or some other protocol as defined by the Internet specifications. In this example, the protocol is TCP.

WebUI

Policy > Policy Elements > Services > Custom > New: Enter the following, then click **OK**:

```
Service Name: cust-telnet
Service Timeout: Custom (select), 30 (type)
Transport Protocol: TCP (select)
Source Port Low: 1
Source Port High: 65535
Destination Port Low: 23000
Destination Port High: 23000
```

CLI

```
set service cust-telnet protocol tcp src-port 1-65535 dst-port 23000-23000
set service cust-telnet timeout 30
save
```

NOTE: The timeout value is in minutes. If you do not set it, the timeout value of a custom service is 180 minutes. If you do not want a service to time out, enter **never**.

Modifying a Custom Service

In this example, you modify the custom service “cust-telnet” by changing the destination port range to 23230-23230.

Use the **set service** *service_name* **clear** command to remove the definition of a custom service without removing the service from the service book:

WebUI

Policy > Policy Elements > Services > Custom > Edit (for cust-telnet): Enter the following, then click **OK**:

```
Destination Port Low: 23230
Destination Port High: 23230
```

CLI

```
set service cust-telnet clear
set service cust-telnet + tcp src-port 1-65535 dst-port 23230-23230
save
```

Removing a Custom Service

In this example, you remove the custom service “cust-telnet”.

WebUI

Policy > Policy Elements > Services > Custom: Click **Remove** in the Configure column for “cust-telnet”.

CLI

```
unset service cust-telnet
save
```

Setting a Service Timeout

The timeout value you set for a service determines the session timeout. You can set the timeout for a predefined or custom service; you can use the service default timeout, specify a custom timeout, or use no timeout at all. Service timeout behavior is the same in virtual systems (vsys) security domains as at the root level.

Service Timeout Configuration and Lookup

Service timeout values are stored in the service entry database and in the corresponding vsys TCP and UDP port-based timeout tables. When you set a service timeout value, the security device updates these tables with the new value. There are also default timeout values in the services entry database, which are taken from predefined services. You can set a timeout but you cannot alter the default values.

Services with multiple rule entries share the same timeout value. If multiple services share the same protocol and destination port range, all services share the last timeout value configured.

For single service entries, service timeout lookup proceeds as follows:

1. The specified timeout in the service entry database, if set.
2. The default timeout in the service entry database, if specified in the predefined service.
3. The protocol-based default timeout table.

Table 21: Protocol-Based Default Timeout Table

Protocol	Default Timeout (minutes)
TCP	30
UDP	1
ICMP	1
OSPF	1
Other	30

For service groups, including hidden groups created in multi-cell policy configurations, and for the predefined service “ANY” (if timeout is not set), service timeout lookup proceeds as follows:

1. The vsys TCP and UDP port-based timeout table, if a timeout is set.
2. The protocol-based default timeout table.

Contingencies

When setting timeouts, be aware of the following:

- If a service contains several service rule entries, all rule entries share the same timeout. The timeout table is updated for each rule entry that matches the protocol (for UDP and TCP—other protocols use the default). You need define the service timeout only once. For example, if you create a service with two rules, the following commands will set the timeout to 20 minutes for both rules:

```
set service test protocol tcp dst-port 1035-1035 timeout 20
set service test + udp src-port 1-65535 dst-port 1111-1111
```

- If multiple services are configured with the same protocol and overlapping destination ports, the latest service timeout configured overrides the others in the port-based table. For example:

```
set service ftp-1 protocol tcp src 0-65535 dst 2121-2121 timeout 10
set service telnet-1 protocol tcp src 0-65535 dst 2100-2148 timeout 20
```

With this configuration, the security device applies the 20-minute timeout for destination port 2121 in a service group, because the destination port numbers for telnet-1 (2100-2148) overlap those for ftp-1 (2121), and you defined telnet-1 after you defined ftp-1.

To modify a service timeout when multiple services use the same protocol and an overlapping destination port range, you must unset the service and reset it with the new timeout value. This is because, during reboot, services are loaded according to creation time, not modification time.

To avoid the unintended application of the wrong timeout to a service, do not create services with overlapping destination port numbers.

- If you unset a service timeout, the default protocol-based timeout in the service entry database is used, and the timeout values in both the service entry and port-based timeout tables are updated with the default value.

If the modified service has overlapping destination ports with other services, the default protocol-based timeout might not be the desired value. In that case, reboot the security device, or set the service timeout again for the desired timeout to take effect.

- When you modify a predefined service and reboot, the modified service might not be the last one in the configuration. This is because predefined services are loaded before custom services, and any change made to a custom service, even if made earlier, will show as the later than the predefined service change when you reboot.

For example, if you create the following service:

```
set service my_service protocol tcp dst-port 179-179 timeout 60
```

and later modify the timeout of the predefined service BGP as follows:

```
set service bgp timeout 75
```

the BGP service will use the 75-minute timeout value, because it is now written to the service entry database. But the timeout for port 179, the port BGP uses, is also changed to 75 in the TCP port-based timeout table. After you reboot, the BGP service will continue to use the 75-minute timeout which, as a single service, it gets from the service entry database. But the timeout in the TCP port-based table for port 179 will now be 60. You can verify this by entering the **get service bgp** command.

This has no effect on single services. But if you add BGP or my_service to a service group, the 60-minute timeout value will be used for destination port 179. This is because service group timeout is taken from the port-based timeout table, if one is set.

To ensure predictability when you modify a predefined service timeout, therefore, you can create a similar service, for example:

```
set service my-bgp protocol tcp dst-port 179-179 timeout 75
```

Example

In the following example, you change the timeout threshold for the FTP predefined service to 75 minutes:

WebUI

Policy > Policy Elements > Services > Predefined > Edit (FTP): Enter the following, then click **OK**:

Service Timeout: Custom (select), 75 (type)

CLI

```
set service FTP timeout 75
save
```

Defining a Custom Internet Control Message Protocol Service

ScreenOS supports Internet Control Message Protocol (ICMP) as well as several ICMP messages, as predefined or custom services. When configuring a custom ICMP service, you must define a type and code. There are different message types within ICMP. For example:

type 0 = Echo Request message

type 3 = Destination Unreachable message

NOTE: For more information about ICMP types and codes, refer to RFC 792, *Internet Control Message Protocol*.

An ICMP message type can also have a message code. The code provides more specific information about the message, as shown in Table 22.

Table 22: Message Descriptions

Message Type	Message Code
5 = Redirect	0 = Redirect Datagram for the Network (or subnet)
	1 = Redirect Datagram for the Host
	2 = Redirect Datagram for the Type of Service and Network
	3 = Redirect Datagram for the Type of Service and Host
11 = Time Exceeded Codes	0 = Time to Live exceeded in Transit
	1 = Fragment Reassembly Time Exceeded

ScreenOS supports any type or code within the 0-255 range.

In this example, you define a custom service named “host-unreachable” using ICMP as the transport protocol. The type is 3 (for Destination Unreachable) and the code is 1 (for Host Unreachable). You set the timeout value at 2 minutes.

WebUI

Policy > Policy Elements > Services > Custom: Enter the following, then click **OK**:

```
Service Name: host-unreachable
Service Timeout: Custom (select), 2 (type)
Transport Protocol: ICMP (select)
ICMP Type: 3
ICMP Code: 1
```

CLI

```

set service host-unreachable protocol icmp type 5 code 0
set service host-unreachable timeout 2
save

```

Remote Shell ALG

Remote Shell Application Layer Gateway (RSH ALG) allows authenticated users to run shell commands on remote hosts. Juniper Networks security devices support the RSH service in Transparent (L2), Route (L3), and NAT modes, but the devices do not support port translation of RSH traffic.

Sun Remote Procedure Call Application Layer Gateway

Sun Remote Procedure Call—also known as Open Network Computing Remote Procedure Call (ONC RPC)—provides a way for a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.

Juniper Networks security devices support Sun RPC as a predefined service and allow and deny traffic based on a policy you configure. The Application Layer Gateway (ALG) provides the functionality for security devices to handle the dynamic transport address negotiation mechanism of Sun RPC and to ensure program number-based firewall policy enforcement. You can define a firewall policy to permit or deny all RPC requests, or to permit or deny by specific program number. The ALG also supports Route and NAT mode for incoming and outgoing requests.

Typical RPC Call Scenario

When a client calls a remote service, it needs to find the transport address of the service—in the case of TCP/UDP, this is a port number. A typical procedure for this case is as follows:

1. The client sends the GETPORT message to the RPCBIND service on the remote machine. The GETPORT message contains the program number, and version and procedure number of the remote service it wants to call.
2. The RPCBIND service replies with a port number.
3. The client calls the remote service using the port number returned.
4. The remote service replies to the client.

A client also can use the CALLIT message to call the remote service directly, without knowing the port number of the service. In this case, the procedure is as follows:

1. The client sends a CALLIT message to the RPCBIND service on the remote machine. The CALLIT message contains the program number, and the version and procedure number of the remote service it wants to call.
2. RPCBIND calls the service for the client.

3. RCPBIND replies to the client if the call has been successful. The reply contains the call result and the services's port number.

Customizing Sun RPC Services

Because Sun RPC services use dynamically negotiated ports, you can not use regular service objects based on fixed TCP/UDP ports to permit them in security policy. Instead, you must create sun rpc service objects using program numbers. For example, NFS uses two program numbers: 100003 and 100227. The corresponding TCP/UDP ports are dynamic. In order to permit the program numbers, you create a sun-rpc-nfs service object that contains these two numbers. The ALG maps the program numbers into dynamically negotiated TCP/UDP ports and permits or denies the service based on a policy you configure.

In this example, you create a service object called my-sunrpc-nfs to use the Sun RPC Network File System, which is identified by two Program IDs: 100003 and 100227.

WebUI

Policy > Policy Elements > Services > Sun RPC Services > New: Enter the following, then click **Apply**:

```
Service Name: my-sunrpc-nfs
Service Timeout: (select)
Program ID Low: 100003
Program ID High: 100003
Program ID Low: 100227
Program ID High: 100227
```

CLI

```
set service my-sunrpc-nfs protocol sun-rpc program 100003-100003
set service my-sunrpc-nfs + sun-rpc program 100227-100227
save
```

Customizing Microsoft Remote Procedure Call ALG

Microsoft Remote Procedure Call (MS RPC) is the Microsoft implementation of the Distributed Computing Environment (DCE) RPC. Like the Sun RPC (see “Sun Remote Procedure Call Application Layer Gateway” on page 127), MS RPC provides a way for a program running on one host to call procedures in a program running on another host. Because of the large number of RPC services and the need to broadcast, the transport address of an RPC service is dynamically negotiated based on the service program's Universal Unique Identifier (UUID). The Endpoint Mapper binding protocol is defined in ScreenOS to map the specific UUID to a transport address.

Juniper Networks security devices support MS RPC as a predefined service; they allow and deny traffic based on a policy you configure. The ALG provides the functionality for security devices to handle the dynamic transport address negotiation mechanism of MS RPC, and to ensure UUID-based firewall policy enforcement. You can define a firewall policy to permit or deny all RPC requests, or to permit or deny by specific UUID number. The ALG also supports Route and NAT mode for incoming and outgoing requests.

Because MS RPC services use dynamically negotiated ports, you can not use regular service objects based on fixed TCP/UDP ports to permit them in a security policy. Instead, you must create MS RPC service objects using UUIDs. The MS Exchange Info Store service, for example, uses the following four UUIDs:

- 0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde
- 1453c42c-0fa6-11d2-a910-00c04f990f3b
- 10f24e8e-0fa6-11d2-a910-00c04f990f3b
- 1544f5e0-613c-11d1-93df-00c04fd7bd09

The corresponding TCP/UDP ports are dynamic. To permit them, you create an ms-exchange-info-store service object that contains these four UUIDs. The ALG maps the program numbers into dynamically negotiated TCP/UDP ports based on these four UUIDs and permits or denies the service based on a policy you configure.

In this example, you create a service object called my-ex-info-store that includes the UUIDs for the MS Exchange Info Store service.

WebUI

Policy > Policy Elements > Services > MS RPC: Enter the following, then click **Apply**:

```
Service Name: my-ex-info-store
UUID: 0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde
UUID: 1453c42c-0fa6-11d2-a910-00c04f990f3b
UUID: 10f24e8e-0fa6-11d2-a910-00c04f990f3b
UUID: 1544f5e0-613c-11d1-93df-00c04fd7bd09
```

CLI

```
set service my-ex-info-store protocol ms-rpc uuid
0e4a0156-dd5d-11d2-8c2f-00c04fb6bcde
set service my-ex-info-store + ms-rpc uuid
1453c42c-0fa6-11d2-a910-00c04f990f3b
set service my-ex-info-store + ms-rpc uuid
10f24e8e-0fa6-11d2-a910-00c04f990f3b
set service my-ex-info-store + ms-rpc uuid
1544f5e0-613c-11d1-93df-00c04fd7bd09
save
```

Real-Time Streaming Protocol Application Layer Gateway

Real-Time Streaming Protocol (RTSP) is an Application Layer protocol used to control delivery of one or more synchronized streams of multimedia, such as audio and video. Although RTSP is capable of delivering the data streams itself—interleaving the continuous media streams with the control stream—it is more typically used as a kind of network remote control for multimedia servers. The protocol was designed as a means for selecting delivery channels, such as UDP, multicast UDP, and TCP, and for selecting delivery mechanism based on the Real Time Protocol (RTP). RTSP may also use the Session Description Protocol (SDP) as a

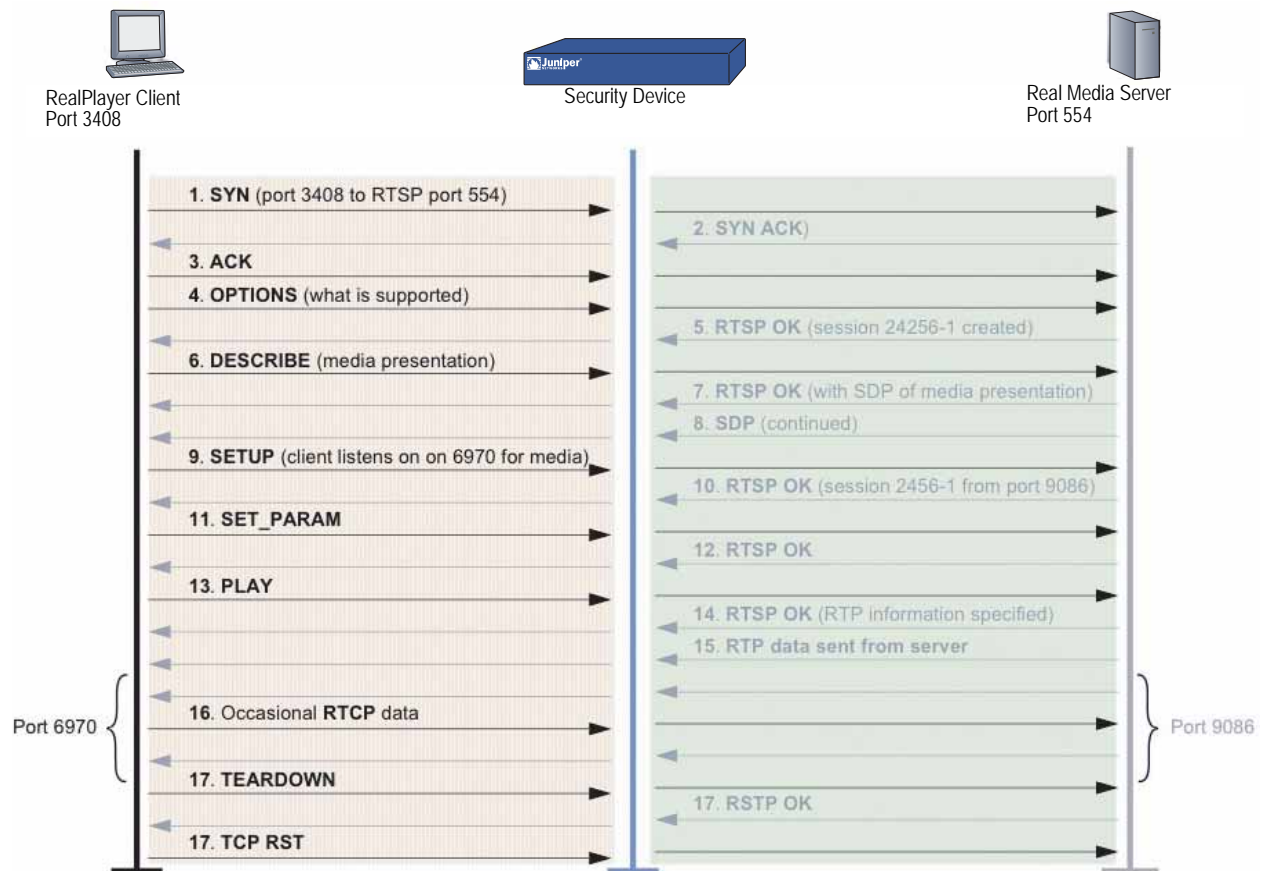
means of providing information to the client for aggregate control of a presentation composed of streams from one or more servers, and non-aggregate control of a presentation composed of multiple streams from a single server. The sources of data can be live feeds or stored clips.

Juniper Networks security devices support RTSP as a service and allow or deny RTSP traffic based on a policy you configure. The ALG is needed because RTSP uses dynamically assigned port numbers that are conveyed in the packet payload during control connection establishment. The ALG keeps track of the dynamically assigned port numbers and opens pinholes accordingly. In NAT mode, the ALG translates IP addresses and ports if necessary. Security devices support RTSP in Route mode, Transparent mode, and in both interface-based and policy-based NAT mode.

Figure 49 on page 131 diagrams a typical RTSP session. The client initiates the session (when the user clicks the Play button on a RealPlayer, for example) and establishes a TCP connection to the RTSP server on port 554, then sends the OPTIONS message (messages are also called methods), to find out what audio and video features the server supports. The server responds to the OPTIONS message by specifying the name and version of the server, and a session identifier, for example, 24256-1. (For more information about methods, see “SIP Request Methods” on page 6-14, and refer to RFC 2326, Section 11.)

The client then sends the DESCRIBE message with the URL of the actual media file it wants. The server responds to the DESCRIBE message with a description of the media using the SDP format. The client then sends the SETUP message, which specifies the transport mechanisms acceptable to the client for streamed media, for example RTP/RTCP or RDT, and the ports on which it will receive the media. When using NAT, the RTSP ALG keeps track of these ports and translates them as necessary. The server responds to the SETUP method and selects one of the transport protocols, and, in this way, both client and server agree on a mechanism for media transport. The client then sends the PLAY method, and the server begins streaming the media to the client.

Figure 49: Typical RTSP Session



RTSP Request Methods

Table 23 on page 132 lists methods that can be performed on a resource (media object), the direction or directions in which information flows, and whether the method is required, recommended, or optional. Presentation refers to information such as network addresses, encoding, and content about a set of one or more streams presented to the client as a complete media feed. A Stream is a single media instance, for example audio or video, as well as all packets created by a source within the session.

Table 23: RTSP Request Methods (page 1 of 2)

Method	Direction	Object	Requirement	Description
OPTIONS	Client to Server	Presentation, Stream	Client to Server required	Client queries the server about what audio or video features it supports, as well as such things as the name and version of the server, and session ID.
	Server to Client	Presentation, Stream	Server to Client optional	
DESCRIBE	Client to Server	Presentation, Stream	Recommended	For exchange of media initialization information, such as clock rates, color tables, and any transport-independent information the client needs for playback of the media stream. Typically the client sends the URL of the of file it is requesting, and the server responds with a description of the media in SDP format.
ANNOUNCE	Client to Server	Presentation, Stream	Optional	Client uses this method to post a description of the presentation or media object identified by the request URL. The server uses this method to update the session description in real-time.
	Server to Client	Presentation, Stream		
SETUP	Client to Server	Stream	Required	Client specifies acceptable transport mechanisms to be used, such as the ports on which it will receive the media stream, and the transport protocol.
GET_PARAMETER	Client to Server	Presentation, Stream	Optional	Retrieves the value of a presentation or stream parameter specified in the URL. This method can be used with no entity body to test client or server aliveness. Ping can also be used to test for aliveness.
	Server to Client			
SET_PARAMETER	Client to Server	Presentation, Stream	Optional	Client uses this method to set the value of a parameter for a presentation or stream specified by the URI. Due to firewall considerations, this method cannot be used to set transport parameters.
	Server to Client			
PLAY	Client to Server	Presentation, Stream	Required	Instructs the server to begin sending data using the mechanism specified in SETUP. The Client does not issue PLAY requests until all SETUP requests are successful. The server queues PLAY requests in order, and delays executing any new PLAY request until an active PLAY request is completed. PLAY requests may or may not contain a specified range. The range may contain a time parameter—specified in Coordinated Universal Time (UTC)—for start of playback, which can also be used to synchronize streams from different sources.
PAUSE	Client to Server	Presentation, Stream	Recommended	Temporarily halts delivery of an active presentation. If the request URL specifies a particular stream, for example audio, this is equivalent to muting. Synchronization of tracks is maintained when playback or recording is resumed, although servers may close the session if PAUSE is for the duration specified in the timeout parameter in SETUP. A PAUSE request discards all queued PLAY requests.

Table 23: RTSP Request Methods (page 2 of 2)

Method	Direction	Object	Requirement	Description
RECORD	Client to Server	Presentation, Stream	Optional	Initiates recording a range of media defined in the presentation description. A UTC timestamp indicates start and end times, otherwise the server uses the start and end times in the presentation description.
REDIRECT	Server to Client	Presentation, Stream	Optional	Informs the client it must connect to a different server, and contains location information and possibly a range parameter for that new URL. To continue to receive media for this URL, the client must issue a TEARDOWN request for the current session and a SETUP for the new session.
TEARDOWN	Client to Server	Presentation, Stream	Required	Stops stream delivery for the given URI and frees the resources associated with it. Unless all transport parameters are defined by the session description, a SETUP request must be issued before the session can be played again.

RTSP Status Codes

RTSP uses status codes to provide information about client and server requests. Status codes include a machine-readable three digit result code, and a human-readable reason phrase. It is at the client's discretion whether to display the reason phrase. Status codes are described in Table 24.

Table 24: RSTP Status Codes

Code	Number	Description
Informational	100 to 199	Request has been received and is being processed
Success	200 to 299	Action has been received successfully, understood, and accepted
Redirection	300 to 399	Further action is necessary to complete the request
Client Error	400 to 499	Request contains bad syntax and cannot be fulfilled
Server Error	500 to 599	Server failed to fulfill an apparently valid request

Table 25 lists all status codes defined for RTSP 1.0, and recommended reason phrases. Reason phrases can be revised or redefined without affecting the operation of the protocol.

Table 25: RTSP 1.0 Status Codes (page 1 of 2)

Status Code	Reason Phrase	Status Code	Reason Phrase
100	Continue	414	Request-URI Too Large
200	OK	415	Unsupported Media Type
201	Created	451	Unsupported Media Type
250	Low on Storage Space	452	Conference Not Found
300	Multiple Choices	453	Not Enough Bandwidth
301	Moved Permanently	454	Session Not Found
303	See Other	455	Method Not Valid in This State

Table 25: RTSP 1.0 Status Codes (page 2 of 2)

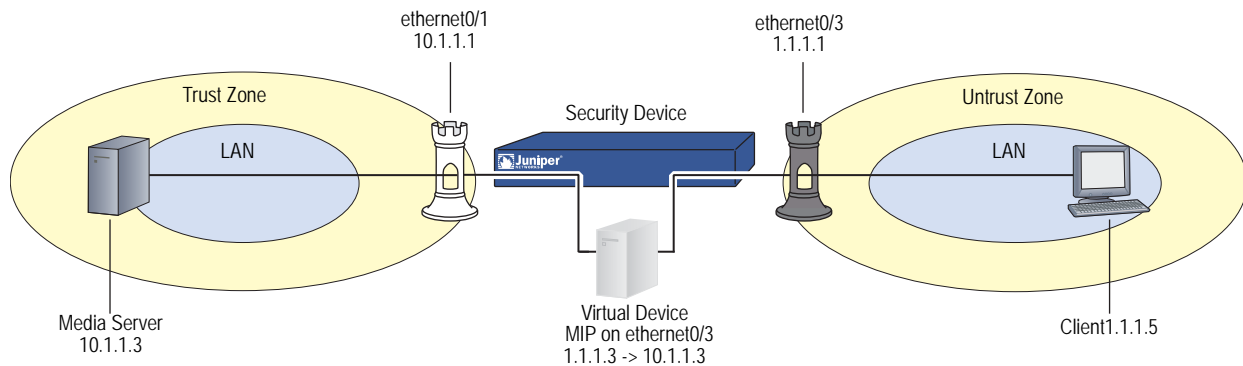
Status Code	Reason Phrase	Status Code	Reason Phrase
304	Not Modified	456	Header Field Not Valid for Resource
305	Use Proxy	457	Invalid Range
400	Bad Request	458	Parameter is Read-Only
401	Unauthorized	459	Aggregate operation not allowed
402	Payment Required	460	Only aggregate operation allowed
403	Forbidden	461	Unsupported transport
404	Not Found	462	Destination unreachable
405	Method Not Allowed	500	Internal Server Error
406	Not Acceptable	501	Not Implemented
407	Proxy Authentication Required	502	Bad Gateway
408	Request Time-out	503	Service Unavailable
410	Gone	504	Gateway Time-out
411	Length Required	505	RTSP Version not supported
412	Precondition Failed	551	Option not supported
413	Request Entity Too Large		

NOTE: For complete definitions of status codes, refer to RFC 2326, *Real Time Streaming Protocol (RTSP)*.

Configuring a Media Server in a Private Domain

In this example, the media server is in the Trust zone and the client is in the Untrust zone. You put a MIP on the ethernet0/3 interface to the media server in the Trust zone, then create a policy to allow RTSP traffic to flow from the client in the Untrust zone to the media server in the Trust zone.

Figure 50: RTSP Private Domain



WebUI**1. Interfaces**

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click

Apply:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24
 Manage IP: 10.1.1.2

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click

Apply:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24
 Manage IP: 1.1.1.2

2. Addresses

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: media_server
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.1.3/24
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: client
 IP Address/Domain Name:
 IP/Netmask: (select), 1.1.1.5/24
 Zone: Untrust

3. MIP

Network > Interfaces > Edit (for ethernet0/3) > MIP > New: Enter the following, then click **OK**:

Mapped IP: 1.1.1.3
 Host IP Address: 10.1.1.5

4. Policy

Policy > Policies > (From: Untrust, To: Trust) > New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), client
 Destination Address:
 Address Book Entry: (select), MIP(1.1.1.3)
 Service: RTSP
 Action: Permit

CLI

1. Interfaces

```
set interface ethernet0/1 trust
set interface ethernet0/1 ip 10.1.1.1
set interface ethernet0/3 untrust
set interface ethernet0/3 ip 1.1.1.1
```

2. Addresses

```
set address trust media_server 10.1.1.3/24
set address untrust client 1.1.1.5
```

3. MIP

```
set interface ethernet0/3 mip (1.1.1.3) host 10.1.1.3
```

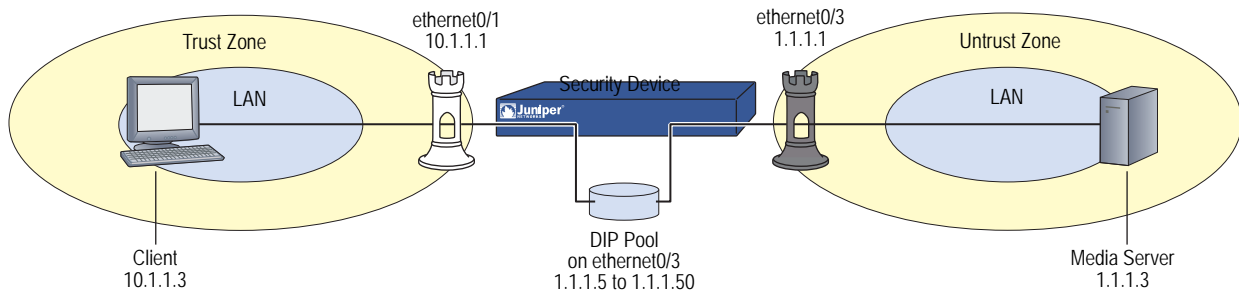
4. Policy

```
set policy from untrust to trust client mip(1.1.1.3) rtsp permit
save
```

Configuring a Media Server in a Public Domain

In this example, the media server is in the Untrust zone and the client is in the Trust zone. You put a DIP pool on the ethernet0/3 interface to do NAT when the media server responds to the client from the Untrust zone, then create a policy to allow RTSP traffic to flow from the Trust to the Untrust zone.

Figure 51: RTSP Public Domain



WebUI

1. Interface

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

```
Zone Name: Trust
Static IP: (select this option when present)
IP Address/Netmask: 10.1.1.1/24
Manage IP: 10.1.1.2
```

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **Apply**:

```
Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 1.1.1.1/24
Manage IP: 1.1.1.2
```

2. Addresses

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: client
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.1.3/24
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: media_server
 IP Address/Domain Name:
 IP/Netmask: (select), 1.1.1.3/24
 Zone: Untrust

3. DIP Pool

Network > Interfaces > Edit (for ethernet0/3) > DIP > New: Enter the following, then click **OK**:

ID: 5
 IP Address Range: (select) 1.1.1.5 ~ 1.1.1.50
 Port Translation: (select)

4. Policy

Policy > Policies > (From: Trust, To: Untrust) > New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry (select): client
 Destination Address:
 Address Book Entry (select): media_server
 Service: RTSP
 Action: Permit

> Advanced: Enter the following, then click **OK**:

NAT:
 Source Translation: (select)
 (DIP on): 5 (1.1.1.5-1.1.1.50)/port-xlate

CLI**1. Interface**

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust client ip 10.1.1.3/24
set address untrust media_server ip 1.1.1.3/24
```

3. DIP Pool

```
set interface ethernet0/3 dip 5 1.1.1.5 1.1.1.50
```

4. Policy

```
set policy from trust to untrust client media_server rtsp nat dip 5 permit
save
```

Service Groups

A service group is a set of services that you have gathered together under one name. After you create a group containing several services, you can then apply services at the group level to policies, thus simplifying administration.

The ScreenOS service group option has the following features:

- Each service book entry can be referenced by one or more service groups.
- Each service group can contain predefined and user-defined service book entries.

Service groups are subject to the following limitations:

- Service groups cannot have the same names as services; therefore, if you have a service named “FTP,” you cannot have a service group named “FTP.”
- If a service group is referenced in a policy, you can edit the group but you cannot remove it until you have first removed the reference to it in the policy.
- If a custom service book entry is deleted from the service book, the entry is also removed from all the groups in which it was referenced.
- One service group cannot contain another service group as a member.
- The all-inclusive service term “ANY” cannot be added to groups.
- A service can be part of only one group at a time.

Creating a Service Group

In this example, you create a service group named grp1 that includes IKE, FTP, and LDAP services.

WebUI

Policy > Policy Elements > Services > Groups > New: Enter the following group name, move the following services, then click **OK**:

Group Name: grp1

Select **IKE** and use the << button to move the service from the Available Members column to the Group Members column.

Select **FTP** and use the << button to move the service from the Available Members column to the Group Members column.

Select **LDAP** and use the << button to move the service from the Available Members column to the Group Members column.

CLI

```
set group service grp1
set group service grp1 add ike
set group service grp1 add ftp
set group service grp1 add ldap
save
```

NOTE: If you try to add a service to a service group that does not exist, the security device creates the group. Also, ensure that groups referencing other groups do not include themselves in the reference list.

Modifying a Service Group

In this example, you change the members in the service group named `grp1` that you created in “Creating a Service Group” on page 138. You remove IKE, FTP, and LDAP services, and add HTTP, FINGER, and IMAP.

WebUI

Policy > Policy Elements > Services > Groups > Edit (for `grp1`): Move the following services, then click **OK**:

Select **IKE** and use the > > button to move the service from the Group Members column to the Available Members column.

Select **FTP** and use the > > button to move the service from the Group Members column to the Available Members column.

Select **LDAP** and use the > > button to move the service from the Group Members column to the Available Members column.

Select **HTTP** and use the < < button to move the service from the Available Members column to the Group Members column.

Select **Finger** and use the < < button to move the service from the Available Members column to the Group Members column.

Select **IMAP** and use the < < button to move the service from the Available Members column to the Group Members column.

CLI

```
unset group service grp1 clear
set group service grp1 add http
set group service grp1 add finger
set group service grp1 add imap
save
```

Removing a Service Group

In this example, you delete the service group named “grp1”.

WebUI

Policy > Policy Elements > Services > Groups: Click **Remove** (for grp1).

CLI

```
unset group service grp1
save
```

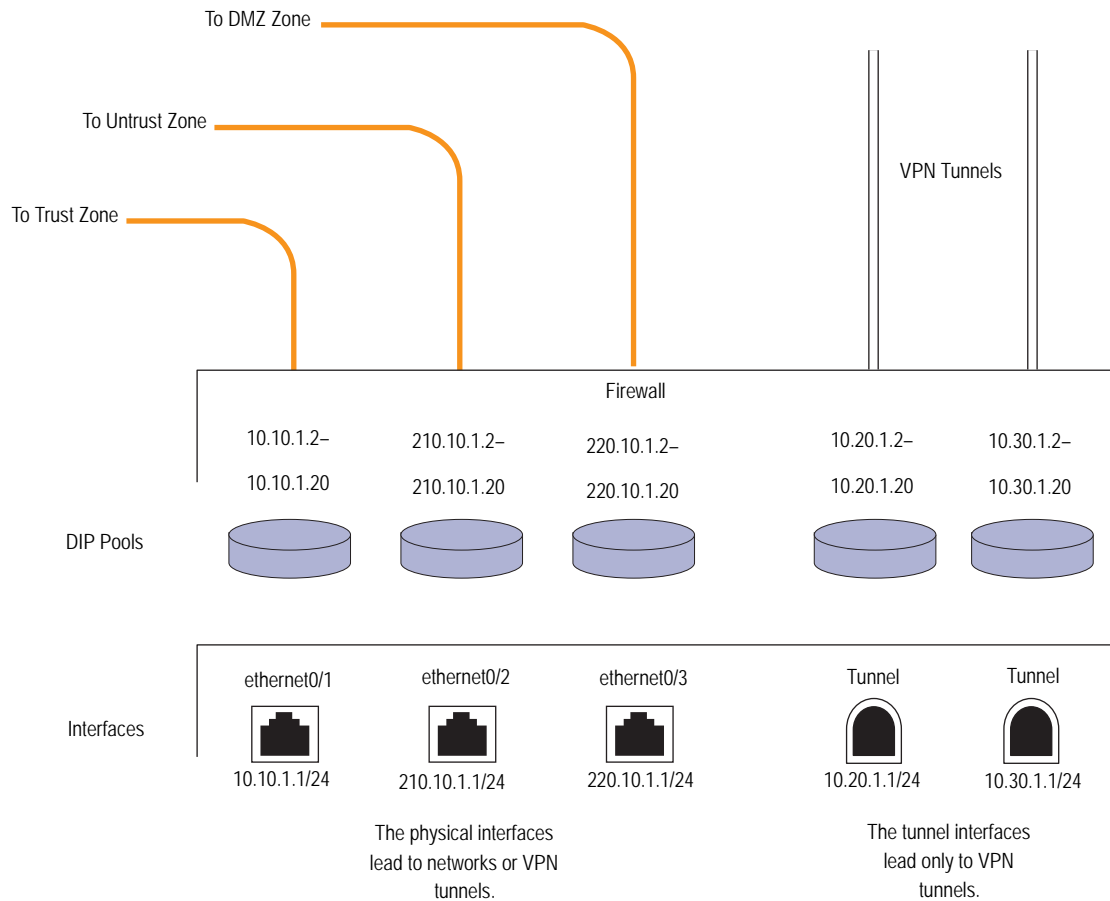
NOTE: The security device does not automatically delete a group from which you have removed all members.

Dynamic IP Pools

A dynamic IP (DIP) pool is a range of IP addresses from which the security device can dynamically or deterministically take addresses to use when performing Network Address Translation on the source IP address (NAT-src) in IP packet headers. (For information about deterministic source address translation, see “NAT-Src from a DIP Pool with Address Shifting” on page 8-20.) If the range of addresses in a DIP pool is in the same subnet as the interface IP address, the pool must exclude the interface IP address, router IP addresses, and any mapped IP (MIP) or virtual IP (VIP) addresses that might also be in that subnet. If the range of addresses is in the subnet of an extended interface, the pool must exclude the extended interface IP address.

There are three kinds of interfaces that you can link to dynamic IP (DIP) pools: physical interfaces and subinterfaces for network and VPN traffic, and tunnel interfaces for VPN tunnels only.

Figure 52: DIP Interfaces



Port Address Translation

Using Port Address Translation (PAT), multiple hosts can share the same IP address, the security device maintaining a list of assigned port numbers to distinguish which session belongs to which host. With PAT enabled, up to ~ 64,500 hosts can share a single IP address.

Some applications, such as NetBIOS Extended User Interface (NetBEUI) and Windows Internet Naming Service (WINS), require specific port numbers and cannot function properly if PAT is applied to them. For such applications, you can specify not to perform PAT (that is, to use a fixed port) when applying DIP. For fixed-port DIP, the security device hashes the original host IP address and saves it in its host hash table, thus allowing the security device to associate the right session with each host.

Creating a DIP Pool with PAT

In this example, you want to create a VPN tunnel for users at the local site to reach an FTP server at a remote site. However, the internal networks at both sites use the same private address space of 10.1.1.0/24. To solve the problem of overlapping addresses, you create a tunnel interface in the Untrust zone on the local security device, assign it IP address 10.10.1.1/24, and associate it with a DIP pool with a range of one address (10.10.1.2–10.10.1.2) and Port Address Translation enabled.

The admin at the remote site, must also create a tunnel interface with an IP address in a neutral address space, such as 10.20.2.1/24, and set up a mapped IP (MIP) address to its FTP server, such as 10.20.2.5 to host 10.1.1.5.

NOTE: This example includes only the configuration of the tunnel interface and its accompanying DIP pool. For a complete example showing all the configuration steps necessary for this scenario, see “VPN Sites with Overlapping Addresses” on page 5-140.

WebUI

Network > Interfaces > New Tunnel IF: Enter the following, then click **OK**:

```
Tunnel Interface Name: tunnel.1
Zone (VR): Untrust (trust-vr)
Fixed IP: (select)
IP Address / Netmask: 10.10.1.1/24
```

Network > Interfaces > Edit (for tunnel.1) > DIP > New: Enter the following, then click **OK**:

```
ID: 5
IP Address Range: 10.10.1.2 ~ 10.10.1.2
Port Translation: (select)
In the same subnet as the interface IP or its secondary IPs: (select)
```

NOTE: You can use the ID number displayed, which is the next available number sequentially, or type a different number.

CLI

```
set interface tunnel.1 zone untrust-tun
set interface tunnel.1 ip 10.10.1.1/24
set interface tunnel.1 dip 5 10.10.1.2 10.10.1.2
save
```

NOTE: Because PAT is enabled by default, there is no argument for enabling it. To create the same DIP pool as defined above but without PAT (that is, with fixed port numbers), do the following:

(WebUI) Network > Interfaces > Edit (for tunnel.1) > DIP > New: Clear the Port Translation checkbox, then click **OK**.

(CLI) set interface tunnel.1 dip 5 10.10.1.2 10.10.1.2 fix-port

Modifying a DIP Pool

In this example, you change the range of addresses in an existing DIP pool (ID 5) from 10.20.1.2 – 10.20.1.2 to 10.20.1.2 – 10.20.1.10. This DIP pool is associated with tunnel.1. Note that to change the DIP pool range through the CLI, you must first remove (or unset) the existing DIP pool and then create a new pool.

NOTE: There are no policies using this particular DIP pool. If a policy uses a DIP pool, you must first delete the policy or modify it to not use the DIP pool before you can modify the DIP pool.

WebUI

Network > Interfaces > Edit (for tunnel.1) > DIP > Edit (for ID 5): Enter the following, then click **OK**:

IP Address Range: 10.20.1.2 ~ 10.20.1.10

CLI

```
unset interface tunnel.1 dip 5
set interface tunnel.1 dip 5 10.20.1.2 10.20.1.10
save
```

Sticky DIP Addresses

When a host initiates several sessions that match a policy requiring Network Address Translation (NAT) and is assigned an address from a DIP pool with port translation enabled, the security device assigns a different source IP address for each session. Such random address assignment can be problematic for services that create multiple sessions that require the same source IP address for each session.

NOTE: For DIP pools that do not perform port translation, the security device assigns one IP address for all concurrent sessions from the same host.

For example, it is important to have the same IP address for multiple sessions when using the AOL Instant Messaging (AIM) client. You create one session when you log in, and another for each chat. For the AIM server to verify that a new chat belongs to an authenticated user, it must match the source IP address of the login session with that of the chat session. If they are different—possibly because they were randomly assigned from a DIP pool during the NAT process—the AIM server rejects the chat session.

To ensure that the security device assigns the same IP address from a DIP pool to a host for multiple concurrent sessions, you can enable the “sticky” DIP address feature by entering the CLI command **set dip sticky**.

Using DIP in a Different Subnet

If circumstances require that the source IP address in outbound firewall traffic be translated to an address in a different subnet from that of the egress interface, you can use the extended interface option. This option allows you to graft a second IP address and an accompanying DIP pool onto an interface that is in a different subnet. You can then enable NAT on a per-policy basis and specify the DIP pool built on the extended interface for the translation.

In this example, two branch offices have leased lines to a central office. The central office requires them to use only the authorized IP addresses it has assigned them. However, the offices receive different IP addresses from their ISPs for Internet traffic. For communication with the central office, you use the extended interface option to configure the security device in each branch office to translate the source IP address in packets it sends to the central office to the authorized address. The authorized and assigned IP addresses for branch offices A and B are as follows:

Table 26: Authorized Office IP Addresses

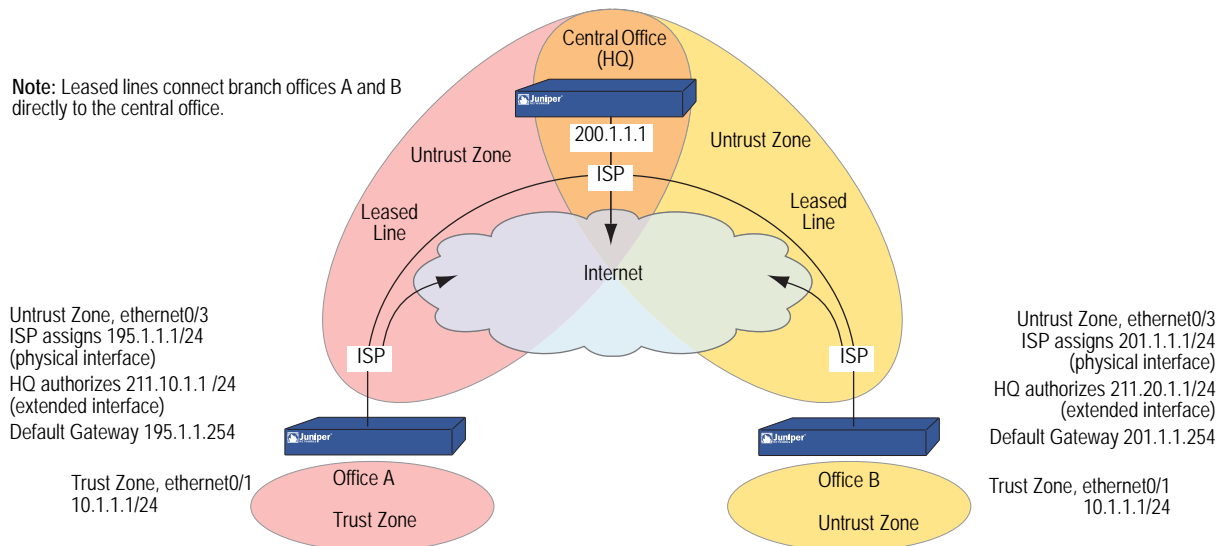
	Assigned IP Address (from ISP) Used for Untrust Zone Physical Interface	Authorized IP Address (from Central Office) Used for Untrust Zone Extended Interface DIP
Office A	195.1.1.1/24	211.10.1.1/24
Office B	201.1.1.1/24	211.20.1.1/24

The security devices at both sites have a Trust zone and an Untrust zone. All security zones are in the trust-vr routing domain. You bind ethernet0/1 to the Trust zone and assign it IP address 10.1.1.1/24. You bind ethernet0/3 to the Untrust zone and give it the IP address assigned by the ISPs: 195.1.1.1/24 for Office A and 201.1.1.1/24 for Office B. You then create an extended interface with a DIP pool containing the authorized IP address on ethernet0/3:

- Office A: extended interface IP 211.10.1.10/24; DIP pool 211.10.1.1 – 211.10.1.1; PAT enabled
- Office B: extended interface IP 211.20.1.10/24; DIP pool 211.20.1.1 – 211.20.1.1; PAT enabled

You set the Trust zone interface in NAT mode. It uses the Untrust zone interface IP address as its source address in all outbound traffic except for traffic sent to the central office. You configure a policy to the central office that translates the source address to an address in the DIP pool in the extended interface. (The DIP pool ID number is 5. It contains one IP address, which, with Port Address Translation (PAT), can handle sessions for ~ 64,500 hosts.) The MIP address that the central office uses for inbound traffic is 200.1.1.1, which you enter as “HQ” in the Untrust zone address book on each security device.

Figure 53: DIP Under Another Subnet



NOTE: Each ISP must set up a route for traffic destined to a site at the end of a leased line to use that leased line. The ISPs route any other traffic they receive from a local security device to the Internet.

WebUI (Branch Office A)

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Zone Name: Trust
Static IP: (select this option when present)
IP Address/Netmask: 10.1.1.1/24
Interface Mode: NAT

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 195.1.1.1/24
Interface Mode: Route

Network > Interfaces > Edit (for ethernet0/3) > DIP > New: Enter the following, then click **OK**:

ID: 5
IP Address Range: 211.10.1.1 ~ 211.10.1.1
Port Translation: (select)
Extended IP/Netmask: 211.10.1.10/255.255.255.0

2. Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: HQ
IP Address/Domain Name:
IP/Netmask: (select), 200.1.1.1/32
Zone: Untrust

3. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet0/3
Gateway IP address: 195.1.1.254

4. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: ANY
Action: Permit

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), HQ
Service: ANY
Action: Permit
Position at Top: (select)

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

NAT:
Source Translation: (select)
(DIP on): 5 (211.10.1.1-211.10.1.1)/X-late

WebUI (Branch Office B)**1. Interfaces**

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24
 Interface Mode: NAT

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 201.1.1.1/24
 Interface Mode: Route

Network > Interfaces > Edit (for ethernet0/3) > DIP > New: Enter the following, then click **OK**:

ID: 5
 IP Address Range: 211.20.1.1 ~ 211.20.1.1
 Port Translation: (select)
 Extended IP/Netmask: 211.20.1.10/255.255.255.0

2. Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: HQ
 IP Address/Domain Name:
 IP/Netmask: (select), 200.1.1.1/32
 Zone: Untrust

3. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/3
 Gateway IP address: 201.1.1.254

4. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: ANY
 Action: Permit

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), HQ
 Service: ANY
 Action: Permit
 Position at Top: (select)

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

NAT:
 Source Translation: (select)
 DIP On: (select), 5 (211.20.1.1-211.20.1.1)/X-late

CLI (Branch Office A)

1. Interfaces

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/1 nat
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 195.1.1.1/24
set interface ethernet0/3 route
set interface ethernet0/3 ext ip 211.10.1.10 255.255.255.0 dip 5 211.10.1.1
```

2. Address

```
set address untrust hq 200.1.1.1/32
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 195.1.1.254
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy top from trust to untrust any hq any nat src dip 5 permit
save
```

CLI (Branch Office B)

1. Interfaces

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/1 nat
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 201.1.1.1/24
set interface ethernet0/3 route
set interface ethernet0/3 ext ip 211.20.1.10 255.255.255.0 dip 5 211.20.1.1
```

2. Address

```
set address untrust hq 200.1.1.1/32
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 201.1.1.254
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy top from trust to untrust any hq any nat src dip 5 permit
save
```

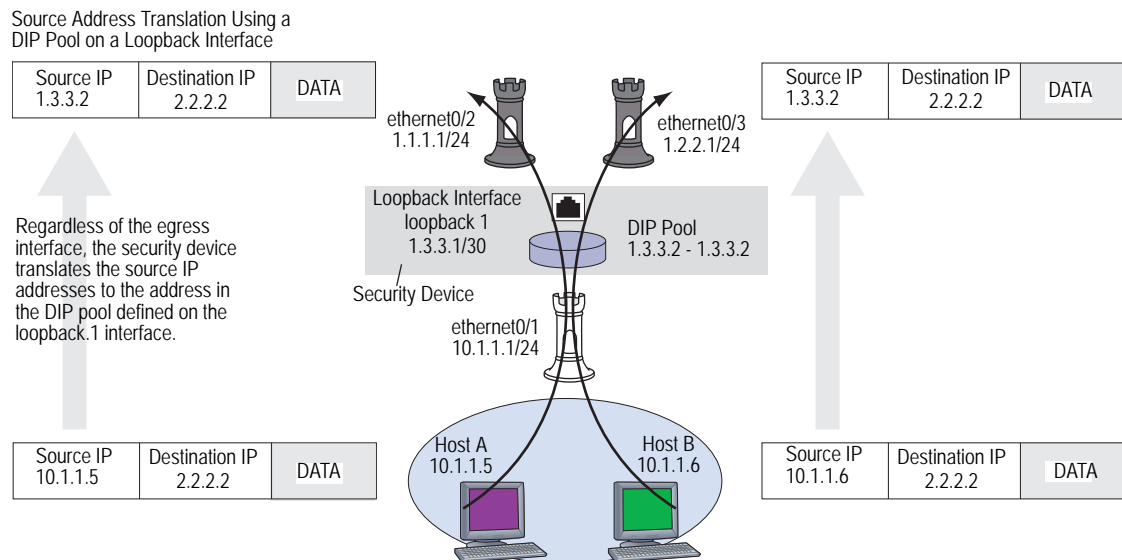
Using a DIP on a Loopback Interface

A loopback interface is a logical interface that is always in the up state as long as the device on which it resides is up. You can create a pool of dynamic IP (DIP) addresses on a loopback interface so that it can be accessed by the group of interfaces belonging to its associated loopback interface group when performing source address translation. The addresses that the security device draws from such a DIP pool are in the same subnet as the loopback interface IP address, not in the subnet of any of the member interfaces. (Note that the addresses in the DIP pool must not overlap with the interface IP address or any MIP addresses also defined on the loopback interface.)

NOTE: For information about loopback interfaces, see “Loopback Interfaces” on page 58.

The primary application for putting a DIP pool on a loopback interface is to translate source addresses to the same address or range of addresses although different packets might use different egress interfaces.

Figure 54: Loopback DIP



In this example, the security device receives the following IP addresses for two Untrust zone interfaces from different Internet service providers (ISPs): ISP-1 and ISP-2:

- ethernet0/2, 1.1.1.1/24, ISP-1
- ethernet0/3, 1.2.2.1/24, ISP-2

You bind these interfaces to the Untrust zone and then assign them the above IP addresses. You also bind ethernet0/1 to the Trust zone and assign it IP address 10.1.1.1/24.

You want the security device to translate the source address in outbound traffic from the Trust zone to a remote office in the Untrust zone. The translated address must be the same IP address (1.3.3.2) because the remote office has a policy permitting inbound traffic only from that IP address. You have previously obtained the public IP addresses 1.3.3.1 and 1.3.3.2 and have notified both ISPs that you are using these addresses in addition to the addresses that they assign the device.

You configure a loopback interface loopback.1 with the IP address 1.3.3.1/30 and a DIP pool of 1.3.3.2 – 1.3.3.2 on that interface. The DIP pool has ID number 10. You then make ethernet0/1 and ethernet0/2 members of the loopback group for loopback.1.


You define an address for the remote office named “r-office” with IP address 2.2.2.2/32. You also define default routes for both ethernet0/1 and ethernet0/2 interfaces pointing to the routers for ISP-1 and ISP-2, respectively.

You define routes to two gateways for outbound traffic to use. Because you do not prefer one route over the other, you do not include any metrics in the routes. Outbound traffic might follow either route.

NOTE: To indicate a route preference, include metrics in both routes, giving your preferred route a higher metric—that is, a value closer to 1.

Finally, you create a policy applying Source Network Address Translation (NAT-src) to outbound traffic to the remote office. The policy references DIP pool ID 10.

Figure 55: Loopback DIP Policy



WebUI**1. Interfaces**

Network > Interfaces > New Loopback IF: Enter the following, then click **OK**:

Interface Name: loopback.1
 Zone: Untrust (trust-vr)
 IP Address/Netmask: 1.3.3.1/30

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

As member of loopback group: loopback.1
 Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24
 Interface Mode: NAT

Network > Interfaces > Edit (for ethernet0/2): Enter the following, then click **OK**:

As member of loopback group: loopback.1
 Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24
 Interface Mode: Route

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.2.2.1/24
 Interface Mode: Route

2. DIP Pool

Network > Interfaces > Edit (for loopback.1) > DIP > New: Enter the following, then click **OK**:

ID: 5
 IP Address Range: 1.3.3.2 ~ 1.3.3.2
 Port Translation: (select)

3. Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: r-office
 IP Address/Domain Name:
 IP/Netmask: (select), 2.2.2.2/32
 Zone: Untrust

4. Routes

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/2
 Gateway IP address: 1.1.1.250

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/3
 Gateway IP address: 1.2.2.250

5. Policy

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), r-office
 Service: ANY
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

NAT:
 Source Translation: (select)
 DIP On: (select), 10 (1.3.3.2-1.3.3.2)/port-xlate

CLI

1. Interfaces

```
set interface loopback.1 zone untrust
set interface loopback.1 ip 1.3.3.1/30
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/1 nat
set interface ethernet0/2 zone untrust
set interface ethernet0/2 ip 1.1.1.1/24
set interface ethernet0/2 loopback-group loopback.1
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 1.2.2.1/24
set interface ethernet0/3 loopback-group loopback.1
```

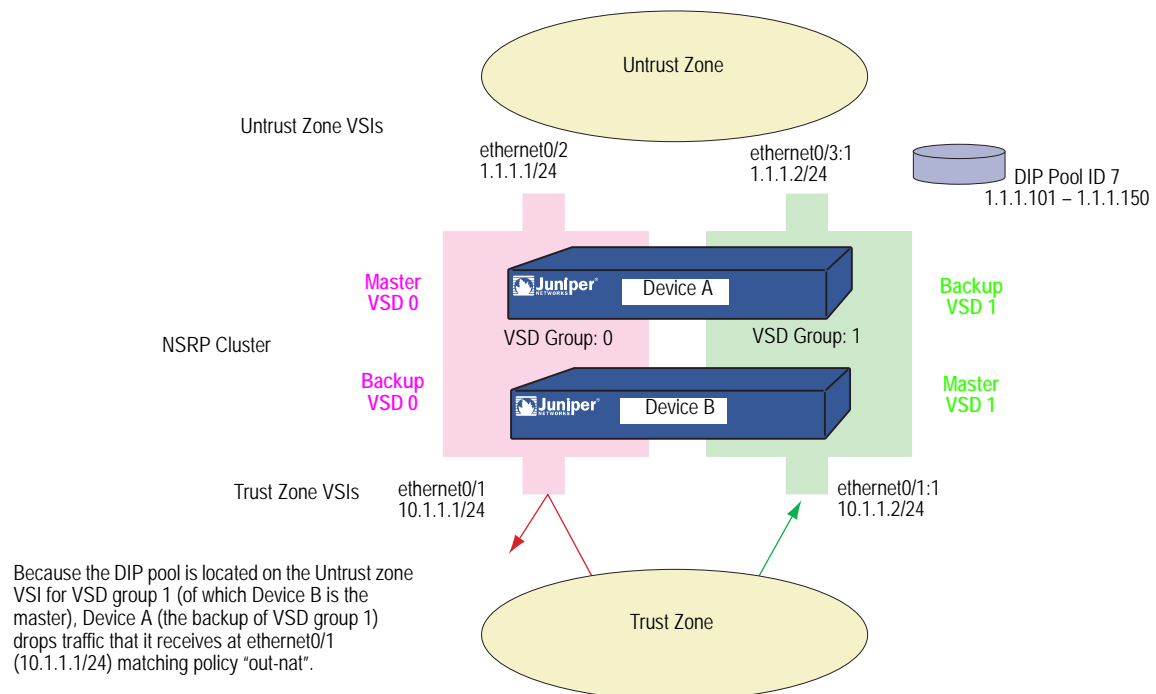
2. **DIP Pool**
set interface loopback.1 dip 10 1.3.3.2 1.3.3.2
3. **Address**
set address untrust r-office 2.2.2.2/32
4. **Routes**
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/2 gateway 1.1.1.250
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 1.2.2.250
5. **Policy**
set policy from trust to untrust any r-office any nat src dip-id 10 permit
save

Creating a DIP Group

When you group two security devices into a redundant cluster to provide high availability (HA) services in an active/active configuration, both devices share the same configuration and both process traffic simultaneously. A problem can arise when you define a policy to perform Network Address Translation (NAT) using a dynamic IP (DIP) pool located on one VSI. Because that VSI is active only on the security device acting as the primary of the VSD group to which the VSI is bound, any traffic sent to the other security device—the one acting as the backup of that VSD group—cannot use that DIP pool and is dropped.

Figure 56: DIP Problems with NAT with One VSI

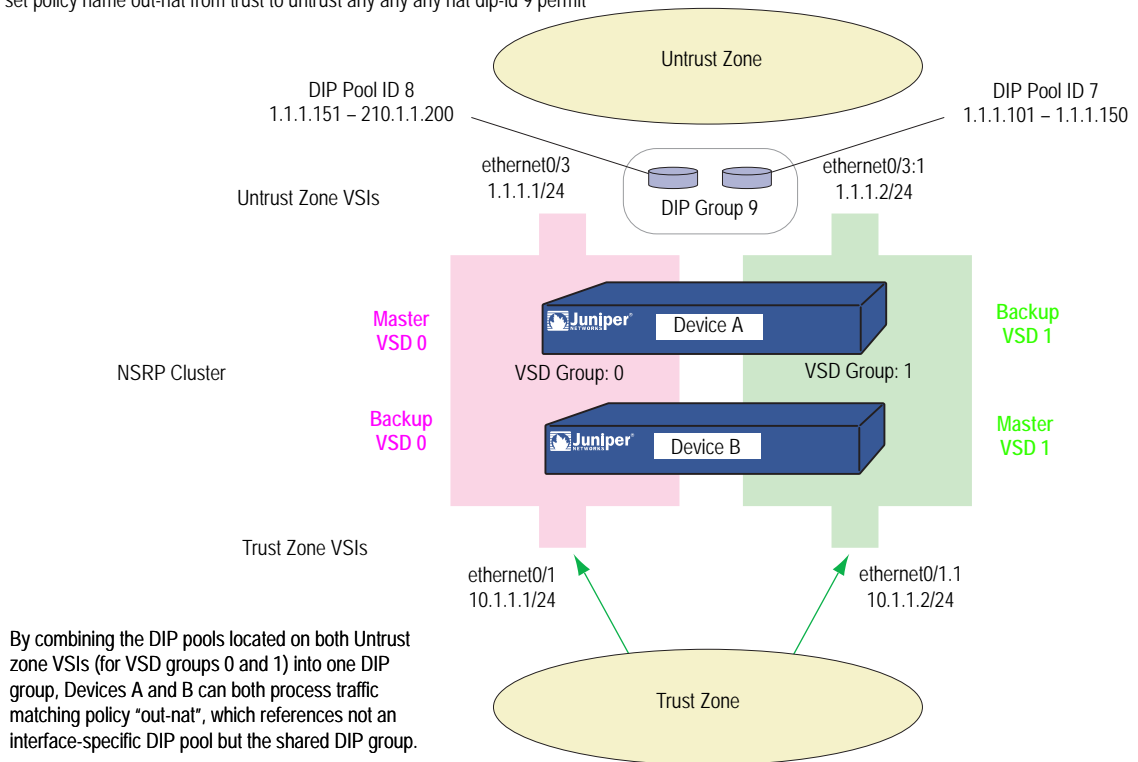
Problematic use of a DIP pool in a policy when in an NSRP cluster:
set policy name out-nat from trust to untrust any any nat src dip-id 7 permit



To solve this problem, you can create two DIP pools—one on the Untrust zone VSI for each VSD group—and combine the two DIP pools into one DIP group, which you reference in the policy. Each VSI uses its own VSD pool even though the policy specifies the DIP group.

Figure 57: Creating Two DIP Pools in One DIP Group

Recommended use of a DIP group in a policy when in an NSRP cluster:
 set policy name out-nat from trust to untrust any any nat dip-id 9 permit



NOTE: For more information about setting up security devices for HA, see *Volume 11: High Availability*.

In this example, you provide NAT services on two security devices (Devices A and B) in an active/active HA pair.

You create two DIP pools—DIP 5 (1.1.1.20 – 1.1.1.29) on ethernet0/3 and DIP 6 (1.1.1.30 – 1.1.1.39) on ethernet0/3:1. You then combine them into a DIP group identified as DIP 7, which you reference in a policy.

The VSIs for VSD groups 0 and 1 are as follows:

- Untrust zone VSI ethernet0/3 1.1.1.1/24 (VSD group 0)
- Untrust zone VSI ethernet0/3:1 1.1.1.2/24 (VSD group 1)
- Trust zone VSI ethernet0/1 10.1.1.1/24 (VSD group 0)
- Trust zone VSI ethernet0/1:1 10.1.1.2/24 (VSD group 1)

Let's assume that you have already set up Devices A and B in an NSRP cluster, created VSD group 1 (ScreenOS automatically creates VSD group 0 when you put a device in an NSRP cluster), and configured the above interfaces. (For information about configuring security devices for NSRP, see *Volume 11: High Availability*.)

WebUI

1. DIP Pools

Network > Interfaces > Edit (for ethernet0/3) > DIP > New: Enter the following, then click **OK**:

ID: 5
 IP Address Range: 1.1.1.20 – 1.1.1.29
 Port Translation: (select)

Network > Interfaces > Edit (for ethernet0/3:1) > DIP > New: Enter the following, then click **OK**:

ID: 6
 IP Address Range: 1.1.1.30 – 1.1.1.39
 Port Translation: (select)

NOTE: At the time of this release, you can only define a DIP group through the CLI.

2. Policy

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Any
 Service: ANY
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

NAT:
 Source Translation: (select)
 DIP On: (select), 7

CLI

1. DIP Pools

```
set interface ethernet0/3 dip 5 1.1.1.20 1.1.1.29
set interface ethernet0/3:1 dip 6 1.1.1.30 1.1.1.39
```

2. DIP Groups

```
set dip group 7 member 5
set dip group 7 member 6
```

3. Policy

```
set policy from trust to untrust any any any nat src dip-id 7 permit
save
```

Setting a Recurring Schedule

A schedule is a configurable object that you can associate with one or more policies to define when they are in effect. Through the application of schedules, you can control network traffic flow and enforce network security.

When you define a schedule, enter values for the following parameters:

- **Schedule Name:** The name that appears in the Schedule drop-down list in the Policy Configuration dialog box. Choose a descriptive name to help you identify the schedule. The name must be unique and is limited to 19 characters.
- **Comment:** Any additional information that you want to add.
- **Recurring:** Enable this when you want the schedule to repeat on a weekly basis.

Start and End Times: You must configure both a start time and an end time. You can specify up to two time periods within the same day.

- **Once:** Enable this when you want the schedule to start and end only once.

mm/dd/yyyy hh:mm: You must enter both start and stop dates and times.

In this example, there is a short-term employee named Tom who is using the company's Internet access for personal pursuits after work. You create a schedule for non-business hours that you can then associate with a policy to deny outbound TCP/IP traffic from that worker's computer (10.1.1.5/32) outside of regular business hours.

WebUI

1. Schedule

Policy > Policy Elements > Schedules > New: Enter the following, then click **OK**:

Schedule Name: After Hours
 Comment: For non-business hours
 Recurring: (select)
 Period 1:

Weekday	Start Time	End Time
Sunday	00:00	23:59
Monday	00:00	06:00
Tuesday	00:00	06:00
Wednesday	00:00	06:00
Thursday	00:00	06:00
Friday	00:00	06:00
Saturday	00:00	23:59

Period 2:

Weekday	Start Time	End Time
Sunday	17:00	23:59
Monday	17:00	23:59
Tuesday	17:00	23:59
Wednesday	17:00	23:59
Thursday	17:00	23:59
Friday	17:00	23:59
Saturday	17:00	23:59

2. Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: Tom
 Comment: Temp
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.1.5/32
 Zone: Trust

3. Policy

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: No Net
 Source Address:
 Address Book Entry: (select), Tom
 Destination Address:
 Address Book Entry: (select), Any
 Service: HTTP
 Action: Deny
 Schedule: After Hours

CLI

1. Schedule

```
set schedule "after hours" recurrent sunday start 00:00 stop 23:59
set schedule "after hours" recurrent monday start 00:00 stop 06:00 start 17:00
stop 23:59
set schedule "after hours" recurrent tuesday start 00:00 stop 06:00 start 17:00
stop 23:59
set schedule "after hours" recurrent wednesday start 00:00 stop 06:00 start
17:00 stop 23:59
set schedule "after hours" recurrent thursday start 00:00 stop 06:00 start 17:00
stop 23:59
set schedule "after hours" recurrent friday start 00:00 stop 06:00 start 17:00
stop 23:59
set schedule "after hours" recurrent saturday start 00:00 stop 23:59 comment
"for non-business hours"
```

2. Address

```
set address trust tom 10.1.1.5/32 "temp"
```

3. Policy

```
set policy from trust to untrust tom any http deny schedule "after hours"
save
```


Chapter 6

Policies

The default behavior of a security device is to deny all traffic between security zones (interzone traffic) and—except for traffic within the Untrust zone—allow all traffic between interfaces bound to the same zone (intrazone traffic). To permit selected interzone traffic to cross a security device you must create interzone policies that override the default behavior. Similarly, to prevent selected intrazone traffic from crossing a security device, you must create intrazone policies.

This chapter describes what policies do and how the various elements that comprise a policy are related. It contains the following sections:

- “Basic Elements” on page 160
- “Three Types of Policies” on page 161
- “Policy Set Lists” on page 163
- “Policies Defined” on page 164
- “Policies Applied” on page 174

NOTE: If you configure multicast routing on a security device, you might have to configure multicast policies. For information about multicast policies, see “Multicast Policies” on page 7-153.

Basic Elements

A policy permits, denies, or tunnels specified types of traffic unidirectionally between two points. The type of traffic (or “service”), the location of the two endpoints, and the invoked action compose the basic elements of a policy. Although there can be other components, the required elements, which together constitute the core section of a policy, are described in Table 27.

Table 27: Basic Policy Elements

Element	Description
Direction	The direction of traffic between two security zones (from a source zone to a destination zone)
Source Address	The address from which traffic initiates
Destination Address	The address to which traffic is sent
Service	The type of traffic transmitted
Action	The action that the security device performs when it receives traffic meeting the first four criteria: deny, permit, reject, or tunnel Note: The “tunnel” action—(VPN or L2TP tunnel)—contains the concept of “permit” implicitly.

For example, the policy stated in the following CLI command permits FTP traffic from any address in the Trust zone to an FTP server named “server1” in the DMZ zone:

```
set policy from trust to untrust any server1 ftp permit
```

- **Direction: from trust to untrust** (that is, from the Trust zone to the Untrust zone)
- **Source Address: any** (that is, any address in the Trust zone. The term “any” stands for a predefined address that applies to any address in a zone)
- **Destination Address: server1** (a user-defined address in the Untrust zone address book)
- **Service: ftp** (File Transfer Protocol)
- **Action: permit** (that security device permits this traffic to traverse its firewall)

Three Types of Policies


You control the flow of traffic with the following three types of policies:

- **Interzone Policies**—Let you regulate the kind of traffic allowed to pass from one security zone to another.
- **Intrazone Policies**— Let you regulate the kind of traffic allowed to cross interfaces bound to the same zone.
- **Global Policies**—Let you regulate traffic between addresses, regardless of their security zones.

Interzone Policies

Interzone policies provide traffic control between security zones. You can set interzone policies to deny, permit, reject, or tunnel traffic from one zone to another. Using stateful inspection techniques, a security device maintains a table of active TCP sessions and active UDP “pseudo” sessions so that it can allow replies to service requests. For example, if you have a policy allowing HTTP requests from host A in the Trust zone to server B in the Untrust zone, when the security device receives HTTP replies from server B to host A, the security device checks the received packet against its table. Finding the packet to be a reply to an approved HTTP request, the security device allows the packet from server B in the Untrust zone to cross the firewall to host A in the Trust zone. To permit traffic initiated by server B to host A (not just replies to traffic initiated by host A), you must create a second policy from server B in the Untrust zone to host A in the Trust zone.

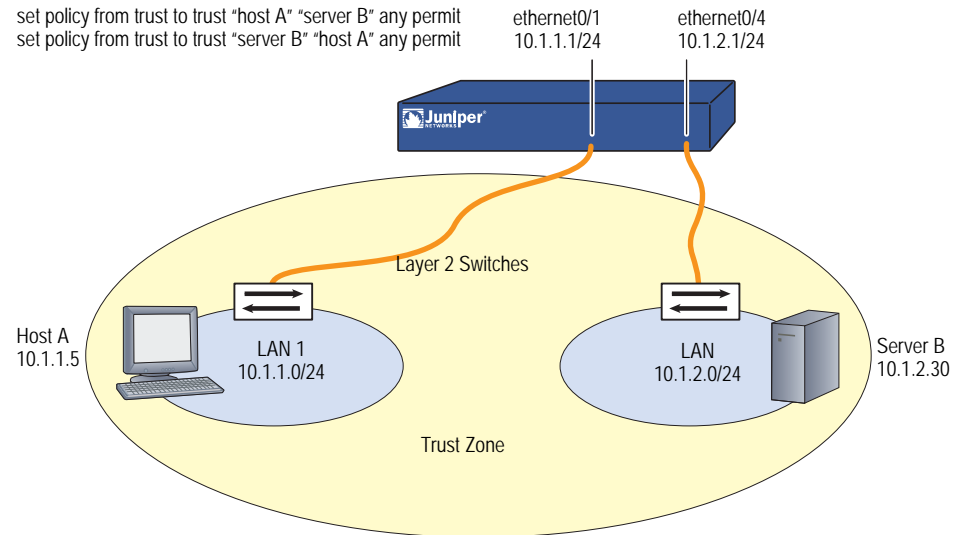
Figure 58: Interzone Policy



Intrazone Policies

Intrazone policies provide traffic control between interfaces bound to the same security zone. The source and destination addresses are in the same security zone but are reached via different interfaces on the security device. Like interzone policies, intrazone policies control traffic flowing unidirectionally. To allow traffic initiated at either end of a data path, you must create two policies—one policy for each direction.

Figure 59: Intrazone Policy



Intrazone policies do not support VPN tunnels or source network address translation (NAT-src) when it is set at the interface level (**set interface interface nat**). However, intrazone policies do support policy-based NAT-src and NAT-dst. They also support destination address translation when the policy references a mapped IP (MIP) as the destination address. (For information about NAT-src, NAT-dst, and MIPs, see *Volume 8: Address Translation*.)

Global Policies

Unlike interzone and intrazone policies, global policies do not reference specific source and destination zones. Global policies reference user-defined Global zone addresses or the predefined Global zone address “any”. These addresses can span multiple security zones. For example, if you want to provide access to or from multiple zones, you can create a global policy with the Global zone address “any,” which encompasses all addresses in all zones.

NOTE: At the time of this release, global policies do not support source network address translation (NAT-src), VPN tunnels, or Transparent mode. You can, however, specify a MIP or VIP as the destination address in a global policy.

Policy Set Lists

A security device maintains three different policy set lists, one each for interzone policies, intrazone policies, and global policies.

When the security device receives a packet initiating a new session, the device notes the ingress interface, and thereby learns the source zone to which that interface is bound. The security device then performs a route lookup to determine the egress interface, and thus determines the destination zone to which that interface is bound. Using the source and destination zones, the security device can perform a policy lookup, consulting the policy set lists in the following order:

1. If the source and destination zones are different, the security device performs a policy lookup in the interzone policy set list.

(or)

If the source and destination zones are the same, the security device performs a policy lookup in the intrazone policy set list.

2. If the security device performs the interzone or intrazone policy lookup and does not find a match, the security device then checks the global policy set list for a match.
3. If the security device performs the interzone and global policy lookups and does not find a match, the security device then applies the default permit/deny policy to the packet: **unset/set policy default-permit-all**.

(or)

If the security device performs the intrazone and global policy lookups and does not find a match, the security device then applies the intrazone blocking setting for that zone to the packet: **unset/set zone zone block**.

The security device searches each policy set list from top to bottom. Therefore, you must position more specific policies above less specific policies in the list. (For information about policy order, see “Reordering Policies” on page 190.)

Policies Defined

A firewall provides a network boundary with a single point of entry and exit. Because all traffic must pass through this point, you can screen and direct that traffic through the implementation of policy set lists—for interzone policies, intrazone policies, and global policies.

Policies allow you to deny, permit, reject (deny and send a TCP RST or ICMP port unreachable message to the source host), encrypt and decrypt, authenticate, prioritize, schedule, filter, and monitor the traffic attempting to cross from one security zone to another. You decide which users and what data can enter and exit, and when and where they can go.

NOTE: For security devices that support virtual systems, policies set in the root system do not affect policies set in virtual systems.

Policies and Rules

A single user-defined policy produces one or more logical rules internally, and each logical rule consists of a set of components—source address, destination address, and service. The components consume memory resources. The logical rules that reference the components do not.

Depending on the use of multiple entries or groups for the source address, destination address, and service components in a policy, the number of logical rules can be much larger than is readily apparent from the creation of the single policy. For example, the following policy produces 125 logical rules:

1 policy: 5 source addresses x 5 destination addresses x 5 services = 125
logical rules

However, the security device does not duplicate components for each logical rule. The rules make use of the same set of components in various combinations. For example, the above policy that produces 125 logical rules results in only 15 components:

5 source addresses + 5 destination addresses + 5 services = 15 components

These 15 components combine in various ways to produce the 125 logical rules generated by the single policy. By allowing multiple logical rules to use the same set of components in different combinations, the security device consumes far fewer resources than if each logical rule had a one-to-one relationship with its components.

Because the installation time of a new policy is proportional to the number of components that the security device adds, removes, or modifies, policy installation becomes faster with fewer components. Also, by allowing a large number of logical rules to share a small set of components, ScreenOS allows you to create more policies—and the security device to create more rules—than would be possible if each rule required dedicated components.

Anatomy of a Policy

A policy must contain the following elements:

- ID (automatically generated, but can be user-defined in the CLI)
- Zones (source and destination)
- Addresses (source and destination)
- Services
- Action (deny, permit, reject, tunnel)

A policy can also contain the following elements:

- Application
- Name
- VPN Tunneling
- L2TP Tunneling
- Deep Inspection
- Placement at the Top of the Policy List
- Source Address Translation
- Destination Address Translation
- User Authentication
- HA Session Backup
- Web Filtering
- Logging
- Counting
- Traffic Alarm Threshold
- Schedules
- Antivirus Scanning
- Traffic Shaping

The remainder of this section examines each of the above elements in turn.

ID

Every policy has an ID number, whether you define one or the security device automatically assigns it. You can only define an ID number for a policy through the `set policy` command in the CLI: **set policy id** *number* ... After you know the ID number, you can enter the policy context to issue further commands to modify the policy. (For more information about policy contexts, see “Entering a Policy Context” on page 185.)

Zones

A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or logical entity that performs a specific function (a function zone). A policy allows traffic to flow between two security zones (interzone policy) or between two interfaces bound to the same zone (intrazone policy). (For more information, see “Zones” on page 25, “Interzone Policies” on page 161, and “Intrazone Policies” on page 161.)

Addresses

Addresses are objects that identify network devices such as hosts and networks by their location in relation to the firewall—in one of the security zones. Individual hosts are specified using the mask 255.255.255.255, indicating that all 32 bits of the address are significant. Networks are specified using their subnet mask to indicate which bits are significant. To create a policy for specific addresses, you must first create entries for the relevant hosts and networks in the address book.

You can also create address groups and apply policies to them as you would to other address book entries. When using address groups as elements of policies, be aware that because the security device applies the policy to each address in the group, the number of available internal logical rules and the components that comprise those rules can become depleted more quickly than expected. This is a danger especially when you use address groups for both the source and destination. (For more information, see “Policies and Rules” on page 164.)

Services

Services are objects that identify application protocols using Layer 4 information such as standard and accepted TCP and UDP port numbers for application services like Telnet, FTP, SMTP, and HTTP. The ScreenOS includes predefined core Internet services. Additionally, you can define custom services.

You can define policies that specify which services are permitted, denied, encrypted, authenticated, logged, or counted.

Action

An action is an object that describes what the firewall does to the traffic it receives.

- **Deny** blocks the packet from traversing the firewall.
- **Permit** allows the packet to pass the firewall.
- **Reject** blocks the packet from traversing the firewall. The security device drops the packet and sends a TCP reset (RST) segment to the source host for TCP traffic and an ICMP “destination unreachable, port unreachable” message (type 3, code 3) for UDP traffic. For types of traffic other than TCP and UDP, the security device drops the packet without notifying the source host, which is also what occurs when the action is “deny.”

NOTE: The security device sends a TCP RST after receiving (and dropping) a TCP segment with any code bit set other than another RST.

When the ingress interface is operating at Layer 2 or 3 and the protocol is TCP, the source IP address in the TCP RST is the destination IP address in the original (dropped) packet. When the ingress interface is operating at Layer 2 and the protocol is UDP, the source IP address in the ICMP message is also the destination IP address in the original packet. However, if the ingress interface is operating at Layer 3 and the protocol is UDP, then the source IP address in the ICMP message is that of the ingress interface.

-
- **Tunnel** encapsulates outgoing IP packets and decapsulates incoming IP packets. For an IPSec VPN tunnel, specify which VPN tunnel to use. For an L2TP tunnel, specify which L2TP tunnel to use. For L2TP-over-IPSec, specify both an IPSec VPN tunnel and an L2TP tunnel.

NOTE: For L2TP-over-IPSec, the source and destination addresses for the IPSec VPN tunnel must be the same as those for the L2TP tunnel.

The security device applies the specified action on traffic that matches the previously presented criteria: zones (source and destination), addresses (source and destination), and service.

Application

The application option specifies the Layer 7 application that maps to the Layer 4 service that you reference in a policy. A predefined service already has a mapping to a Layer 7 application. However, for custom services, you must link the service to an application explicitly, especially if you want the policy to apply an Application Layer Gateway (ALG) or Deep Inspection to the custom service.

NOTE: ScreenOS supports ALGs for numerous services, including DNS, FTP, H.323, HTTP, RSH, SIP, Telnet, and TFTP.

Applying an ALG to a custom service involves the following two steps:

- Define a custom service with a name, timeout value, transport protocol, and source and destination ports
- When configuring a policy, reference that service and the application type for the ALG that you want to apply

For information about applying Deep Inspection to a custom service, see “Mapping Custom Services to Applications” on page 4-155.

Name

You can give a policy a descriptive name to provide a convenient means for identifying its purpose.

NOTE: For information regarding ScreenOS naming conventions—which apply to the names you create for policies—see “Command Line Interface Conventions” on page x.

VPN Tunneling

You can apply a single policy or multiple policies to any VPN tunnel that you have configured. In the WebUI, the VPN Tunnel option provides a drop-down list of all such tunnels. In the CLI, you can see all available tunnels with the **get vpn** command. (For more information, see “Site-to-Site Virtual Private Networks” on page 5-79 and “Dialup Virtual Private Networks” on page 5-159.)

When the VPN configurations at both ends of a VPN tunnel are using policy-based-NAT, then the administrators of both gateway devices each need to create an inbound and an outbound policy (four policies in total). When the VPN policies constitute a matching pair (that is, everything in the inbound and outbound policy configurations is the same except that the source and destination addresses are reversed), you can configure one policy and then select the **Modify matching bidirectional VPN policy** checkbox to create a second policy automatically for the opposite direction. For the configuration of a new policy, the matching VPN policy checkbox is cleared by default. For the modification of an existing policy that is a member of a matching pair, the checkbox is selected by default, and any changes made to one policy are propagated to the other.

NOTE: This option is only available through the WebUI. It is not supported when there are multiple entries for any of the following policy components: source address, destination address, or service.

L2TP Tunneling

You can apply a single policy or multiple policies to any Layer 2 Tunneling Protocol (L2TP) tunnel that you have configured. In the WebUI, the L2TP option provides a drop-down list of all such tunnels. In the CLI, you can display status of active L2TP tunnels with the **get l2tp tunn_str active** command, and see all available tunnels with the **get l2tp all** command. You can also combine a VPN tunnel and an L2TP tunnel—if both have the same endpoints—to create a tunnel combining the characteristics of each. This is called L2TP-over-IPSec.

NOTE: A security device in Transparent mode does not support L2TP.

Deep Inspection

Deep Inspection (DI) is a mechanism for filtering the traffic permitted at the Network and Transport Layers by examining not only these layers but the content and protocol characteristics at the Application Layer. The goal of DI is the detection and prevention any attacks or anomalous behavior that might be present in traffic that the Juniper Networks firewall permits.

NOTE: In the Open Systems Interconnection (OSI) Model, the Network Layer is Layer 3, the Transport Layer is Layer 4, and the Application Layer is Layer 7. The OSI Model is a networking industry standard model of network protocol architecture. The OSI Model consists of seven layers.

To configure a policy for attack protection, you must make two choices: which attack group (or groups) to use and which attack action to take if an attack is detected. (For more information, see “Deep Inspection” on page 4-119.)

Placement at the Top of the Policy List

By default, ScreenOS positions a newly created policy at the bottom of a policy set list. If you need to reposition the policy, you can use either of the policy reordering methods explained in “Reordering Policies” on page 190. To avoid the extra step of repositioning a newly created policy to the top of a policy set list, you can select the **Position at Top** option in the WebUI, or use the keyword **top** in the **set policy** command (**set policy top ...**) in the CLI.

Source Address Translation

You can apply source address translation (NAT-src) at the policy level. With NAT-src, you can translate the source address on either incoming or outgoing network and VPN traffic. The new source address can come from either a dynamic IP (DIP) pool or the egress interface. NAT-src also supports source port address translation (PAT). To learn about all the NAT-src options that are available, see “Source Network Address Translation” on page 8-13.

NOTE: You can also perform source address translation at the interface level, known as network address translation (NAT). For information about interface level NAT-src, or simply NAT, see “NAT Mode” on page 92.

Destination Address Translation

You can apply destination address translation (NAT-dst) at the policy level. With NAT-dst, you can translate the destination address on either incoming or outgoing network and VPN traffic. NAT-dst can also support destination port mapping. To learn about all the NAT-dst options that are available, see “Destination Network Address Translation” on page 8-27.

User Authentication

Selecting this option requires the auth user at the source address to authenticate his/her identity by supplying a username and password before traffic is allowed to traverse the firewall or enter the VPN tunnel. The security device can use the local database or an external RADIUS, SecurID, or LDAP auth server to perform the authentication check.

NOTE: If a policy requiring authentication applies to a subnet of IP addresses, authentication is required for each IP address in that subnet.

If a host supports multiple auth user accounts (as with a UNIX host running Telnet), then after the security device authenticates the first user, all other users from that host can pass traffic through the security device without being authenticated, having inherited the privileges of the first user.

ScreenOS provides two authentication schemes:

- Run-time authentication, in which the security device prompts an auth user to log on when it receives HTTP, FTP or Telnet traffic matching a policy that has authentication enabled
- WebAuth, in which a user must authenticate himself or herself before sending traffic through the security device

Run-Time Authentication

The run-time authentication process proceeds as follows:

1. When the auth user sends an HTTP, FTP or Telnet connection request to the destination address, the security device intercepts the packet and buffers it.
2. The security device sends the auth user a login prompt.
3. The auth user responds to this prompt with his/her username and password.
4. The security device authenticates the auth user's login information.

If the authentication is successful, a connection is established between the auth user and the destination address.

For the initial connection request, a policy must include one or all of the three following services: Telnet, HTTP, or FTP. Only a policy with one or all of these services is capable of initiating the authentication process. You can use any of the following services in a policy involving user authentication:

- Any (because "any" includes all three required services).
- Telnet, or FTP, or HTTP.
- A service group that includes the service or services you want, plus one or more of the three services required to initiate the authentication process (Telnet, FTP, or HTTP). For example, you can create a custom service group named "Login" that supports FTP, NetMeeting, and H.323 services. Then, when you create the policy, specify "Login" as the service.

For any connection following a successful authentication, all services specified in the policy are valid.

NOTE: A policy with authentication enabled does not support DNS (port 53) as the service.

Pre-Policy Check Authentication (WebAuth)

The WebAuth authentication process proceeds as follows:

1. The auth user makes an HTTP connection to the IP address of the WebAuth server.
2. The security device sends the auth user a login prompt.
3. The auth user responds to this prompt with his/her username and password.
4. The security device or an external auth server authenticates the auth user's login information.

If the authentication attempt is successful, the security device permits the auth user to initiate traffic to destinations as specified in policies that enforce authentication via the WebAuth method.

NOTE: For more information about these two user authentication methods, see “Referencing Auth Users in Policies” on page 9-46.

You can restrict or expand the range of auth users to which the policy applies by selecting a specific user group, local or external user, or group expression. (For information about group expressions, see “Group Expressions” on page 9-5.) If you do not reference an auth user or user group in a policy (in the WebUI, select the **Allow Any** option), the policy applies to all auth users in the specified auth server.

NOTE: ScreenOS links authentication privileges with the IP address of the host from which the auth user logs on. If the security device authenticates one user from a host behind a NAT device that uses a single IP address for all NAT assignments, then users at other hosts behind that NAT device automatically receive the same privileges.

HA Session Backup

When two security devices are in an NSRP cluster for high availability (HA), you can specify which sessions to backup and which not to backup. For traffic whose sessions you do not want backed up, apply a policy with the HA session backup option disabled. In the WebUI, clear the **HA Session Backup** checkbox. In the CLI, use the **no-session-backup** argument in the **set policy** command. By default, security devices in an NSRP cluster back up sessions.

Web Filtering

Web filtering, also called URL filtering, enables you to manage Internet access and prevent access to inappropriate web content. For more information, see “Web Filtering” on page 4-100.

Logging

When you enable logging in a policy, the security device logs all connections to which that particular policy applies. You can view the logs through either the WebUI or CLI. In the WebUI, click **Reports > Policies > Logging** (for the policy whose log you want to see). In the CLI, use the **get log traffic policy id_num** command.

NOTE: For more information about viewing logs and graphs, see “Monitoring Security Devices” on page 3-55.

Counting

When you enable counting in a policy, the security device counts the total number of bytes of traffic to which this policy applies and records the information in historical graphs. To view the historical graphs for a policy in the WebUI, click **Reports > Policies > Counting** (for the policy whose traffic count you want to see).

Traffic Alarm Threshold

You can set a threshold that triggers an alarm when the traffic permitted by the policy exceeds a specified number of bytes per second, bytes per minute, or both. Because the traffic alarm requires the security device to monitor the total number of bytes, you must also enable the counting feature.

NOTE: For more information about traffic alarms, see “Traffic Alarms” on page 3-68.

Schedules

By associating a schedule to a policy, you can determine when the policy is in effect. You can configure schedules on a recurring basis and as a one-time event. Schedules provide a powerful tool in controlling the flow of network traffic and in enforcing network security. For an example of the latter, if you were concerned about employees transmitting important data outside the company, you might set a policy that blocked outbound FTP-Put and MAIL traffic after normal business hours.

In the WebUI, define schedules on the **Policy > Policy Elements > Schedules** page. In the CLI, use the **set schedule** command.

NOTE: In the WebUI, scheduled policies appear with a gray background to indicate that the current time is not within the defined schedule. When a scheduled policy becomes active, it appears with a white background.

Antivirus Scanning

Some Juniper Networks security devices support an internal AV scanner that you can configure to filter FTP, HTTP, IMAP, POP3, and SMTP traffic. If the embedded AV scanner detects a virus, it drops the packet and sends a message reporting the virus to the client initiating the traffic.

NOTE: For more information about antivirus scanning, see “Antivirus Scanning” on page 173.

Traffic Shaping

You can set parameters for the control and shaping of traffic for each policy. Table 28 describes the traffic shaping parameters.

Table 28: Traffic Shaping Parameters

Parameter	Description
Guaranteed Bandwidth	Guaranteed throughput in kilobits per second (kbps). Traffic below this threshold passes with the highest priority without being subject to any traffic management or shaping mechanism.
Maximum Bandwidth	Secured bandwidth available to the type of connection in kilobits per second (kbps). Traffic beyond this threshold is throttled and dropped. Note: It is advised that you do not use rates less than 10 Kbps. Rates below this threshold lead to dropped packets and excessive retries that defeat the purpose of traffic management.
Traffic Priority	When traffic bandwidth falls between the guaranteed and maximum settings, the security device passes higher priority traffic first, and lower priority traffic only if there is no other higher priority traffic. There are eight priority levels.
DiffServ Codepoint Marking	Differentiated Services (DiffServ) is a system for tagging (or “marking”) traffic at a position within a hierarchy of priority. You can map the eight ScreenOS priority levels to the DiffServ system. By default, the highest priority (priority 0) in the ScreenOS system maps to the first three bits (111) in the DiffServ field (see RFC 2474), or the IP precedence field in the TOS byte (see RFC 1349), in the IP packet header. The lowest priority (priority 7) in ScreenOS maps to (000) in the TOS DiffServ system. When you enable DSCP, ScreenOS overwrites the first 3 bits in the ToS byte with the IP precedence priority. When you enable DSCP and set a <i>dscp-byte value</i> , ScreenOS overwrites the first 6 bits of the ToS byte with the DSCP value. Note: Some devices require that you explicitly enable DSCP marking by setting a system-wide environmental variable. Refer to the installation and configuration guide for your device to find out if it requires that you explicitly enable DSCP marking before using it in policies. If your device requires it, use the following command to enable DSCP marking system wide: set envvar ipsec-dscp-mark= yes . This variable cannot be set using the WebUI. Use the unset envvar ipsec-dscp-mark to disable DSCP marking system wide.

NOTE: For a more detailed discussion of traffic management and shaping, see “Traffic Shaping” on page 193.

To change the mapping between the ScreenOS priority levels and the DiffServ system, use the following CLI command:

```
set traffic-shaping ip_precedence number0 number1 number2 number3 number4
number5 number6 number7
```

where *number0* is the mapping for priority 0 (the highest priority in the TOS DiffServ system), *number1* is the mapping for priority 1, and so on.

To subsume IP precedence into class selector codepoints—that is, to zero out the second three bits in the DiffServ field and thus insure that priority levels you set with **ip_precedence** are preserved and handled correctly by downstream routers—use the following CLI command:

```
set traffic-shaping dscp-class-selector
```

Policies Applied

This section describes the management of policies: viewing, creating, modifying, ordering and reordering, and removing policies.

Viewing Policies

To view policies through the WebUI, click **Policies**. You can sort the displayed policies by source and destination zones by choosing zone names from the **From** and **To** drop-down lists and then clicking **Go**. In the CLI, use the **get policy [all | from zone to zone | global | id number]** command.

Creating Policies

To allow traffic to flow between two zones, you create policies to deny, permit, reject, or tunnel traffic between those zones. You can also create policies to control traffic within the same zone if the security device is the only network device that can route the intrazone traffic between the source and destination addresses referenced in the policy. You can also create global policies, which make use of source and destination addresses in the Global zone address book.

To allow bidirectional traffic between two zones—for example, between the Trust and Untrust zones—you need to create a policy that goes from Trust to Untrust and then create a second policy from Untrust to Trust. Depending on your needs, the policies can use the same or different IP addresses, only the source and destination addresses are reversed.

You can define policies between any zones that are located within the same system—root or virtual. To define a policy between the root system and a vsys, one of the zones must be a shared zone. (For information about shared zones in relation to virtual systems, see *Volume 10: Virtual Systems*.)

Creating Interzone Policies Mail Service

In this example, you create three policies to control the flow of email traffic.

The first policy allows internal users in the Trust zone to send and retrieve email from a local mail server in the DMZ zone. This policy permits the services MAIL (that is, SMTP) and POP3 originating from the internal users to traverse the Juniper Networks firewall to reach the local mail server.

The second and third policies permit the service MAIL to traverse the firewall between the local mail server in the DMZ zone and a remote mail server in the Untrust zone.

However, before creating policies to control traffic between different security zones, you must first design the environment in which to apply those policies. First, you first bind interfaces to zones and assign the interfaces IP addresses:

- Bind ethernet0/1 to the Trust zone and assign it IP address 10.1.1.1/24.
- Bind ethernet0/2 to the DMZ zone and assign it IP address 1.2.2.1/24.
- Bind ethernet0/3 to the Untrust zone and assign it IP address 1.1.1.1/24.

All security zones are in the trust-vr routing domain.

Second, you create addresses for use in the policies:

- Define an address in the Trust zone named “corp_net” and assign it IP address 10.1.1.0/24.
- Define an address in the DMZ zone named “mail_svr” and assign it IP address 1.2.2.5/32.
- Define an address in the Untrust zone named “r-mail_svr” and assign it IP address 2.2.2.5/32.

Third, you create a service group named “MAIL-POP3” containing the two predefined services MAIL and POP3.

Fourth, you configure a default route in the trust-vr routing domain pointing to the external router at 1.1.1.250 through ethernet0/3.

After completing steps 1 through 4, you can then create the policies necessary to permit the transmission, retrieval, and delivery of email in and out of your protected network.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24

Enter the following, then click **OK**:
 Interface Mode: NAT

Network > Interfaces > Edit (for ethernet0/2): Enter the following, then click **OK**:

Zone Name: DMZ
 Static IP: (select this option when present)
 IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone Name: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24

2. Addresses

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: corp_net
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.1.0/24
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: mail_svr
 IP Address/Domain Name:
 IP/Netmask: (select), 1.2.2.5/32
 Zone: DMZ

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: r-mail_svr
 IP Address/Domain Name:
 IP/Netmask: (select), 2.2.2.5/32
 Zone: Untrust

3. Service Group

Policy > Policy Elements > Services > Groups: Enter the following group name, move the following services, then click **OK**:

Group Name: MAIL-POP3

Select **MAIL** and use the << button to move the service from the Available Members column to the Group Members column.

Select **POP3** and use the << button to move the service from the Available Members column to the Group Members column.

4. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/3
 Gateway IP Address: 1.1.1.250

5. Policies

Policy > Policies > (From: Trust, To: DMZ) > New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), corp_net
 Destination Address:
 Address Book Entry: (select), mail_svr
 Service: Mail-POP3
 Action: Permit

Policy > Policies > (From: DMZ, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), mail_svr
 Destination Address:
 Address Book Entry: (select), r-mail_svr
 Service: MAIL
 Action: Permit

Policy > Policies > (From: Untrust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), r-mail_svr
 Destination Address:
 Address Book Entry: (select), mail_svr
 Service: MAIL
 Action: Permit

CLI**1. Interfaces**

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/2 zone dmz
set interface ethernet0/2 ip 1.2.2.1/24
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust corp_net 10.1.1.0/24
set address dmz mail_svr 1.2.2.5/32
set address untrust r-mail_svr 2.2.2.5/32
```

3. Service Group

```
set group service MAIL-POP3
set group service MAIL-POP3 add mail
```

```
set group service MAIL-POP3 add pop3
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 1.1.1.250
```

5. Policies

```
set policy from trust to dmz corp_net mail_svr MAIL-POP3 permit
set policy from dmz to untrust mail_svr r-mail_svr MAIL permit
set policy from untrust to dmz r-mail_svr mail_svr MAIL permit
save
```

Creating an Interzone Policy Set

A small software firm, ABC Design, has divided its internal network into two subnets, and both are in the Trust zone. These two subnets are:


- Engineering (with the defined address “Eng”)
- The rest of the company (with the defined address “Office”)

ABC Design also has a DMZ zone for its web and mail servers.

The following example presents a typical set of policies for the following users:

- “Eng” can use all the services for outbound traffic except FTP-Put, IMAP, MAIL, and POP3.
- “Office” can use email and access the Internet, provided they authenticate themselves via WebAuth. (For information about WebAuth user authentication, see “Authentication Users” on page 9-45.)
- Everyone in the Trust zone can access the Web and mail servers in the DMZ zone.
- A remote mail server in the Untrust zone can access the local mail server in the DMZ zone.
- There is also a group of system administrators (with the user-defined address “sys-admins”) who have complete user and administrative access to the servers in the DMZ zone.

Figure 60: Interzone Policy Set



It is assumed that you have already configured the interfaces, addresses, service groups, and routes that must be in place. For more information about configuring these, see “Interfaces” on page 35, “Addresses” on page 103, “Service Groups” on page 138, and *Volume 7: Routing*.

Table 29: Configured Policies

From Zone - Src Addr	To Zone - Dest Addr	Service	Action
Trust - Any	Untrust - Any	Com (service group: FTP-Put, IMAP, MAIL, POP3)	Reject
Trust - Eng	Untrust - Any	Any	Permit
Trust - Office	Untrust - Any	Internet (service group: FTP-Get, HTTP, HTTPS)	Permit (+ WebAuth)
Untrust - Any	DMZ - mail.abc.com	MAIL	Permit
Untrust - Any	DMZ - www.abc.com	Web (service group: HTTP, HTTPS)	Permit
Trust - Any	DMZ - mail.abc.com	Email (service group: IMAP, MAIL, POP3)	Permit
Trust - Any	DMZ - www.abc.com	Internet (service group: FTP-Get, HTTP, HTTPS)	Permit
Trust - sys-admins	DMZ - Any	Any	Permit
DMZ - mail.abc.com	Untrust - Any	MAIL	Permit

NOTE: The default policy is to deny all.

WebUI

1. From Trust to Untrust

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Eng
Destination Address:
Address Book Entry: (select), Any
Service: ANY
Action: Permit

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Office
Destination Address:
Address Book Entry: (select), Any
Service: Internet
Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)
WebAuth: (select)

NOTE: “Internet” is a service group with the following members: FTP-Get, HTTP, and HTTPS.

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
Address Book Entry: (select), Any
Destination Address:
Address Book Entry: (select), Any
Service: Com
Action: Reject
Position at Top: (select)

NOTE: “Com” is a service group with the following members: FTP-Put, MAIL, IMAP, and POP3.

For traffic from the Untrust zone to the Trust zone, the default deny policy denies everything.

2. From Untrust to DMZ

Policy > Policies > (From: Untrust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), mail.abc.com
 Service: MAIL
 Action: Permit

Policy > Policies > (From: Untrust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), www.abc.com
 Service: Web
 Action: Permit

NOTE: “Web” is a service group with the following members: HTTP and HTTPS.

3. From Trust to DMZ

Policy > Policies > (From: Trust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), mail.abc.com
 Service: e-mail
 Action: Permit

NOTE: “e-mail” is a service group with the following members: MAIL, IMAP, and POP3.

Policy > Policies > (From: Trust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), www.abc.com
 Service: Internet
 Action: Permit

Policy > Policies > (From: Trust, To: DMZ) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), sys-admins
 Destination Address:
 Address Book Entry: (select), Any
 Service: ANY

Action: Permit

4. From DMZ to Untrust

Policy > Policies > (From: DMZ, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), mail.abc.com
 Destination Address:
 Address Book Entry: (select), Any
 Service: MAIL
 Action: Permit

CLI

1. From Trust to Untrust

```
set policy from trust to untrust eng any any permit
set policy from trust to untrust office any Internet permit webauth
set policy top from trust to untrust any any Com reject
```

NOTE: “Internet” is a service group with the following members: FTP-Get, HTTP, and HTTPS.

“Com” is a service group with the following members: FTP-Put, MAIL, IMAP, and POP3.

2. From Untrust to DMZ

```
set policy from untrust to dmz any mail.abc.com mail permit
set policy from untrust to dmz any www.abc.com Web permit
```

NOTE: “Web” is a service group with the following members: HTTP and HTTPS.

3. From Trust to DMZ

```
set policy from trust to dmz any mail.abc.com e-mail permit
set policy from trust to dmz any www.abc.com Internet permit
set policy from trust to dmz sys-admins any any permit
```

NOTE: “e-mail” is a service group with the following members: MAIL, IMAP, and POP3.

“Internet” is a service group with the following members: FTP-Get, HTTP, and HTTPS.

4. From DMZ to Untrust

```
set policy from dmz to untrust mail.abc.com any mail permit
save
```

Creating Intrazone Policies

In this example, you create an intrazone policy to permit a group of accountants access to a confidential server on the corporate LAN in the Trust zone. You first bind ethernet0/1 to the Trust zone and give it IP address 10.1.1.1/24. You then bind ethernet0/2 to the Trust zone and assign it IP address 10.1.5.1/24. You enable intrazone blocking in the Trust zone. Next, you define two addresses—one for a

server on which the company stores its financial records (10.1.1.100/32) and another for the subnet on which hosts for the accounting department are located (10.1.5.0/24). You then create an intrazone policy to permit access to the server from those hosts.

WebUI

1. Trust Zone—Interfaces and Blocking

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.1.1/24
 Select the following, then click **OK**:
 Interface Mode: NAT

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.5.1/24
 Select the following, then click **OK**:
 Interface Mode: NAT

Network > Zones > Edit (for Trust): Enter the following, then click **OK**:

Block Intra-Zone Traffic: (select)

2. Addresses

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: Hamilton
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.1.100/32
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: accounting
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.5.0/24
 Zone: Trust

3. Policy

Policy > Policies > (From: Trust, To: Trust) > New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), accounting
 Destination Address:
 Address Book Entry: (select), Hamilton
 Service: ANY
 Action: Permit

CLI

1. Trust Zone—Interfaces and Blocking

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/1 nat
set interface ethernet0/2 zone trust
set interface ethernet0/2 ip 10.1.5.1/24
set interface ethernet0/2 nat
set zone trust block
```

2. Addresses

```
set address trust Hamilton 10.1.1.100/32
set address trust accounting 10.1.5.0/24
```

3. Policy

```
set policy from trust to trust accounting Hamilton any permit
save
```

Creating a Global Policy

In this example, you create a global policy so that every host in every zone can access the company website, which is `www.juniper.net`. Using a global policy is a convenient shortcut when there are many security zones. In this example, one global policy accomplishes what n interzone policies would have accomplished (where n = number of zones).

NOTE: To use a domain name instead of an IP address, be sure to have DNS service configured on the security device.

WebUI

1. Global Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

```
Address Name: server1
IP Address/Domain Name:
    Domain Name: (select), www.juniper.net
Zone: Global
```

2. Policy

Policy > Policies > (From: Global, To: Global) > New: Enter the following, then click **OK**:

```
Source Address:
    Address Book Entry: (select), Any
Destination Address:
    Address Book Entry: (select), server1
Service: HTTP
Action: Permit
```

CLI**1. Global Address**

```
set address global server1 www.juniper.net
```

2. Policy

```
set policy global any server1 http permit
save
```

Entering a Policy Context

When configuring a policy through the CLI, after you first create a policy, you can then enter the context of the policy to make additions and modifications. For example, perhaps you first create the following policy:

```
set policy id 1 from trust to untrust host1 server1 HTTP permit attack
HIGH:HTTP:SIGS action close
```

If you want to make some changes to the policy, such as adding another source or destination address, another service, or another attack group, you can enter the context for policy 1 and then enter the pertinent commands:

```
set policy id 1
device(policy:1)-> set src-address host2
device(policy:1)-> set dst-address server2
device(policy:1)-> set service FTP
device(policy:1)-> set attack CRITICAL:HTTP:SIGS
```

You can also remove multiple items for a single policy component as long as you do not remove them all. For example, you can remove server2 from the above configuration, but not server2 and server1 because then no destination address would remain.

Multiple Items per Policy Component

ScreenOS allows you to add multiple items to the following components of a policy:

- Source address
- Destination address
- Service
- Attack group

In ScreenOS releases prior to 5.0.0, the only way to have multiple source and destination addresses or services is to first create an address or service group with multiple members and then reference that group in a policy. You can still use address and service groups in policies in ScreenOS 5.0.0. In addition, you can simply add extra items directly to a policy component.

NOTE: If the first address or service referenced in a policy is “Any,” you cannot logically add anything else to it. ScreenOS prevents this kind of misconfiguration and displays an error message should it occur.

To add multiple items to a policy component, do either of the following:

WebUI

To add more addresses and services, click the **Multiple** button next to the component to which you want to add more items. To add more attack groups, click the **Attack Protection** button. Select an item in the “Available Members” column, then use the < < key to move it to the “Active Members” column. You can repeat this action with other items. When finished, click **OK** to return to the policy configuration page.

CLI

Enter the policy context with the following command:

```
set policy id number
```

Then use one of the following commands as appropriate:

```
device(policy:number)-> set src-address string
device(policy:number)-> set dst-address string
device(policy:number)-> set service string
device(policy:number)-> set attack string
```

Setting Address Negation

You can configure a policy so that it applies to all addresses except the one specified as either the source or destination. For example, you might want to create a policy that permits Internet access to everyone except the “P-T_contractors” address group. To accomplish this, you can use the address negation option.

In the WebUI, this option is available on the pop-up that appears when you click **Multiple** next to either **Source Address** or **Destination Address** on the policy configuration page.

In the CLI, you insert an exclamation point (!) immediately before source or destination address.

NOTE: Address negation occurs at the policy component level, applying to all items in the negated component.

In this example, you create an intrazone policy that allows all addresses in the Trust zone access to all FTP servers except to an FTP server named “vulcan”, which engineering uses to post functional specifications for one another.

However, before creating the policy, you must first design the environment in which to apply it. First, you enable intrazone blocking for the Trust zone. Intrazone blocking requires a policy lookup before the security device passes traffic between two interfaces bound to the same zone.

Second, you bind two interfaces to the Trust zone and assign them IP addresses:

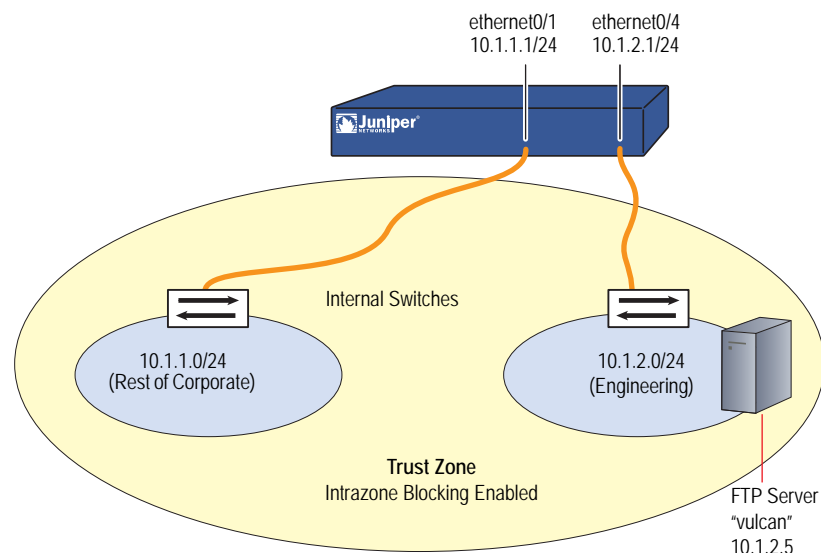
- You bind ethernet0/1 to the Trust zone and assign it IP address 10.1.1.1/24.
- You bind ethernet0/4 to the Trust zone and assign it IP address 10.1.2.1/24.

Third, you create an address (10.1.2.5/32) for the FTP server named “vulcan” in the Trust zone.

After completing these two steps, you can then create the intrazone policies.

NOTE: You do not have to create a policy for the engineering department to reach their FTP server because the engineers are also in the 10.1.2.0/24 subnet and do not have to cross the Juniper Networks firewall to reach their own server.

Figure 61: Intrazone Policies Negation



WebUI

1. Intrazone Blocking

Network > Zones > Edit (for Trust): Enter the following, then click **OK**:

Virtual Router Name: trust-vr
Block Intra-Zone Traffic: (select)

2. Trust Zone Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone Name: Trust
Static IP: (select this option when present)
IP Address/Netmask: 10.1.1.1/24
Select the following, then click **OK**:
Interface Mode: NAT

Network > Interfaces > Edit (for ethernet0/4): Enter the following, then click **Apply**:

Zone Name: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 10.1.2.1/24
 Select the following, then click **OK**:
 Interface Mode: NAT

3. Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: vulcan
 IP Address/Domain Name:
 IP/Netmask: (select), 10.1.2.5/32
 Zone: Trust

4. Policy

Policy > Policies > (From: Trust, To: Trust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), vulcan

> Click **Multiple**, select **Negate Following**, then click **OK** to return to the basic policy configuration page.

Service: FTP
 Action: Permit

CLI

1. Intrazone Blocking

set zone trust block

2. Trust Zone Interfaces

set interface ethernet0/1 zone trust
 set interface ethernet0/1 ip 10.1.1.1/24
 set interface ethernet0/1 nat
 set interface ethernet0/4 zone trust
 set interface ethernet0/4 ip 10.1.2.1/24
 set interface ethernet0/4 nat

3. Address

set address trust vulcan 10.1.2.5/32

4. Policy

set policy from trust to trust any !vulcan ftp permit
 save

Modifying and Disabling Policies

After you create a policy, you can always return to it to make modifications. In the WebUI, click the **Edit** link in the Configure column for the policy that you want to change. In the Policy configuration page that appears for that policy, make your changes, then click **OK**. In the CLI, use the **set policy** command.

ScreenOS also provides a means for enabling and disabling policies. By default, a policy is enabled. To disable it, do the following:

WebUI

Policies: Clear the **Enable** checkbox in the Configure column for the policy that you want to disable.

The row of text for a disabled policy appears as grey.

CLI

```
set policy id id_num disable
save
```

NOTE: To enable the policy again, select **Enable** in the Configure column for the policy that you want to enable (WebUI), or type **unset policy id *id_num* disable** (CLI).

Policy Verification

ScreenOS offers a tool for verifying that the order of policies in the policy list is valid. It is possible for one policy to eclipse, or “shadow,” another policy. Consider the following example:

```
set policy id 1 from trust to untrust any any HTTP permit
set policy id 2 from trust to untrust any dst-A HTTP deny
```

Because the security device performs a policy lookup starting from the top of the list, when it finds a match for traffic received, it does not look any lower in the policy list. In the above example, the security device never reaches policy 2 because the destination address “any” in policy 1 includes the more specific “dst-A” address in policy 2. When an HTTP packet arrives at the security device from an address in the Trust zone bound for dst-A in the Untrust zone, the security device always first finds a match with policy 1.

To correct the above example, you can simply reverse the order of the policies, putting the more specific one first:

```
set policy id 2 from trust to untrust any dst-A HTTP deny
set policy id 1 from trust to untrust any any HTTP permit
```

Of course, this example is purposefully simple to illustrate the basic concept. In cases where there are dozens or hundreds of policies, the eclipsing of one policy by another might not be so easy to spot. To check if there is any policy shadowing in your policy list, you can use the following CLI command:

```
exec policy verify
```

This command reports the shadowing and shadowed policies. It is then the admin’s responsibility to correct the situation.

NOTE: The concept of policy “shadowing” refers to the situation where a policy higher in the policy list always takes effect before a subsequent policy. Because the policy lookup always uses the first policy it finds that matches the five-part tuple of source and destination zone, source and destination address, and service type, if another policy applies to the same tuple (or a subset of the tuple), the policy lookup uses the first policy in the list and never reaches the second one.

The policy verification tool cannot detect the case where a combination of policies shadows another policy. In the following example, no single policy shadows policy 3; however, policies 1 and 2 together do shadow it:

```
set group address trust grp1 add host1
set group address trust grp1 add host2
set policy id 1 from trust to untrust host1 server1 HTTP permit
set policy id 2 from trust to untrust host2 server1 HTTP permit
set policy id 3 from trust to untrust grp1 server1 HTTP deny
```

Reordering Policies

The security device checks all attempts to traverse the firewall against policies, beginning with the first one listed in the policy set for the appropriate list (see “Policy Set Lists” on page 163) and moving through the list. Because the security device applies the action specified in the policy to the first matching policy in the list, you must arrange them from the most specific to the most general. (Whereas a specific policy does not preclude the application of a more general policy located down the list, a general policy appearing before a specific one does.)

By default, a newly created policy appears at the bottom of a policy set list. There is an option that allows you to position a policy at the top of the list instead. In the Policy configuration page in the WebUI, select **Position at Top**. In the CLI, add the key word **top** to the **set policy** command: **set policy top ...**

To move a policy to a different position in the list, do either of the following:

WebUI

There are two ways to reorder policies in the WebUI: by clicking the circular arrows or by clicking the single arrow in the Configure column for the policy you want to move.

If you click the circular arrows:

A User Prompt dialog box appears.

To move the policy to the very end of the list, enter `< -1 >` . To move it up in the list, enter the ID number of the policy above which you want to move the policy in question.

Click **OK** to execute the move.

If you click the single arrow:

A Policy Move page appears displaying the policy you want to move and a table displaying the other policies.

In the table displaying the other policies, the first column, Move Location, contains arrows pointing to various locations where you can move the policy. Click the arrow that points to the location in the list where you want to move the policy.

The Policy List page reappears with the policy you moved in its new position.

CLI

```
set policy move id_num { before | after } number  
save
```

Removing a Policy

In addition to modifying and repositioning a policy, you can also delete it. In the WebUI, click **Remove** in the Configure column for the policy that you want to remove. When the system message prompts for confirmation to proceed with the removal, click **Yes**. In the CLI, use the **unset policy *id_num*** command.

Chapter 7

Traffic Shaping

This chapter discusses the various ways you can use your Juniper Networks security device to manage limited bandwidth without compromising quality and availability of the network to all of your users. It contains the following sections:

- “Managing Bandwidth at the Policy Level” on page 194
- “Setting Traffic Shaping” on page 195
- “Setting Service Priorities” on page 198
- “Setting Priority Queuing” on page 199
- “Ingress Policing” on page 203
- “Shaping Traffic on Virtual Interfaces” on page 203
- “DSCP Marking and Shaping” on page 214

Traffic shaping is the allocation of the appropriate amount of network bandwidth to every user and application on an interface. The appropriate amount of bandwidth is defined as cost-effective carrying capacity at a guaranteed Quality of Service (QoS). You use a security device to shape traffic by creating policies and by applying appropriate rate controls to each class of traffic going through the device.

Managing Bandwidth at the Policy Level

To classify traffic, you create policies and specify the amount of guaranteed bandwidth and maximum bandwidth, and the priority for each class of traffic. Guaranteed bandwidth and maximum bandwidth are not strictly policy-based. With multiple physical interfaces in the egress zone, bandwidth is based both on policy and on total available egress physical interface bandwidth. The physical bandwidth of every interface is allocated to the guaranteed bandwidth parameter for all traffic-shaping policies. If there is any bandwidth left over, it is sharable by any other traffic. In other words, each policy gets its guaranteed bandwidth and shares whatever is left over, on a priority basis (up to the limit of its maximum bandwidth specification), with all other policies.

The traffic-shaping function applies to traffic from all policies. If you turn off traffic shaping for a specific policy, while traffic shaping is still turned on for other policies, the system applies a default traffic-shaping policy to that particular policy, with the following parameters:

- Guaranteed bandwidth 0
- Unlimited maximum bandwidth
- Priority of 7 (the lowest priority setting)

NOTE: You can enable a mapping of priority levels to the DiffServ Codepoint Marking system. For more information about DS Codepoint Marking, see “Traffic Shaping” on page 173.


If you do not want the system to assign this default traffic-shaping policy to policies for which you have turned off traffic shaping, you can turn off traffic shaping system wide via the CLI command: **set traffic-shaping mode off**. Use the CLI command: **set traffic-shaping mode on** to turn on shaping on an interface. Or, you can set traffic shaping to automatic in the WebUI: **Configuration > Advanced > Traffic Shaping**. This allows the system to turn on traffic shaping when a policy requires it and to turn off traffic shaping when policies do not require it.

NOTE: You cannot perform traffic-shaping on ASIC-based ISG 1000, ISG 2000, and NS5000 devices. ScreenOS supports traffic-shaping on non-ASIC devices.

Setting Traffic Shaping

In this example, you partition 45Mbps of bandwidth on a T3 interface among three departments on the same subnet. The interface ethernet0/1 is bound to the Trust zone, and ethernet0/3 is bound to the Untrust zone.

Figure 62: Traffic Shaping



WebUI

1. Bandwidth on Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Traffic Bandwidth: 45000

NOTE: If you do not specify bandwidth settings on an interface, the security device uses the available physical bandwidth.

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Traffic Bandwidth: 45000

2. Bandwidth in Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: Marketing Traffic Shaping
 Source Address:
 Address Book Entry: (select), Marketing
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit
 VPN Tunnel: None

NOTE: You can also enable traffic shaping in policies referencing VPN tunnels.

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 10000
 Maximum Bandwidth: 15000

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: Sales Traffic Shaping Policy
 Source Address:
 Address Book Entry: (select), Sales
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 10000
 Maximum Bandwidth: 10000

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: Support Traffic Shaping Policy
 Source Address:
 Address Book Entry: (select), Support
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 5000
 Maximum Bandwidth: 10000

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Name: Allow Incoming Access to Marketing
 Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Marketing
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 10000
 Maximum Bandwidth: 10000

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Name: Allow Incoming Access to Sales
 Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Sales
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 5000
 Maximum Bandwidth: 10000

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Name: Allow Incoming Access to Support
 Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Support
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 5000
 Maximum Bandwidth: 5000

CLI

To enable traffic shaping by policy, do the following:

1. Bandwidth on Interfaces

```
set interface ethernet0/1 bandwidth 45000
set interface ethernet0/3 bandwidth 45000
```

NOTE: If you do not specify bandwidth settings on an interface, the security device uses the available physical bandwidth.

2. Bandwidth in Policies

```

set policy name "Marketing Traffic Shaping" from trust to untrust marketing any
any permit traffic gbw 10000 priority 0 mbw 15000
set policy name "Sales Traffic Shaping Policy" from trust to untrust sales any any
permit traffic gbw 10000 priority 0 mbw 10000
set policy name "Support Traffic Shaping Policy" from trust to untrust support any
any permit traffic gbw 5000 priority 0 mbw 10000
set policy name "Allow Incoming Access to Marketing" from untrust to trust any
marketing any permit traffic gbw 10000 priority 0 mbw 10000
set policy name "Allow Incoming Access to Sales" from untrust to trust any sales
any permit traffic gbw 5000 priority 0 mbw 10000
set policy name "Allow Incoming Access to Support" from untrust to trust any
support any permit traffic gbw 5000 priority 0 mbw 5000
save

```

Setting Service Priorities

The traffic-shaping feature on Juniper Networks security devices allows you to perform priority queuing on the bandwidth that is not allocated to guaranteed bandwidth, or unused guaranteed bandwidth. Priority queuing is a feature that allows all your users and applications to have access to available bandwidth as they need it, while ensuring that important traffic can get through, if necessary at the expense of less important traffic. Queuing allows the security device to buffer traffic in up to eight different priority queues. These eight queues are:

- High priority
- 2nd priority
- 3rd priority
- 4th priority
- 5th priority
- 6th priority
- 7th priority
- Low priority (default)

The priority setting for a policy means that the bandwidth not already guaranteed to other policies is queued on the basis of high priority first and low priority last. Policies with the same priority setting compete for bandwidth in a round robin fashion. The security device processes all of the traffic from all of the policies with high priority before processing any traffic from policies with the next lower priority setting, and so on, until all traffic requests have been processed. If traffic requests exceed available bandwidth, the lowest priority traffic is dropped.



CAUTION: Be careful not to allocate more bandwidth than the interface can support. The policy configuration process does not prevent you from creating unsupported policy configurations. You can lose data if the guaranteed bandwidth on contending policies surpasses the traffic bandwidth set on the interface.

If you do not allocate any guaranteed bandwidth, then you can use priority queuing to manage all of traffic on your network. That is, all high priority traffic is sent before any 2nd priority traffic is sent, and so on. The security device processes low priority traffic only after all other traffic has been processed.

Setting Priority Queuing

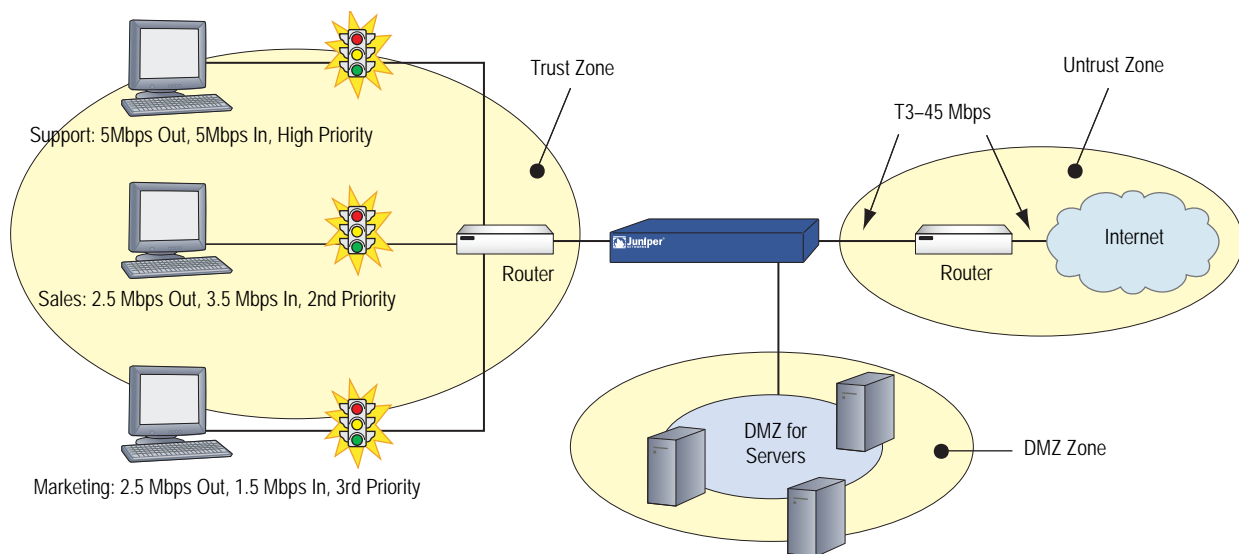
In this example, you configure the guaranteed and maximum bandwidth (in Mbps) for three departments—Support, Sales, and Marketing—as shown in Table 30.

Table 30: Maximum Bandwidth Configuration

	Outbound Guaranteed	Inbound Guaranteed	Combined Guaranteed	Priority
Support	5	5	10	High
Sales	2.5	3.5	6	2
Marketing	2.5	1.5	4	3
Total	10	10	20	

If all three departments send and receive traffic concurrently through the firewall, the security device must allocate 20 Mbps of bandwidth to fulfill the guaranteed policy requirements. The interface ethernet0/1 is bound to the Trust zone, and ethernet0/3 is bound to the Untrust zone.

Figure 63: Priority Queuing



WebUI

1. Bandwidth on Interfaces

Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Traffic Bandwidth: 40000

Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Traffic Bandwidth: 40000

2. Bandwidth in Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: Sup-out
 Source Address:
 Address Book Entry: (select), Support
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 5000
 Maximum Bandwidth: 40000
 Traffic Priority: High priority
 DiffServ Codepoint Marking: (select)

NOTE: Differentiated Services (DS) is a system for tagging (or “marking”) traffic at a position within a hierarchy of priority. DS Codepoint Marking maps the ScreenOS priority level of the policy to the first three bits of codepoint in the DS field in the IP packet header. For more information about DS Codepoint Marking, see “Traffic Shaping” on page 173.

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: Sal-out
 Source Address:
 Address Book Entry: (select), Sales
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 2500
 Maximum Bandwidth: 40000
 Traffic Priority: 2nd priority
 DiffServ Codepoint Marking: Enable

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: Mar-out
 Source Address:
 Address Book Entry: (select), Marketing
 Destination Address:
 Address Book Entry: (select), Any
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 2500
 Maximum Bandwidth: 40000
 Traffic Priority: 3rd priority
 DiffServ Codepoint Marking: (select)

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Name: Sup-in
 Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Support
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 5000
 Maximum Bandwidth: 40000
 Traffic Priority: High priority
 DiffServ Codepoint Marking: (select)

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Name: Sal-in
 Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Sales
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 3500
 Maximum Bandwidth: 40000
 Traffic Priority: 2nd priority
 DiffServ Codepoint Marking: (select)

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Name: Mar-in
 Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), Marketing
 Service: Any
 Action: Permit

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)
 Guaranteed Bandwidth: 1500
 Maximum Bandwidth: 40000
 Traffic Priority: 3rd priority
 DiffServ Codepoint Marking: (select)

CLI

1. Bandwidth on Interfaces

```
set interface ethernet0/1 bandwidth 40000
set interface ethernet0/3 bandwidth 40000
```

2. Bandwidth in Policies

```
set policy name sup-out from trust to untrust support any any permit traffic gbw
5000 priority 0 mbw 40000 enable
set policy name sal-out from trust to untrust sales any any permit traffic gbw 2500
priority 2 mbw 40000 dscp enable
```

NOTE: Some devices require that you explicitly enable DSCP marking by setting a system-wide environmental variable. Refer to the installation and configuration guide for your device to find out if it requires that you explicitly enable DSCP marking before using it in policies. If your device requires it, use the following command to enable DSCP marking system wide: **set envvar ipsec-dscp-mark= yes**. This variable cannot be set using the WebUI. Use the **unset envvar ipsec-dscp-mark** to disable DSCP marking system wide.

```
set policy name mar-out from trust to untrust marketing any any permit traffic gbw
2500 priority 3 mbw 40000 dscp enable
set policy name sup-in from untrust to trust any support any permit traffic gbw
5000 priority 0 mbw 40000 dscp enable
set policy name sal-in from untrust to trust any sales any permit traffic gbw 3500
priority 2 mbw 40000 dscp enable
set policy name mar-in from untrust to trust any marketing any permit traffic gbw
1500 priority 3 mbw 40000 dscp enable
save
```


Ingress Policing

Ingress policing is traffic control at the ingress side of the security device. By constraining the flow of traffic at the point of ingress, traffic exceeding your bandwidth setting is dropped with minimal processing, conserving system resources. You can configure ingress policing at the interface level and in security policies.

You configure ingress policing on an interface by setting a maximum bandwidth (the **mbw** keyword). The following command, for example, limits bandwidth on ethernet0/1, the ingress interface, to 22 Mbps:

```
set interface ethernet0/1 bandwidth ingress mbw 22000
```

Incoming traffic on ethernet0/1 exceeding this bandwidth is dropped. If you set traffic shaping at the interface, you must also set traffic-shaping mode to on (**set traffic-shaping mode on**).

To apply ingress policing to a specific application, however, requires a policy. The following command creates a policy called *my_ftp* that limits FTP bandwidth on the ingress side of the security device to 10 Mbps:

```
set policy my_ftp from untrust to trust any any ftp permit traffic pbw 10000
```

Incoming FTP traffic exceeding the configured policing bandwidth (the **pbw** keyword) is dropped. You can also set **mbw** in the policy, but at the policy level **mbw** applies only to the egress side of traffic flow—traffic exceeding your configured rate is still processed, and is dropped only at the egress side (see Figure 65, “Traffic-Shaping Packet Flow” on page 206). You can configure **mbw** or **pbw** in a policy, but not both.

Configuration and enforcement of ingress policing on virtual interfaces is the same as on physical interfaces, with the one exception that you can also configure guaranteed bandwidth (the **gbw** keyword) on virtual interfaces (see “Policy-Level Traffic Shaping” on page 205). On physical interfaces, guaranteed bandwidth is the same as maximum bandwidth.

NOTE: Ingress policing on tunnel interfaces is enforced after the encrypted packets are decrypted by the VPN engine.

Shaping Traffic on Virtual Interfaces

In the context of traffic shaping, the term *virtual interfaces* refers only to subinterfaces and tunnel interfaces—not to other types of virtual interfaces, such as virtual security interfaces (VSI), or aggregate or redundant interfaces. You cannot configure shaping parameters in policies created in a vsys. Similarly, bandwidth cannot be shaped on interfaces owned (inherited) by a user-created vsys. See *Volume 10: Virtual Systems* for more information.


Traffic shaping (as distinct from ingress policing) concerns traffic management at the egress side of the security device. As with physical interfaces, you shape traffic on virtual interfaces by setting bandwidth values at the interface level, and in policies.

Interface-Level Traffic Shaping

Traffic shaping at the interface level is control of the minimum and maximum rate of traffic flow on a specific interface. You control minimum bandwidth by specifying a guaranteed bandwidth (**gbw**). This means that no matter what else happens on the device, this minimum rate is guaranteed to the appropriate traffic. The maximum bandwidth (**mbw**) you set establishes the rate traffic can never exceed. By default, guaranteed bandwidth on a physical interface is the carrying capacity of the interface; therefore, you cannot set guaranteed bandwidth on the physical interface.

In the context of traffic shaping, the term virtual interfaces refers to subinterfaces bound to physical interfaces and, by extension, tunnel interfaces bound to those subinterfaces—creating a hierarchy of interfaces. A subinterface bound to a physical interface is said to be the *child* of the physical interface, its *parent*. Accordingly, a tunnel interface bound to a subinterface is the child of that subinterface, the physical interface being its *grandparent*. Figure 64 on page 204 illustrates these dependencies.

Figure 64: Interface Hierarchy



When working with virtual interfaces, bear in mind the following rules of interface hierarchy:

- Guaranteed bandwidth allocated to subinterfaces cannot be greater than the carrying capacity of the physical interface they are bound to. In Figure 64, for example, the combined **gbw** of ethernet0/1.1 and ethernet0/1.2 is 7000 Kbps, 3000 Kbps below the **mbw** of ethernet0/1. Note, however, that the combined maximum bandwidth of these two subinterfaces exceeds the carrying capacity of the physical interface they are bound to by 2000 Kbps. This is acceptable because the **mbw** keyword is used only to limit traffic to a maximum rate. If traffic falls below a maximum setting on a subinterface, that bandwidth is available to any other subinterface bound to the same physical interface.
- Guaranteed bandwidth allocated to tunnel interfaces cannot be greater than the guaranteed bandwidth of the subinterface they are bound to.
- If guaranteed bandwidth is not configured for the immediate parent, bandwidth is taken from the grandparent interface.
- Total guaranteed bandwidth of children cannot exceed parent guaranteed bandwidth.

- Child maximum bandwidth cannot exceed parent maximum bandwidth.

As already stated, you cannot configure guaranteed bandwidth on physical interfaces because guaranteed bandwidth is the same as maximum bandwidth, which is the link speed of the interface. On virtual interfaces, however, you can configure egress **gbw** and **mbw**. You can also configure ingress **mbw**, which is ingress policing at the interface level. The following command guarantees a minimum out-going bit rate of 1000 Kbps on ethernet0/4.1, and a maximum rate, both incoming and outgoing, of 2000 Kbps:

```
set interface ethernet0/4.1 bandwidth egress gbw 1000 mbw 2000 ingress mbw 2000
```

You set bandwidth in the WebUI on the Network > Interfaces > Edit page.

After setting bandwidth, you use the **get traffic-shaping interface** command to see the actual bandwidth flowing through the security device. For example, you might have traffic entering on ethernet0/1 and exiting on ethernet0/3. If you set ingress bandwidth on ethernet0/1, the command **get traffic-shaping interface ethernet0/3** will show actual throughput on the device.

If you set traffic shaping at the interface, you must also set traffic-shaping mode to on (**set traffic-shaping mode on**).

Policy-Level Traffic Shaping

You shape traffic at the policy level to allocate bandwidth for particular types of traffic. The following command guarantees a minimum 1Mbps bandwidth to FTP traffic, and drops any traffic exceeding 2 Mbps:

```
set policy from trust to untrust any any ftp permit traffic gbw 1000 pbw 2000
```

Note that this command uses the policing bandwidth (**pbw**) keyword. You can use **pbw** or **mbw** in a policy, but not both. The advantage to using **pbw** is that traffic is dropped at the ingress side of the security device, reducing throughput processing and conserving system resources. (See “Ingress Policing” on page 203.)

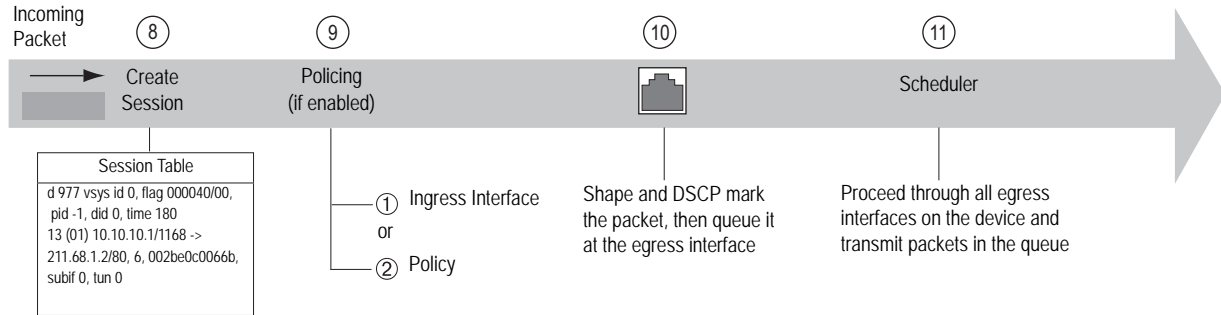
In the WebUI, after creating a policy, click the Advanced button to configure traffic-shaping parameters.

Although you must set traffic-shaping mode to **on** to shape traffic on interfaces, it is not necessary to turn on traffic shaping when shaping traffic in policies. This is because traffic-shaping mode is set to **auto** by default. When a session becomes active and policy lookup discovers traffic shaping, ScreenOS turns on traffic shaping for that session.

Packet Flow

Figure 65 illustrates the part of the packet flow through the security device that is affected by traffic shaping and policing. (See “Packet-Flow Sequence” on page 10 for a complete picture of packet flow.) Packets exceeding **pbw** (or **mbw** configured at the interface) are dropped at step 9; shaping and DSCP marking occur at step 10, and packets exceeding **mbw** (configured in a policy) are dropped at step 11.

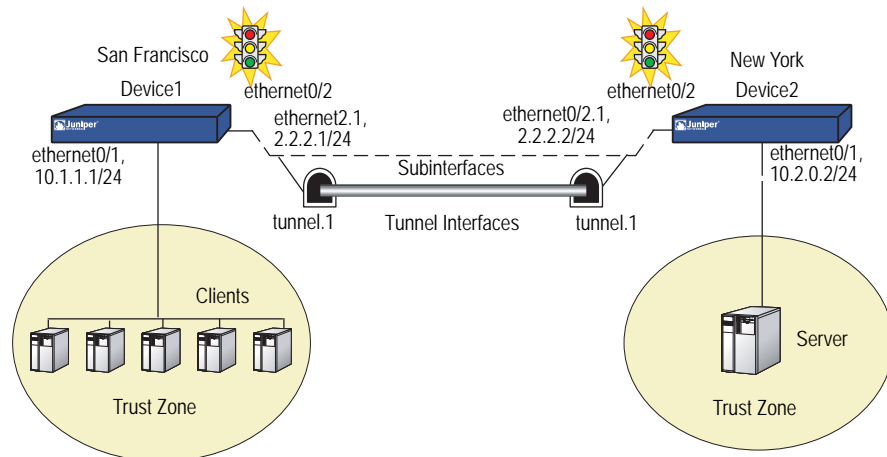
Figure 65: Traffic-Shaping Packet Flow



Example: Route-Based VPN with Ingress Policing

This example illustrates how to enforce ingress policing at the interface level for encrypted traffic. Ingress policing is configured on both the subinterface (ethernet0/2.1, maximum bandwidth:1200 Kbps) and the tunnel interface (tunnel.1, maximum bandwidth:1000 Kbps). You set the policing rate on the subinterface higher than on the tunnel interface bound to it to allow for the overhead of encryption (assuming, in this example, that all traffic received on the subinterface is meant for the tunnel interface). Policing on the subinterface is applied to the encrypted packets, while policing on the tunnel interface is applied to the decrypted inner packets. All encrypted traffic over 1200 Kbps on ethernet0/2.1 is dropped. And all decrypted (clear text) traffic over 1000 Kbps. on the tunnel.1 interface is dropped.

Figure 66: Route-Based VPN



WebUI (Configuration for Device1)**1. Interfaces**

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

IP Address/Netmask: 10.1.1.1/24
Zone: Trust

Network > Interfaces > Sub-IF > New: Enter the following, then click **Apply**:

Interface Name: (Select) ethernet0/2 and enter: 1
Zone: Untrust
IP Address/Netmask: 2.2.2.1/24
VLAN Tag: 128

Network > Interfaces > New Tunnel IF: Enter the following, then click **Apply**:

Tunnel Interface Name: 1
Zone: Untrust
Unnumbered (select) ethernet0/2.1
Interface: ethernet0/2.1

2. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 10.2.0.0/24
Interface (select): Tunnel.1

3. IKE

VPNs > AutoKey Advanced > Gateway > New: Enter the following, then click **OK**:

Gateway Name: device1_ike
Security Level: Standard
Remote Gateway Type:
Static IP Address: (select), IP Address/Hostname: 2.2.2.2
Preshared Key

Preshared Key: secret
Outgoing Interface: ethernet0/2.1

VPNs > AutoKey IKE New: Enter the following, then click **OK**:

VPN Name: device1_vpn
Gateway Name: device1_ike

> Advanced: Enter the following advanced settings, then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: (Select) Tunnel Interface, (Select) tunnel.1

CLI (Configuration for the Device1)

1. Interfaces

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.1.1.1/24
set interface ethernet0/2.1 tag 128 zone untrust
set interface tunnel.1 zone trust
set interface ethernet0/2.1 ip 2.2.2.1/24
set interface tunnel.1 ip unnumbered interface ethernet0/2.1
set route 10.2.0.0/24 int tunnel.1
```

2. IKE

```
set ike gateway device1_ike address 2.2.2.2 outgoing-interface ethernet0/2.1
  preshare sec-level standard
set vpn device1_vpn gateway 208a_ike sec-level standard
set vpn device1_vpn bind interface tunnel.1
save
```

WebUI (Configuration for Device2)

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

IP Address/Netmask: 10.2.0.2/24
Zone: Trust

Network > Interfaces > Sub-IF > New: Enter the following, then click **Apply**:

Interface Name: (Select) ethernet0/2 and enter: 1
Zone: Untrust
IP Address/Netmask: 2.2.2.2/24
VLAN Tag: 128

Network > Interfaces > Tunnel IF > New: Enter the following, then click **Apply**:

Tunnel Interface Name: 1
Unnumbered: (select) ethernet0/2.1
Interface: ethernet0/2.1

2. Bandwidth on Interfaces

Network > Interfaces > Edit (for ethernet0/2.1): Enter the following, then click **OK**:

Traffic Bandwidth, Ingress: 1200

Network > Interfaces > Edit (for tunnel.1): Enter the following, then click **OK**:

Traffic Bandwidth, Ingress: 1000

3. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 10.1.1.0/24
Interface (select): Tunnel.1

4. IKE

VPNs > AutoKey Advanced > Gateway > New: Enter the following, then click **OK**:

Gateway Name: device2_ike
 Security Level: Standard
 Remote Gateway Type:
 Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key

Preshared Key: secret
 Outgoing Interface: ethernet0/2.1

VPNs > AutoKey IKE New: Enter the following, then click **OK**:

VPN Name: device2_vpn
 Gateway Name: device2_ike

> Advanced: Enter the following advanced settings, then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: (Select) Tunnel Interface, (Select) tunnel.1

5. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Service: Any
 Action: Permit

Policy > Policies > (From: Untrust, To: Trust) New: Enter the following, then click **OK**:

Service: Any
 Action: Permit

CLI (Configuration for the Device2)**1. Interfaces**

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 10.2.0.2/24
set interface ethernet0/2.1 tag 128 zone untrust
set interface ethernet0/2.1 ip 2.2.2.2/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet0/2.1
set route 10.1.1.0/24
```

2. Bandwidth on interfaces

```
set interface ethernet0/2.1 bandwidth ingress mbw 1200
set interface tunnel.1 bandwidth ingress mbw 1000
```

3. IKE

```
set ike gateway device2_ike address 2.2.2.1 preshare secret sec-level standard
set vpn device2_vpn gateway 208b_ike sec-level standard
set vpn device2_vpn bind interface tunnel.1
```


4. Policy

```
set policy from trust to untrust any any any permit
set policy from untrust to trust any any any permit
save
```

Example: Policy-Based VPN with Ingress Policing

This example illustrates how to enforce ingress policing at both the interface level and in policies. On the ethernet0/1 interface on *Device1*, you set the ingress maximum bandwidth at 20000 Kbps. With this setting, all traffic over 20000 Kbps from clients connected to *Device1* on the ethernet0/1 interface, is dropped. Ingress policing at the interface applies to all the traffic that arrives on that interface. For finer granularity, you can apply ingress policing at the policy level. In this example, you create policies to restrict all ingress FTP protocol traffic on *Device1* by creating policies between the trust and untrust zones, and set the policing bandwidth to 5000 Kbps. All FTP traffic over 5000 Kbps from the trust zone to the untrust zone is dropped.

Figure 67: Policy-Based VPN



WebUI (Configuration for Device1)

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

```
IP Address/Netmask: 10.1.1.1/24
Zone: Trust
Interface mode: (select) NAT
```

Network > Interfaces > Edit (for ethernet0/2): Enter the following, then click **OK**:

```
IP Address/Netmask: 2.1.1.1/24
Zone: Untrust
Interface mode: (select) Route
```


2. IKE VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, then click **OK**:

Gateway Name: device2_ike
 Security Level: Standard
 Remote Gateway Type:
 Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key

Preshared Key: secret
 Outgoing Interface: ethernet0/2

> Advanced: Enter the following advanced settings, then click **OK** to return to basic Gateway configuration page:

Phase 1 Proposal: pre-g2-3des-sha

VPNs > AutoKey IKE New: Enter the following, then click **OK**:

VPN Name: device2_vpn
 Gateway Name: device2_ike

3. Interface-Based Policing

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Traffic Bandwidth, Ingress: 20000

4. Routing

Network > Routing > Destination > New: Enter the following, then click **OK**:

Network IP Address/Netmask: 10.2.1.0/24
 Interface: (select), ethernet0/2
 Gateway IP Address: 2.2.2.2

5. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: 1
 Service: FTP
 Action: Tunnel
 Tunnel VPN: (select), device2_vpn
 Modify matching bidirectional VPN policy: (select)

> Advanced: Enter the following advanced settings, then click **Return** to return to the basic Policies configuration page:

Traffic Shaping (select) Policing Bandwidth: 5000

CLI (Configuration for Device1)**1. Interfaces**

```
set interface ethernet0/1 zone trust
set interface ethernet0/2 zone untrust
set interface ethernet0/1 ip 10.1.1.1/24
```

```
set interface ethernet0/1 nat
set interface ethernet0/2 ip 2.1.1.1/24
set interface ethernet0/1 route
```

2. IKE VPN

```
set ike gateway device2_ike address 2.2.2.2 main outgoing interface ethernet0/2
  preshare secret proposal pre-g2-3des-sha
set vpn device2_vpn gateway device2_ike no-replay tunnel idletime 0 sec-level
  standard
```

3. Routing

```
set route 10.2.1.0/24 interface ethernet0/2 gateway 2.2.2.2
```

4. Policies

```
set policy from trust to untrust any any ftp tunnel vpn device2_vpn pair-policy 2
  traffic pbw 5000
set policy from untrust to trust any any ftp tunnel vpn netscreen2_vpn pair-policy 1
  traffic pbw 5000
```

5. Interface-Based Policing

```
set interface ethernet0/1 bandwidth ingress mbw 20000
save
```

WebUI (Configuration for Device2)

1. Interfaces

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

```
IP Address/Netmask: 1.1.1.1/24
Zone: Trust
Interface mode: (select) Route
```

Network > Interfaces > Edit (for ethernet0/2): Enter the following, then click **OK**:

```
IP Address/Netmask: 10.2.2.1/24
Zone: Untrust
Interface mode: (select) NAT
```

2. IKE VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, then click **OK**:

```
Gateway Name: device1_ike
Security Level: Standard
Remote Gateway Type:
  Static IP Address: (select), IP Address/Hostname: 2.1.1.1
```

3. Preshared Key

Preshared Key: secret
 Outgoing Interface: ethernet0/2

> Advanced: Enter the following advanced settings, then click **OK** to return to basic Gateway configuration page:

Phase 1 Proposal: pre-g2-3des-sha

VPNs > AutoKey IKE New: Enter the following, then click **OK**:

VPN Name: device1_vpn
 Gateway Name: device1_ike

4. Routing

Network > Routing > Destination > New: Enter the following, then click **OK**:

Network IP Address/Netmask: 10.1.1.0/24
 Interface: (select), ethernet0/2
 Gateway IP Address: 1.1.1.1

5. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Name: 1
 Service: FTP
 Action: Tunnel
 Tunnel VPN: (select), device1_vpn
 Modify matching bidirectional VPN policy: (select)

CLI (Configuration for Device2)**1. Interfaces**

```
set interface ethernet0/1 1.1.1.2/24
set interface ethernet0/1 route
set interface ethernet0/2 ip 10.2.2.1/24
set interface ethernet0/2 nat
```

2. IKE VPN

```
set ike gateway device1_ike address 2.1.1.1 main outgoing interface ethernet0/2
  preshare secret proposal pre-g2-3des-sha
set vpn device1_vpn gateway device1_ike no-replay tunnel idletime 0 sec-level
  standard
```

3. Routing

```
set route 10.1.1.0/24 interface ethernet0/1 gateway 1.1.1.1
```

4. Policies

```
set policy id 1 from trust to untrust any any ftp tunnel vpn device1_vpn pair-policy 2
set policy id 2 from untrust to trust any any ftp tunnel vpn device1_vpn pair-policy 1
save
```

Traffic Shaping Using a Loopback Interface

Traffic shaping is not supported on loopback interfaces, because no traffic is actually transmitted on a loopback interface. However, a loopback interface is often used as an anchor point (for example in the case of a VPN, to derive the source IP address), while the data is transmitted on an actual egress interface. When using a loopback interface in a VPN, therefore, you configure traffic shaping on the outgoing interface. ScreenOS then associates the session with the real outgoing interface, which it deduces from the routing table, dynamically updating the association as the routing table changes.

DSCP Marking and Shaping

As stated earlier in this chapter, Differentiated Services (DS) is a system for tagging (or “marking”) traffic at a position within a hierarchy of priority. Differentiated Services Codepoint (DSCP) marking maps the ScreenOS priority level of the policy to the first three bits of codepoint in the DS field in the IP packet header. (See “Setting Service Priorities” on page 198 for more information).

You can shape traffic in a policy that uses DSCP marking, or you can use DSCP marking independent of traffic shaping. Traffic shaping governs how traffic is processed on the security device and can be configured at the interface level or in policies. DSCP marking, which you set at the policy level, governs how traffic is processed by downstream routers.

NOTE: Some devices require that you explicitly enable DSCP marking by setting a system-wide environmental variable. Refer to the installation and configuration manual for your device to find out if it requires that you explicitly enable DSCP marking before using it in policies. If your device requires it, use the following command to enable DSCP marking system wide: **set envar ipsec-dscp-mark= yes**. This variable cannot be set using the WebUI. Use the **unset envar ipsec-dscp-mark** to disable DSCP marking system wide.

The DSCP marking feature is disabled in IDP-capable security devices. For information about IDP-capable security devices, see “Intrusion Detection and Prevention” on page 4-173.

If you specify DSCP marking in a policy but do not set a value, ScreenOS maps the policy priority to an equivalent IP precedence priority in the DSCP system. It does this by overwriting the first 3 bits in the ToS byte with the IP precedence priority. For example, if you create a policy that gives all traffic a priority of, for example, 2 (0 is the highest priority), and you enable DSCP marking, ScreenOS queues traffic for that policy with level 2 priority at the egress interface and marks it with an equivalent IP precedence priority. The following command creates a policy that gives priority 2 to all traffic, and enables DSCP marking:

```
set policy from trust to untrust any any any permit traffic priority 2 dscp enable
```

But if you give DSCP a *dscp-byte value* of, for example, 46 (the highest priority), the security device still queues traffic at the egress interface at priority 2 but overwrites the first 6 bits of the ToS byte with the DSCP value.

```
set policy from trust to untrust any any permit traffic priority 2 dscp enable
value 46
```

DSCP marking is supported on all platforms and can be configured with traffic shaping or independently. Table 31 shows how DSCP marking works for clear packets in policies, Figure 66 shows how DSCP marking works for clear packets in policy-based VPNs, Table 33 shows how DSCP marking works for clear packets in route-based VPNs.

Table 31: DSCP Marking for Clear-Text Traffic

Description	Action
Clear packet with no marking on the policy.	No marking.
Clear packet with marking on the policy.	The packet is marked based on the policy.
Premarked packet with no marking on the policy.	Retain marking in the packet.
Premarked packet with marking on the policy.	Overwrite marking in the packet based on the policy.

Table 32: DSCP Marking for Policy-Based VPNs

Description	Action
Clear packet into policy-based VPN with no marking on the policy.	No marking.
Clear packet into policy-based VPN with marking on the policy	The inner packet and IP header of the ESP are both marked, based on the policy.
Premarked packet into policy-based VPN with no marking on the policy.	Copy the inner packet marking to the IP header of the ESP, retain marking in the inner packet.
Premarked packet into policy-based VPN with marking on the policy.	Overwrite the marking in the inner packet based on the policy, and copy the inner packet marking to the IP header of the ESP.

Table 33: DSCP Marking for Route-Based VPNs

Description	Action
Clear packet into route-based VPN with no marking on the policy.	No marking.
Clear packet into route-based VPN with marking on the policy.	The inner packet and IP header of the ESP are both marked, based on the policy.
Premarked packet into route-based VPN with no marking on the policy.	Copy the inner packet marking to the IP header of the ESP, retain marking in the inner packet.
Premarked packet into route-based VPN with marking on the policy.	Overwrite the marking in the inner packet based on the policy, and copy the inner packet marking to the IP header of the ESP.

Chapter 8

System Parameters

This chapter focuses on the concepts involved in establishing system parameters affecting the following areas of a security appliance. It contains the following sections:

- “Domain Name System Support” on page 217
- “Dynamic Host Configuration Protocol” on page 225
- “Point-to-Point Protocol over Ethernet” on page 243
- “License Keys” on page 250
- “Registration and Activation of Subscription Services” on page 251
- “System Clock” on page 253

Domain Name System Support

The Juniper Networks security device incorporates Domain Name System (DNS) support, which allows you to use domain names as well as IP addresses for identifying locations. A DNS server keeps a table of the IP addresses associated with domain names. Using DNS makes it possible to reference locations by domain name (such as `www.juniper.net`) in addition to using the routable IP address, which for `www.juniper.net` is `207.17.137.68`. DNS translation is supported in all the following programs:

- Address Book
- Syslog
- Email
- WebTrends
- Websense
- LDAP
- SecurID

- RADIUS
- NetScreen-Security Manager

Before you can use DNS for domain name/address resolution, you must enter the addresses for DNS servers in the security device.

NOTE: When enabling the security device as a Dynamic Host Configuration Protocol (DHCP) server (see “Dynamic Host Configuration Protocol” on page 225), you must also enter the IP addresses for DNS servers in the DHCP page on the WebUI or through the **set interface *interface* dhcp** command in the CLI.

DNS Lookup

The security device refreshes all the entries in its DNS table by checking them with a specified DNS server at the following times:

- After an HA failover occurs
- At a regularly scheduled time of day and at regularly scheduled intervals throughout the day
- When you manually command the device to perform a DNS lookup
 - WebUI: Network > DNS > Host: Click **Refresh**.
 - CLI: **exec dns refresh**

In addition to the existing method of setting a time for a daily automatic refresh of the DNS table, you can also define an interval of time from 4 to 24 hours.

NOTE: When you add a fully qualified domain name (FQDN) such as an address or IKE gateway through the WebUI, the security device resolves it when you click **Apply** or **OK**. When you type a CLI command that references an FQDN, the security device attempts to resolve it when you enter it.

When the security device connects to the DNS server to resolve a domain name/IP address mapping, it stores that entry in its DNS status table. The following list contains some of the details involved in a DNS lookup:

- When a DNS lookup returns multiple entries, the address book accepts all entries. The other programs listed on page 217 accept only the first one.
- The security device reinstalls all policies if it finds that anything in the domain name table has changed when you refresh a lookup using the **Refresh** button in the WebUI or enter the **exec dns refresh** CLI command.
- If a DNS server fails, the security device looks up everything again.

- If a lookup fails, the security device removes it from the cache table.
- If the domain name lookup fails when adding addresses to the address book, the security device displays an error message stating that you have successfully added the address but the DNS name lookup failed.

The security device must do a new lookup once a day, which you can schedule it to do at a specified time.

WebUI

Network > DNS > Host: Enter the following, then click **Apply**:

DNS refresh every day at: Select checkbox and enter time <hh:mm>

CLI

```
set dns host schedule time_str
save
```

DNS Status Table

The DNS status table reports all the domain names looked up, their corresponding IP addresses, whether the lookup was successful, and when each domain name/IP address was last resolved.

Table 34: DNS Status Table

Name	IP Address	Status	Last Lookup
www.yahoo.com	204.71.200.74	Success	8/13/2000 16:45:33
	204.71.200.75		
	204.71.200.67		
	204.71.200.68		
www.hotbot.com	209.185.151.28	Success	8/13/2000 16:45:38
	209.185.151.210		
	216.32.228.18		

To view the DNS status table, do either of the following:

WebUI

Network > DNS > Host > Show DNS Lookup Table


CLI

```
get dns host report
```

Setting the DNS Server and Refresh Schedule

To implement DNS functionality, the IP addresses for the DNS servers at 24.1.64.38 and 24.0.0.3 are entered in the security device, protecting a single host in a home office. The security device is scheduled to refresh the DNS settings stored in its DNS status table every day at 11:00 P.M.

Figure 68: DNS Refresh



WebUI

Network > DNS > Host: Enter the following, then click **Apply**:

Primary DNS Server: 24.0.0.3
 Secondary DNS Server: 24.1.64.38
 DNS Refresh: (select)
 Every Day at: 23:00

CLI

```
set dns host dns1 24.0.0.3
set dns host dns2 24.1.64.38
set dns host schedule 23:00
save
```

Setting a DNS Refresh Interval

In this example, you configure the security device to refresh its DNS table every 4 hours beginning at 00:01 AM every day.

WebUI

Network > DNS > Host: Enter the following, then click **Apply**:

DNS Refresh: (select)
 Every Day at: 00:01
 Interval: 4

CLI

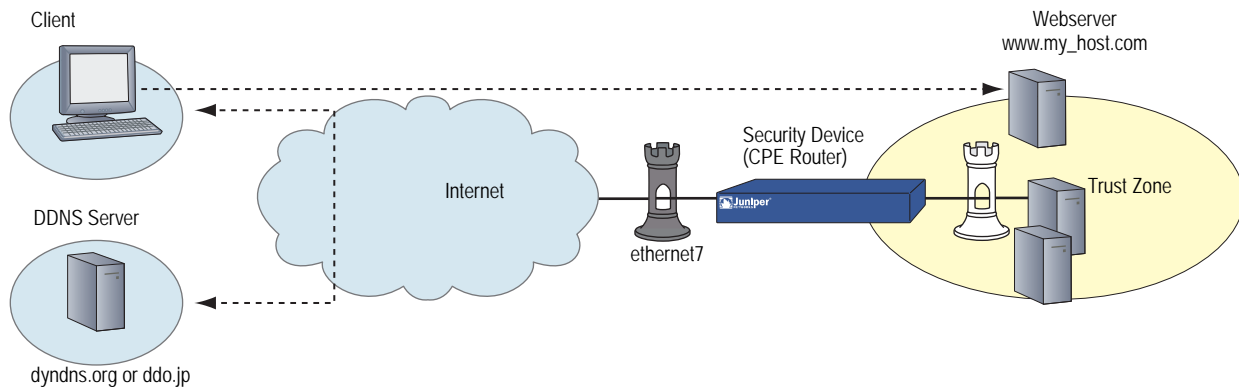
```
set dns host schedule 00:01 interval 4
save
```

Dynamic Domain Name System

Dynamic Domain Name System (DDNS) is a mechanism that allows clients to dynamically update IP addresses for registered domain names. This update is useful when an ISP uses PPP, DHCP, or XAuth to dynamically change the IP address for a CPE router (such as a security device) that protects a web server. Clients from the Internet can access the web server using a domain name, even if the IP address of the CPE router previously changed. This change is made possible by a DDNS server such as dyndns.org or ddo.jp, which contains the dynamically-changed addresses and their associated domain names. The CPE updates the DDNS servers with this information, periodically or in response to IP address changes.

To use DDNS, create an account (username and password) on the DDNS server. The server uses this account information to configure the client device.

Figure 69: Dynamic DNS



In Figure 69, it is possible that the IP address for interface ethernet7 might change. When a change happens, the client can still access the protected web server using the host name (www.my_host.com), either through the dyndns.org server or the ddo.jp server. Each of these servers require different configurations on the security device.

Setting Up DDNS for a DynDNS Server

In the following example, you configure a security device for DDNS operation. The device uses the dyndns.org server to resolve changed addresses. For this server, you specify the protected host using the Host Name setting, which explicitly binds to the DNS interface (ethernet7). When the device sends an update to the ddo.jp server, it associates the Host Name with the IP address of the interface.

WebUI

Network > DNS > DDNS > New: Enter the following, then click **OK**:

```
ID: 12
Server Settings
  Server Type: dyndns
  Server Name: dyndns.org
  Refresh Interval: 24
  Minimum Update Interval: 15
Account Settings
  Username: swordfish
  Password: ad93lvb
Bind to Interface: ethernet7
Host Name: www.my_host.com
```

NOTE: Minimum Update Interval specifies the minimum time interval (expressed in minutes) between DDNS updates. The default is 10 minutes, and the allowable range is 1-1440. In some cases, the device might not update the interval because the DNS server first needs to timeout the DDNS entry from its cache. In addition, if you set the Minimum Update Interval to a low value, then the security device might lock you out. The recommended value is 10 minutes or greater.

CLI

```
set dns ddns
set dns ddns enable
```

```

set dns ddns id 12 server dyndns.org server-type dyndns refresh-interval 24
  minimum-update-interval 15
set dns ddns id 12 src-interface ethernet7 host-name www.my_host.com
set dns ddns id 12 username swordfish password ad93lvb
save

```

Setting Up DDNS for a DDO Server

In the following example, you configure a security device for DDNS. The device uses the ddo.jp server to resolve addresses. For the ddo.jp server, you specify the protected host FQDN as the Username for the DDNS entry, instead of specifying the protected host using the Host Name setting. The service automatically derives the host name from the Username value. For example, the ddo.jp server translates a username of my_host to my_host.ddo.jp. You need to make sure that the registered domain name on ddo.jp matches the derived DNS.

WebUI

Network > DNS > DDNS > New: Enter the following, then click **OK**:

```

ID: 25
Server Settings
  Server Type: ddo
  Server Name: juniper.net
  Refresh Interval: 24
  Minimum Update Interval: 15
Account Settings
  Username: my_host
  Password: ad93lvb
Bind to Interface: ethernet7

```

CLI

```

set dns ddns
set dns ddns enable
set dns ddns id 25 server ddo.jp server-type ddo refresh-interval 24
  minimum-update-interval 15
set dns ddns id 25 src-interface ethernet7
set dns ddns id 25 username my_host password ad93lvb
save

```

Proxy DNS Address Splitting

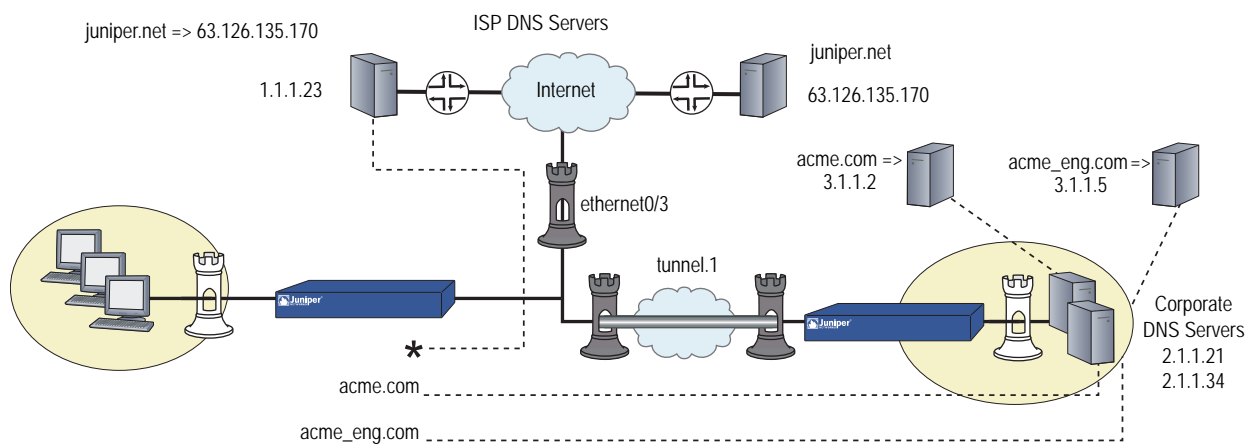
The proxy DNS feature provides a transparent mechanism that allows clients to make split DNS queries. Using this technique, the proxy selectively redirects the DNS queries to specific DNS servers, according to partial or complete domain names. This is useful when multiple VPN tunnels or PPPoE virtual links provide network connectivity and it is necessary to direct some DNS queries to one network and other queries to another network.

A DNS proxy has the following advantages:

- Domain lookups are usually more efficient. For example, DNS queries meant for the corporate domain (such as acme.com) could go to the corporate DNS server exclusively, while all others go to the ISP DNS server, which reduces the load on the corporate server. This can also prevent corporate domain information from leaking into the Internet.
- DNS proxy allows you to transmit selected DNS queries through a tunnel interface, which prevents malicious users from learning about the internal configuration of a network. For example, DNS queries bound for the corporate server can pass through a tunnel interface to use security features such as authentication, encryption, and anti-replay.

The following commands create two proxy-DNS entries that selectively forward DNS queries to different servers.

Figure 70: Splitting DNS Requests



- Any DNS query with a FQDN containing the domain name `acme.com` goes out through tunnel interface `tunnel.1`, to the corporate DNS server at IP address `2.1.1.21`.

For example, if a host sends a DNS query for `www.acme.com`, the device automatically directs the query to this server. (Let's assume that the server resolves the query to IP address `3.1.1.2`.)

- Any DNS query with a FQDN containing the domain name `acme_eng.com` goes out through tunnel interface `tunnel.1` to the DNS server at IP address `2.1.1.34`.

For example, if a host sends a DNS query for `intranet.acme_eng.com`, the device directs the query to this server. (Let's assume that the server resolves the query to IP address `3.1.1.5`.)

- All other DNS queries (denoted by an asterisk) bypass the corporate servers and go out through interface ethernet0/3 to the DNS server at IP address 1.1.1.23.

For example, if the host and domain name is www.juniper.net, the device automatically bypasses the corporate servers and directs the query to this server, which resolves the query to IP address 207.17.137.68.

WebUI

1. **Network > DNS > Proxy: Enter the following, then click Apply:**
 Initialize DNS Proxy: Enable
 Enable DNS Proxy: Enable
2. **Network > DNS > Proxy > New: Enter the following, then click OK:**
 Domain Name: acme.com
 Outgoing Interface: tunnel.1
 Primary DNS Server: 2.1.1.21
3. **Network > DNS > Proxy > New: Enter the following, then click OK:**
 Domain Name: acme_eng.com
 Outgoing Interface: tunnel.1
 Primary DNS Server: 2.1.1.34
4. **Network > DNS > Proxy > New: Enter the following, then click OK:**
 Domain Name: *
 Outgoing Interface: ethernet0/3
 Primary DNS Server: 1.1.1.23

CLI

```
set dns proxy
set dns proxy enable
set interface ethernet0/3 proxy dns
set dns server-select domain acme.com outgoing-interface tunnel.1 primary-server
  2.1.1.21
set dns server-select domain acme_eng.com outgoing-interface tunnel.1
  primary-server 2.1.1.34
set dns server-select domain * outgoing-interface ethernet0/3 primary-server
  1.1.1.23
save
```

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) was designed to reduce the demands on network administrators by automatically assigning the TCP/IP settings for the hosts on a network. Instead of requiring administrators to assign, configure, track, and change (when necessary) all the TCP/IP settings for every machine on a network, DHCP does it all automatically. Furthermore, DHCP ensures that duplicate addresses are not used, reassigns unused addresses, and automatically assigns IP addresses appropriate for the subnet on which a host is connected.

Different security devices support the different DHCP roles described in Table 35.

Table 35: DHCP Roles

Role	Description
DHCP Client	Some security devices can act as DHCP clients, receiving a dynamically assigned IP address for any physical interface in any zone.
DHCP Server	Some security devices can also act as DHCP servers, allocating dynamic IP addresses to hosts (acting as DHCP clients) on any physical or VLAN interface in any zone. Note: While using the DHCP server module to assign addresses to hosts such as workstations in a zone, you can still use fixed IP addresses for other machines such as mail servers and WINS servers.
DHCP Relay Agent	Some security devices can also act as DHCP relay agents, receiving DHCP information from a DHCP server and relaying that information to hosts on any physical or VLAN interface in any zone.
DHCP Client/Server/Relay Agent	Some security devices can simultaneously act as a DHCP client, server, and relay agent. You can only configure one DHCP role on a single interface. For example, you cannot configure the DHCP client and server on the same interface. Optionally, you can configure the DHCP client module to forward TCP/IP settings that it receives to the DHCP server module, for use when providing TCP settings to hosts in the Trust zone acting as DHCP clients.

DHCP consists of two components: a protocol for delivering host-specific TCP/IP configuration settings and a mechanism for allocating IP addresses. When the security device acts as a DHCP server, it provides the following TCP/IP settings to each host when that host boots up:

- Default gateway IP address and netmask. If you leave these settings as 0.0.0.0/0, the DHCP server module automatically uses the IP address and netmask of the default Trust zone interface.

NOTE: On devices that can have multiple interfaces bound to the Trust zone, the default interface is the first interface bound to that zone and assigned an IP address.

- The IP addresses of the following servers:
 - WINS servers (2): A Windows Internet Naming Service (WINS) server maps a NetBIOS name used in a Windows NT network environment to an IP address used on an IP-based network. The number in parentheses indicates the number of servers supported.
 - NetInfo servers (2): NetInfo is an Apple network service used for the distribution of administrative data within a LAN.
 - NetInfo tag (1): The identifying tag used by the Apple NetInfo database.
 - DNS servers (3): A Domain Name System (DNS) server maps a uniform resource locator (URL) to an IP address.
 - SMTP server (1): A Simple Mail Transfer Protocol (SMTP) server delivers SMTP messages to a mail server, such as a POP3 server, which stores the incoming mail.
 - POP3 server (1): A Post Office Protocol version 3 (POP3) server stores incoming mail. A POP3 server must work conjointly with an SMTP server.
 - News server (1): A news server receives and stores postings for news groups.

NOTE: If a DHCP client to which the security device is passing the above parameters has a specified IP address, that address overrides all the dynamic information received from the DHCP server.

Configuring a DHCP Server

A security device can support up to eight DHCP servers on any physical or VLAN interface in any zone. When acting as a DHCP server, a security device allocates IP addresses and subnet masks in two modes:

- In Dynamic mode, the security device, acting as a DHCP server, assigns (or “leases”) an IP address from an address pool to a host DHCP client. The IP address is leased for a determined period of time or until the client relinquishes the address. (To define an unlimited lease period, enter 0.)
- In Reserved mode, the security device assigns a designated IP address from an address pool exclusively to a specific client every time that client goes online.

NOTE: An address pool is a defined range of IP addresses within the same subnet from which the security device can draw DHCP address assignments. You can group up to 255 IP addresses.

The DHCP server supports up to 64 entries, which can include both single IP addresses and IP address ranges, for dynamic and reserved IP addresses.

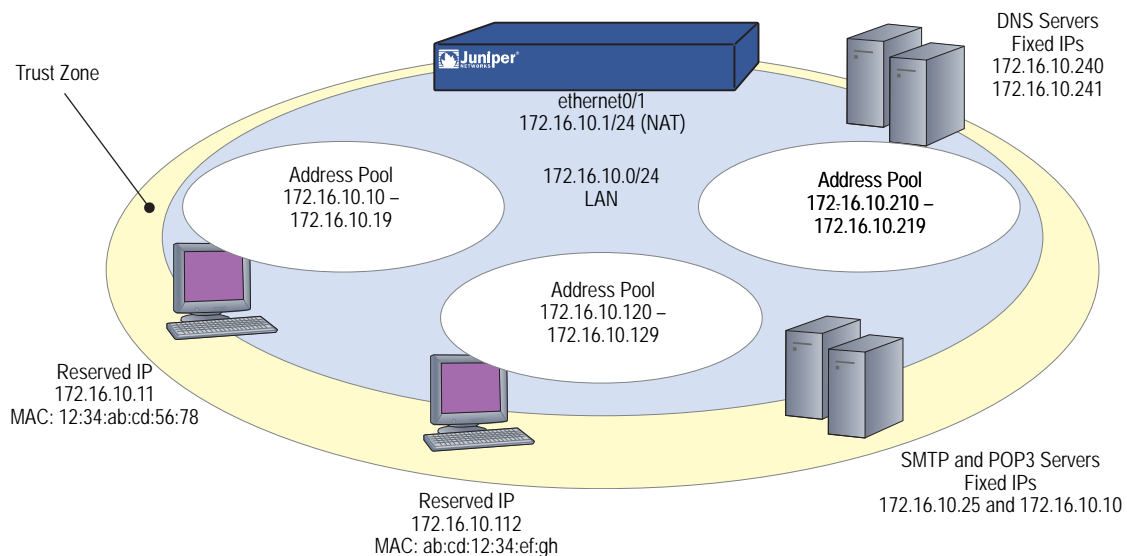
The security device saves every IP address assigned through DHCP in flash memory. Consequently, rebooting the security device does not affect address assignments.

In this example, using DHCP, the 172.16.10.0/24 network in the Trust zone is sectioned into three IP address pools.

- 172.16.10.10 through 172.16.10.19
- 172.16.10.120 through 172.16.10.129
- 172.16.10.210 through 172.16.10.219

The DHCP server assigns all IP addresses dynamically, except for two workstations with reserved IP addresses and four servers with static IP addresses. The interface ethernet0/1 is bound to the Trust zone, has IP address 172.16.10.1/24, and is in NAT mode. The domain name is dynamic.com.

Figure 71: Device as DHCP Server



WebUI**1. Addresses**

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: DNS#1
 Comment: Primary DNS Server
 IP Address/Domain Name:
 IP/Netmask: (select), 172.16.10.240/32
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: DNS#2
 Comment: Secondary DNS Server
 IP Address/Domain Name:
 IP/Netmask: (select), 172.16.10.241/32
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: SMTP
 Comment: SMTP Server
 IP Address/Domain Name:
 IP/Netmask: (select), 172.16.10.25/32
 Zone: Trust

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: POP3
 Comment: POP3 Server
 IP Address/Domain Name:
 IP/Netmask: (select), 172.16.10.110/32
 Zone: Trust

2. DHCP Server

Network > DHCP > Edit (for ethernet0/1) > DHCP Server: Enter the following, then click **Apply**:

Lease: Unlimited (select)
 WINS#1: 0.0.0.0
 DNS#1: 172.16.10.240

NOTE: If you leave the Gateway and Netmask fields as **0.0.0.0**, the DHCP server module sends the IP address and netmask set for ethernet0/1 to its clients (172.16.10.1 and 255.255.255.0 in this example). However, if you enable the DHCP client module to forward TCP/IP settings to the DHCP server module (see “Propagating TCP/IP Settings” on page 240), then you must manually enter **172.16.10.1** and **255.255.255.0** in the Gateway and Netmask fields, respectively.

> Advanced Options: Enter the following, then click **OK** to set the advanced options and return to the basic configuration page:

```
WINS#2: 0.0.0.0
DNS#2: 172.16.10.241
DNS#3: 0.0.0.0
SMTP: 172.16.10.25
POP3: 172.16.10.110
NEWS: 0.0.0.0
NetInfo Server #1: 0.0.0.0
NetInfo Server #2: 0.0.0.0
NetInfo Tag: (leave field empty)
Domain Name: dynamic.com
```

> Addresses > New: Enter the following, then click **OK**:

```
Dynamic: (select)
IP Address Start: 172.16.10.10
IP Address End: 172.16.10.19
```

> Addresses > New: Enter the following, then click **OK**:

```
Dynamic: (select)
IP Address Start: 172.16.10.120
IP Address End: 172.16.10.129
```

> Addresses > New: Enter the following, then click **OK**:

```
Dynamic: (select)
IP Address Start: 172.16.10.210
IP Address End: 172.16.10.219
```

> Addresses > New: Enter the following, then click **OK**:

```
Reserved: (select)
IP Address: 172.16.10.11
Ethernet Address: 1234 abcd 5678
```

> Addresses > New: Enter the following, then click **OK**:

```
Reserved: (select)
IP Address: 172.16.10.112
Ethernet Address: abcd 1234 efgh
```

CLI

1. Addresses

```
set address trust dns1 172.16.10.240/32 "primary dns server"
set address trust dns2 172.16.10.241/32 "secondary dns server"
set address trust snmp 172.16.10.25/32 "snmp server"
set address trust pop3 172.16.10.110/32 "pop3 server"
```

2. DHCP Server

```
set interface ethernet0/1 dhcp server option domainname dynamic.com
set interface ethernet0/1 dhcp server option lease 0
set interface ethernet0/1 dhcp server option dns1 172.16.10.240
set interface ethernet0/1 dhcp server option dns2 172.16.10.241
set interface ethernet0/1 dhcp server option smtp 172.16.10.25
set interface ethernet0/1 dhcp server option pop3 172.16.10.110
```

```

set interface ethernet0/1 dhcp server ip 172.16.10.10 to 172.16.10.19
set interface ethernet0/1 dhcp server ip 172.16.10.120 to 172.16.10.129
set interface ethernet0/1 dhcp server ip 172.16.10.210 to 172.16.10.219
set interface ethernet0/1 dhcp server ip 172.16.10.11 mac 1234abcd5678
set interface ethernet0/1 dhcp server ip 172.16.10.112 mac abcd1234efgh
set interface ethernet0/1 dhcp server service
save
    
```

NOTE: If you do not set an IP address for the gateway or a netmask, the DHCP server module sends its clients the IP address and netmask for ethernet0/1 (172.16.10.1 and 255.255.255.0 in this example). However, if you enable the DHCP client module to forward TCP/IP settings to the DHCP server module (see “Propagating TCP/IP Settings” on page 240), then you must manually set these options: **set interface ethernet0/1 dhcp server option gateway 172.16.10.1** and **set interface ethernet0/1 dhcp server option netmask 255.255.255.0**.

Customizing DHCP Server Options

When you specify DHCP servers for an interface, you might need to specify options that identify the servers or provide information used by the servers. For example, you can specify the IP address of the primary and secondary DNS servers, or set the IP address lease time.

The following are predefined DHCP services, as described in RFC 2132, *DHCP Options and BOOTP Vendor Extensions*.

Table 36: Predefined DHCP Services

Terminology	ScreenOS CLI Terminology	Option Code
Subnet Mask	netmask	1
Router Option	gateway	3
Domain Name System (DNS) server	dns1, dns2, dns3	6
Domain Name	domainname	15
NetBIOS over TCP/IP Name Server Option	wins1, wins2	44
IP Address Lease Time	lease	51
SMTP Server Option	smtp	69
POP3 Server Option	pop3	70
NNTP Server Option	news	71
(N/A)	nis1, nis2	112
(N/A)	nistag	113

In situations where the predefined server options are inadequate, you can define custom DHCP server options. For example, for certain Voice-over IP (VoIP) configurations, it is necessary send extra configuration information, which is not supported by predefined server options. In such cases, you must define suitable custom options.

In the following example, you create DHCP server definitions for IP phones which act as DHCP clients. The phones use the following custom options:

- Option code 444, containing string “Server 4”
- Option code 66, containing IP address 1.1.1.1
- Option code 160, containing integer 2004

CLI

1. Addresses

```
set address trust dns1 172.16.10.240/32 "primary dns server"
set address trust dns2 172.16.10.241/32 "secondary dns server"
```

2. DHCP Server

```
set interface ethernet0/1 dhcp server option domainname dynamic.com
set interface ethernet0/1 dhcp server option lease 0
set interface ethernet0/1 dhcp server option dns1 172.16.10.240
set interface ethernet0/1 dhcp server option dns2 172.16.10.241
set interface ethernet0/1 dhcp server option custom 444 string "Server 4"
set interface ethernet0/1 dhcp server option custom 66 ip 1.1.1.1
set interface ethernet0/1 dhcp server option custom 160 integer 2004
set interface ethernet0/1 dhcp server ip 172.16.10.10 to 172.16.10.19
```

Placing the DHCP Server in an NSRP Cluster

When the primary unit in a redundant NSRP cluster functions as a DHCP server, all members in the cluster maintain all DHCP configurations and IP address assignments. In the event of a failover, the new primary unit maintains all the DHCP assignments. However, termination of HA communication disrupts synchronization of existing DHCP assignments among the cluster members. After restoring HA communication, you can resynchronize the DHCP assignments by using the following CLI command on both units in the cluster: **set nsrp rto-mirror sync**.

DHCP Server Detection

When a DHCP server on a security device starts up, the system can first check to see if there is already a DHCP server on the interface. ScreenOS automatically stops the local DHCP server process from starting if another DHCP server is detected on the network. To detect another DHCP server, the device sends out DHCP boot requests at two-second intervals. If the device does not receive any response to its boot requests, it then starts its local DHCP server process.

If the security device receives a response from another DHCP server, the system generates a message indicating that the DHCP service is enabled on the security device but not started because another DHCP server is present on the network. The log message includes the IP address of the existing DHCP server.

You can set one of three operational modes for DHCP server detection on an interface: Auto, Enable, or Disable. Auto mode causes the security device to always check for an existing DHCP server at bootup. You can configure the device to not attempt to detect another DHCP server on an interface by setting the security DHCP server to Enable or Disable mode. In Enable mode, the DHCP server is always on and the device does not check if there is an existing DHCP server on the network. In Disable mode, the DHCP server is always off.

Enabling DHCP Server Detection

In this example, you set the DHCP server on the ethernet0/1 interface to check for an existing DHCP server on the interface first before starting up.

WebUI

Network > DHCP > Edit (for ethernet0/1) > DHCP Server: Enter the following, then click **OK**:

Server Mode: Auto (select)

CLI

```
set interface ethernet0/1 dhcp server auto
save
```

Disabling DHCP Server Detection

In this example, you set the DHCP server on the ethernet0/1 interface to start up without checking to see if there is an existing DHCP server on the network.

WebUI

Network > DHCP > Edit (for ethernet0/1) > DHCP Server: Enter the following, then click **OK**:

Server Mode: Enable (select)

CLI

```
set interface ethernet0/1 dhcp server enable
save
```

NOTE: Issuing the CLI command **set interface** interface **dhcp server service** command activates the DHCP server. If the DHCP server detection mode for the interface is set to Auto, the DHCP server on the security device starts only if it does not find an existing server on the network. Issuing the **unset interface** interface **dhcp server service** command disables the DHCP server on the security device and also deletes any existing DHCP configuration.

Assigning a Security Device as a DHCP Relay Agent

When acting as a DHCP relay agent, the security device forwards DHCP requests and assignments between DHCP clients directly attached to one interface and one or more DHCP servers accessible through another interface. The clients and servers may be in the same security zone or in separate security zones.

You can configure a DHCP relay agent on one or more physical or VLAN interfaces on a security device, but you cannot configure a DHCP relay agent and DHCP server or client functions on the same interface.

When the security device functions as a DHCP relay agent, its interfaces must be in either Route mode or function as a Layer 3 device. For interfaces in Layer 3 mode (that is have IP addresses assigned to the interfaces), you must configure a security policy (from zone to zone or intrazone) to permit the predefined service DHCP-Relay before forwarding occurs.

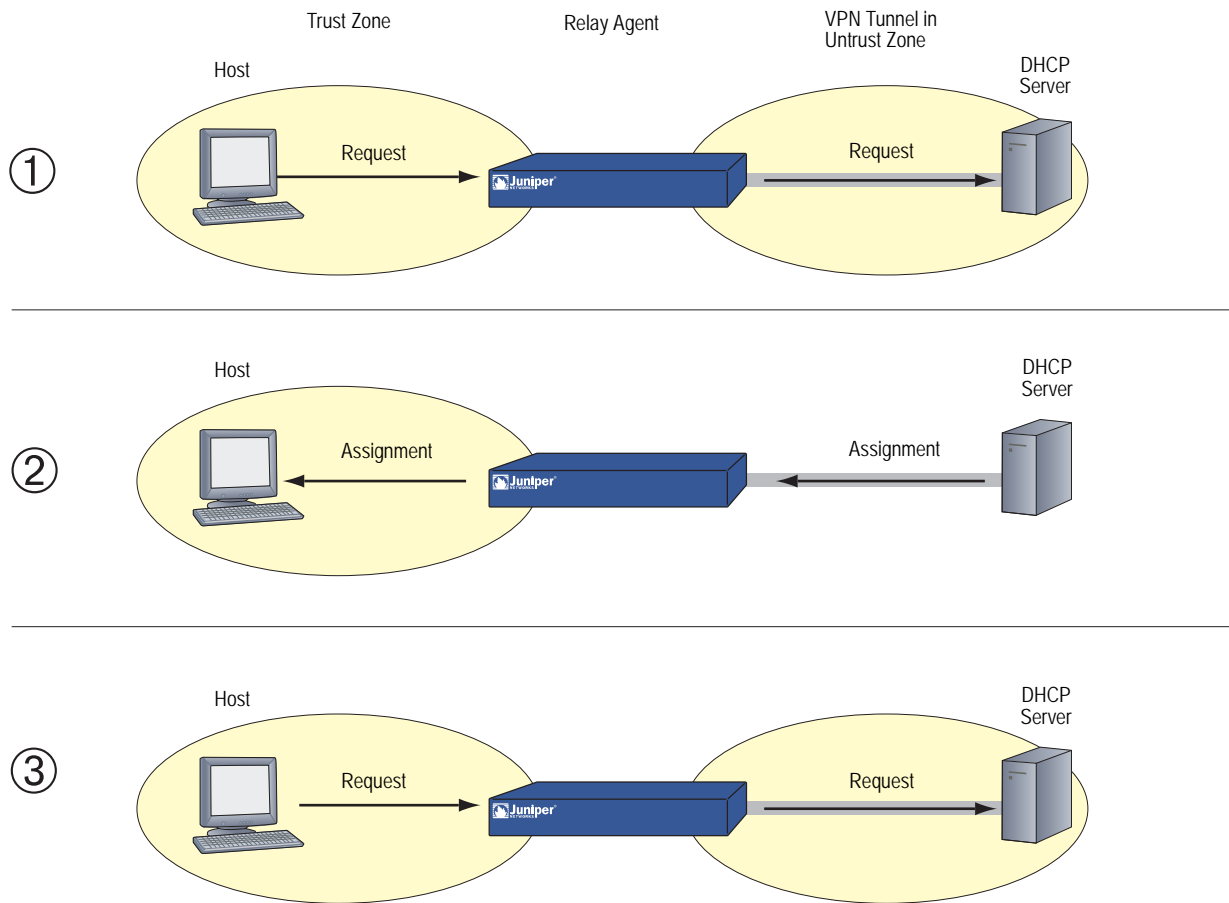
You can configure up to three DHCP servers for each DHCP relay agent. The relay agent unicasts an address request from a DHCP client to all configured DHCP servers. The relay agent forwards to the client all DHCP packets received from all servers. See “Forwarding All DHCP Packets” on page 237.

NOTE: When a security device acts as a DHCP relay agent, the device does not generate DHCP allocation status reports because the remote DHCP server controls all the IP address allocations.

ScreenOS supports DHCP relay in different vsys and for VLAN-tagged subinterfaces.


Figure 72 illustrates the process involved in using a security device as a DHCP relay agent. To ensure security, the DHCP messages pass through a VPN tunnel when traveling through the untrusted network.

Figure 72: DHCP Relay Agent Traffic



In Figure 73, a security device receives its DHCP information from a DHCP server at 194.2.9.10 and relays it to hosts in the Trust zone. The hosts receive IP addresses from an IP pool defined on the DHCP server. The address range is 180.10.10.2—180.10.10.254. The DHCP messages pass through a VPN tunnel between the local security device and the DHCP server, located behind a remote security device whose Untrust zone interface IP address is 2.2.2.2/24. The interface ethernet0/1 is bound to the Trust zone, has the IP address 180.10.10.1/24, and is in Route mode. The interface ethernet0/3 is bound to the Untrust zone and has the IP address 1.1.1.1/24. All security zones are in the trust-vr routing domain.

Figure 73: Device as DHCP Relay Agent



WebUI**1. Interfaces**

Interfaces > Edit (for ethernet0/1): Enter the following, then click **Apply**:

Zone: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 180.10.10.1/24

Enter the following, then click **OK**:

Interface Mode: Route

Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone: Untrust
 Static IP: (select this option when present)
 IP Address/Netmask: 1.1.1.1/24

2. Address

Policy > Policy Elements > Addresses > List > New: Enter the following, then click **OK**:

Address Name: DHCP Server
 IP Address/Domain Name:
 IP/Netmask: (select), 194.2.9.10/32
 Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, then click **OK**:

Gateway Name: dhcp server
 Security Level: Custom
 Remote Gateway Type:
 Static IP: (select), Address/Hostname: 2.2.2.2
 Outgoing Interface: ethernet0/3

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Security Level:
 User Defined: Custom (select)
 Phase1 Proposal: rsa-g2-3des-sha
 Mode (Initiator): Main (ID Protection)

VPNs > AutoKey IKE > New: Enter the following, then click **OK**:

VPN Name: to_dhcp
 Security Level: Compatible
 Remote Gateway:
 Predefined: (select), to_dhcp

> Advanced: Enter the following, then click **Return** to set the advanced options and return to the basic configuration page:

Bind to: None

4. DHCP Relay Agent

Network > DHCP > Edit (for ethernet0/1) > DHCP Relay Agent: Enter the following, then click **Apply**:

Relay Agent Server IP or Domain Name: 194.2.9.10
 Use Trust Zone Interface as Source IP for VPN: (select)

5. Route

Network > Routing > Destination > trust-vr New: Enter the following, then click **OK**:

Network Address/Netmask: 0.0.0.0/0
 Gateway: (select)
 Interface: ethernet0/3
 Gateway IP Address: 1.1.1.250

NOTE: Setting a route to the external router designated as the default gateway is essential for both outbound VPN and network traffic. In this example, the security device sends encapsulated VPN traffic to this router as the first hop along its route to the remote security device. In Figure 73, the concept is presented by depicting the tunnel passing through the router.

6. Policies

Policy > Policies > (From: Trust, To: Untrust) New: Enter the following, then click **OK**:

Source Address:
 Address Book Entry: (select), Any
 Destination Address:
 Address Book Entry: (select), DHCP Server
 Service: DHCP-Relay
 Action: Tunnel
 Tunnel VPN: to_dhcp
 Modify matching outgoing VPN policy: (select)

CLI

1. Interfaces

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 180.10.10.1/24
set interface ethernet0/1 route
set interface ethernet0/3 zone untrust
set interface ethernet0/3 ip 1.1.1.1/24
```

2. Address

```
set address untrust dhcp_server 194.2.9.10/32
```

3. VPN

```
set ike gateway "dhcp server" ip 2.2.2.2 main outgoing-interface ethernet0/3
proposal rsa-g2-3des-sha
set vpn to_dhcp gateway "dhcp server" proposal g2-esp-3des-sha
```

4. DHCP Relay Agent

```
set interface ethernet0/1 dhcp relay server-name 194.2.9.10
set interface ethernet0/1 dhcp relay vpn
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet0/3 gateway 1.1.1.250
```

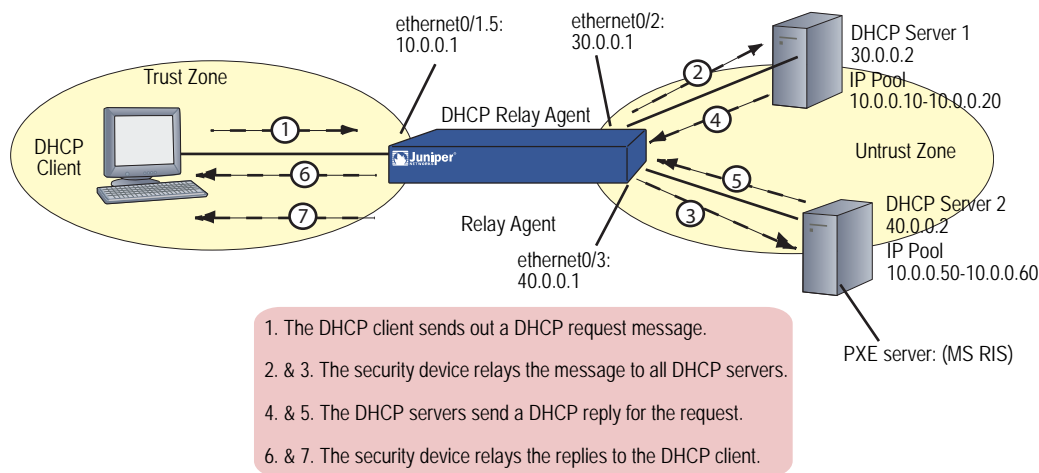
6. Policies

```
set policy from trust to untrust any dhcp_server dhcp-relay tunnel vpn to_dhcp
set policy from untrust to trust dhcp_server any dhcp-relay tunnel vpn to_dhcp
save
```

Forwarding All DHCP Packets

ScreenOS lets your security device relay all DHCP responses from multiple servers to a client. Some environments require multiple servers to respond with identical data and their clients only process the first-received response; other environments have multiple servers replying with unique data and the clients process appropriate data from each, for example in a Pre-Boot Execution Environment (PXE) scenario.

In common PXE cases (as shown in Figure 74), at least two DHCP servers serve clients. When the DHCP servers receive a request from the DHCP client, one of the servers, DHCP Server1, provides DHCP address information to the client while DHCP Server 2 (such as MS RIS) provides PXE information. This release of ScreenOS allows the security device to forward all DHCP packets to the client.

Figure 74: Relaying All DHCP Packets from Multiple DHCP Servers

Typically, a PXE server provides a boot-image-server for diskless PXE clients, which are diskless PC machines. When a PXE client powers on, it sends out a broadcast DHCP-DISCOVER (a kind of request), which means that the client requests the IP and boot-image path. In most cases, two kinds of servers serve the PXE: a PXE server (like a Microsoft RIS server) and a DHCP server. Both servers receive the DISCOVER request. The PXE server replies to the DISCOVER request with boot-image-server information. At the same time, the DHCP server replies to the DISCOVER request with an IP-assignment information. Both the responses from the two servers are forwarded to the DHCP client (diskless PC).

Configuring Next-Server-IP

If a security device receives conflicting or confusing information from the DHCP server, the device uses the IP address in the Next-Server-IP field. This DHCP configuration parameter has traditionally been used as the address of the TFTP server in the bootstrap process.

For example, in PXE scenarios, the first DHCP server serves the IP address, and the second DHCP server provides OS information. The Next-Server-IP field is configured to specify the next server in the chain. The chain and each member can vary from site to site. However, typically, it is a DHCP server chaining to a TFTP server. The chain is terminated either by supplying all zeroes (0.0.0.0) or by specifying the device interface IP into this field as shown in Table 37.

This Next-Server-IP information is returned in the siaddr field of the DHCP header and is often used to chain several bootstrap servers together, with each serving a specific function. The siaddr field is mandatory, if available, because some DHCP servers will place their own IP address in this field when it is not configured.

Table 37: Specifying Next-Server-IP

Next Server IP	Description
None (default)	siaddr = 0.0.0.0 (default)
Interface	siaddr = the IP interface bound to the DHCP server
Option66	siaddr = option66 (identifies the TFTP server for supporting diskless PCs)
Input	siaddr = custom IP address

If the Next-Server-IP is non-zero and not equal to this server's address, then it is interpreted by the client as the address of the next server in a chain that supplies additional boot information. In the following example, the Next-Server-IP is configured for Option66.

WebUI

Network > DHCP > Edit (DHCP Server): Select one of the following, then click **Apply**:

Next Server IP: From Option66

CLI

```
set interface ethernet0/1 dhcp server enable
set interface ethernet0/1 dhcp server option custom 66 ip 10.10.10.1
set interface ethernet0/1 dhcp server config next-server-ip option66
save
```

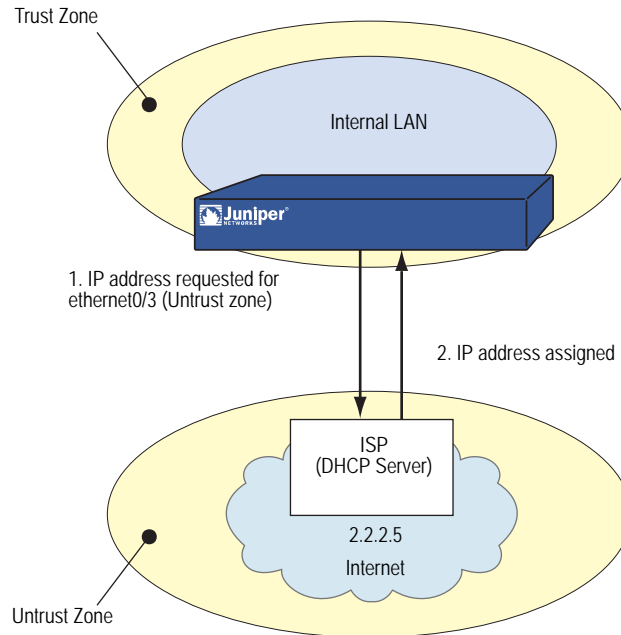
Using a Security Device as a DHCP Client

When acting as a DHCP client, the security device receives an IP address dynamically from a DHCP server for any physical interface in any security zone. If multiple interfaces bound to a single security zone exist, you can configure a DHCP client for each interface as long as each interface is not connected to the same network segment. If you configure a DHCP client for two interfaces that are connected to the same network segment, the first address assigned by a DHCP server is used. (If the DHCP client receives an address update to the same IP address, IKE is not rekeyed.)

NOTE: While some security devices can act as DHCP servers, a DHCP relay agents, or DHCP clients at the same time, you cannot configure more than one DHCP role on a single interface.

In this example, the interface bound to the Untrust zone has a dynamically assigned IP address. When the security device requests its IP address from its ISP, it receives its IP address, subnet mask, gateway IP address, and the length of its lease for the address. The IP address of the DHCP server is 2.2.2.5.

Figure 75: Device as DHCP Client



NOTE: Before setting up a site for DHCP service, you must have a Digital Subscriber Line (DSL) and an account with an Internet Service Provider (ISP).

WebUI

Network > Interfaces > Edit (for ethernet0/3): Select **Obtain IP using DHCP**, then click **OK**.

NOTE: You cannot specify the IP address of the DHCP server through the WebUI; however, you can do so through the CLI.

CLI

```
set interface ethernet0/3 dhcp client
set interface ethernet0/3 dhcp settings server 2.2.2.5
save
```

Propagating TCP/IP Settings

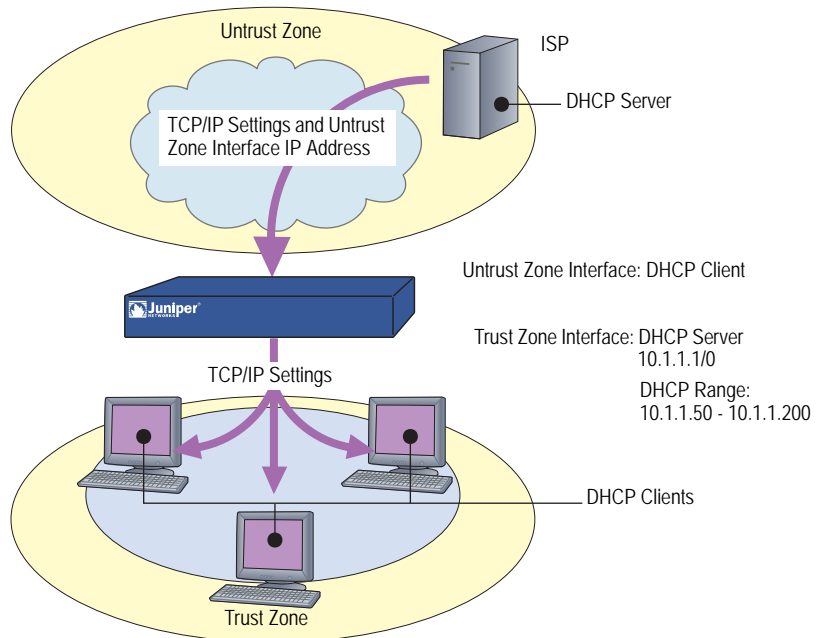
Some security devices can act as a Dynamic Host Control Protocol (DHCP) client, receiving its TCP/IP settings and the IP address for any physical interface in any security zone from an external DHCP server. Some security devices can act as a DHCP server, providing TCP/IP settings and IP addresses to clients in any zone. When a security device acts both as a DHCP client and a DHCP server simultaneously, it can transfer the TCP/IP settings learned through its DHCP client module to its default DHCP server module.

TCP/IP settings include the IP address of the default gateway and a subnet mask, and IP addresses for any or all of the following servers:

- DNS (3)
- WINS (2)
- NetInfo (2)
- SMTP (1)
- POP3 (1)
- News (1)

In Figure 76, the security device is both a client of the DHCP server in the Untrust zone and a DHCP server to the clients in the Trust zone. The device takes the TCP/IP settings that it receives as a DHCP client and forwards them as a DHCP server to the clients in the Trust zone. The Untrust Zone Interface is the DHCP client and receives IP addresses dynamically from an ISP.

Figure 76: DHCP Propagation



You can configure the DHCP server module to propagate all TCP/IP settings that it receives from the DHCP client module using the **set interface *interface* dhcp-client settings update-dhcpserver** command. You can also override any setting with a different one.

In this example, you configure the security device to act both as a DHCP client on the ethernet0/3 interface and as a DHCP server on the ethernet0/1 interface. (The default DHCP server is on the ethernet0/1 interface.)

As a DHCP client, the security device receives an IP address for the ethernet0/3 interface and its TCP/IP settings from an external DHCP server at 211.3.1.6. You enable the DHCP client module in the security device to transfer the TCP/IP settings it receives to the DHCP server module.

You configure the DHCP server module to do the following with the TCP/IP settings that it receives from the DHCP client module:

- Forward the DNS IP addresses to its DHCP clients in the Trust zone.
- Override the default gateway, netmask, SMTP server, and POP3 server IP addresses with the following:
 - 10.1.1.1 (this is the IP address of the ethernet0/1 interface)
 - 255.255.255.0 (this is the netmask for the ethernet0/1 interface)
 - SMTP: 211.1.8.150
 - POP3: 211.1.8.172

NOTE: If the DHCP server is already enabled on the Trust interface and has a defined pool of IP addresses (which is default behavior for certain platforms), you must first delete the IP address pool before you can change the default gateway and netmask.

You also configure the DHCP server module to deliver the following TCP/IP settings that it does not receive from the DHCP client module:

- Primary WINS server: 10.1.2.42
- Secondary WINS server: 10.1.5.90

Finally, you configure the DHCP server module to assign IP addresses from the following IP Pool to the hosts acting as DHCP clients in the Trust zone: 10.1.1.50 – 10.1.1.200.

WebUI

NOTE: You can set this feature only through the CLI.

CLI

1. DHCP Client

```
set interface ethernet0/3 dhcp-client settings server 211.3.1.6
set interface ethernet0/3 dhcp-client settings update-dhcpserver
set interface ethernet0/3 dhcp-client settings autoconfig
set interface ethernet0/3 dhcp-client enable
```

2. DHCP Server

```
set interface ethernet0/1 dhcp server option gateway 10.1.1.1
set interface ethernet0/1 dhcp server option netmask 255.255.255.0
set interface ethernet0/1 dhcp server option wins1 10.1.2.42
set interface ethernet0/1 dhcp server option wins2 10.1.5.90
set interface ethernet0/1 dhcp server option pop3 211.1.8.172
set interface ethernet0/1 dhcp server option smtp 211.1.8.150
set interface ethernet0/1 dhcp server ip 10.1.1.50 to 10.1.1.200
set interface ethernet0/1 dhcp server service
save
```

Configuring DHCP in Virtual Systems

DHCP: ScreenOS now fully supports DHCP relay for virtual systems. You can configure DHCP relay for a specific vsys and relay all packets from multiple DHCP servers to a client.

Setting DHCP Message Relay in Virtual Systems

ScreenOS allows you to configure Dynamic Host Configuration Protocol (DHCP) message relay from one or multiple DHCP servers to clients within a virtual system (vsys). You can configure DHCP message relay on an interface that is available to a virtual system.

If you have two DHCP servers, server 1 and server 2, a security device sitting between the DHCP servers and a client individually passes DHCP requests to each DHCP server on different outgoing interfaces. As each DHCP reply is received, the security device passes them to the root vsys and then forwards them to the appropriate DHCP client within a vsys.

Figure 77: DHCP Relay Services Within a Vsys



To configure DHCP with vsys:

1. Create a virtual system.
2. Enable DHCP for that vsys.
3. Configure a static route to allow the DHCP server in the root system to access the vsys.
4. Set security policies in the virtual system.

Point-to-Point Protocol over Ethernet

PPP-over-Ethernet (PPPoE) merges the Point-to-Point Protocol (PPP), which is usually used for dialup connections, with the Ethernet protocol, which can connect multiple users at a site to the same customer premises equipment. Many users can share the same physical connection, but access control, billing, and type of service are handled on a per-user basis. Some security devices support a PPPoE client, allowing them to operate compatibly on DSL, Ethernet Direct, and cable networks run by ISPs using PPPoE for their clients' Internet access.

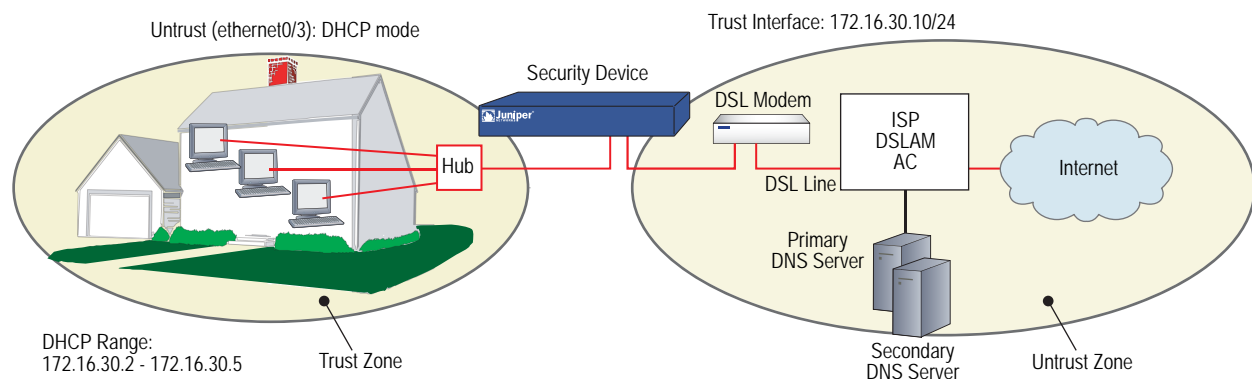
On devices that support PPPoE, you can configure a PPPoE client instance on any or all interfaces. You configure a specific instance of PPPoE with a username, password, and other parameters, and then you bind the instance to an interface. When two Ethernet interfaces (a primary and a backup) are bound to the Untrust zone, you can configure one or both interfaces for PPPoE.

Setting Up PPPoE

The following example illustrates how to define the untrusted interface of a security device for PPPoE connections and how to initiate PPPoE service.

In this example, the security device receives a dynamically assigned IP address for its Untrust zone interface (ethernet0/3) from the ISP, and the security device also dynamically assigns IP addresses for the three hosts in its Trust zone. In this case, the security device acts both as a PPPoE client and a DHCP server. The Trust zone interface must be in either NAT or Route mode. In this example, it is in NAT mode.

Figure 78: PPPoE



Before setting up the site in this example for PPPoE service, you must have the following:

- Digital subscriber line (DSL) modem and line
- Account with ISP
- Username and password (obtained from the ISP)

WebUI

1. Interfaces and PPPoE

Network > Interfaces > Edit (for ethernet0/1): Enter the following, then click **OK**:

Zone: Trust
 Static IP: (select this option when present)
 IP Address/Netmask: 172.16.30.10/24

Network > Interfaces > Edit (for ethernet0/3): Enter the following, then click **OK**:

Zone: Untrust
 Obtain IP using PPPoE: (select)
 Username/Password: *name/password*

Network > Interfaces > Edit (for ethernet0/3): To test your PPPoE connection, click **Connect**.

NOTE: When you initiate a PPPoE connection, your ISP automatically provides the IP addresses for the Untrust zone interface and for the Domain Name System (DNS) servers. When the security device receives DNS addresses by PPPoE, the new DNS settings overwrite the local settings by default. If you do not want the new DNS settings to replace the local settings, you can use the CLI command **unset pppoe dhcp-updateserver** to disable this behavior.

If you use a static IP address for the Untrust zone interface, you must obtain the IP addresses of the DNS servers and manually enter them on the security device and on the hosts in the Trust zone.

2. DHCP Server

Network > Interfaces > Edit (for ethernet0/1) > DHCP: Select **DHCP Server**, then click **Apply**.

Network > Interfaces > Edit (for ethernet0/1) > DHCP: Enter the following, then click **Apply**:

Lease: 1 hour
 Gateway: 0.0.0.0
 Netmask: 0.0.0.0
 DNS#1: 0.0.0.0

> Advanced: Enter the following, then click **Return**:

DNS#2: 0.0.0.0
 Domain Name: (leave blank)

Network > Interfaces > DHCP (for ethernet0/1) > New Address: Enter the following, then click **OK**:

Dynamic: (select)
 IP Address Start: 172.16.30.2
 IP Address End: 172.16.30.5

3. Activating PPPoE on the Security Device

1. Turn off the power to the DSL modem, the security device, and the three workstations.
2. Turn on the DSL modem.
3. Turn on the security device.

The security device makes a PPPoE connection to the ISP and, through the ISP, gets the IP addresses for the DNS servers.

4. Activating DHCP on the Internal Network

Turn on the workstations.

The workstations automatically receive the IP addresses for the DNS servers. They get an IP address for themselves when they attempt a TCP/IP connection.

NOTE: When you use DHCP to assign IP addresses to hosts in the Trust zone, the security device automatically forwards the IP addresses of the DNS servers that it receives from the ISP to the hosts.

If the IP addresses for the hosts are not dynamically assigned through DHCP, you must manually enter the IP addresses for the DNS servers on each host.

Every TCP/IP connection that a host in the Trust zone makes to the Untrust zone automatically goes through the PPPoE encapsulation process.

CLI

1. Interfaces and PPPoE

```
set interface ethernet0/1 zone trust
set interface ethernet0/1 ip 172.16.30.10/24
set interface ethernet0/3 zone untrust
set pppoe interface ethernet0/3
set pppoe username name_str password pswd_str
```

To test your PPPoE connection:

```
exec pppoe connect
get pppoe
```

2. DHCP Server

```
set interface ethernet0/1 dhcp server service
set interface ethernet0/1 dhcp server ip 172.16.30.2 to 172.16.30.5
set interface ethernet0/1 dhcp server option lease 60
save
```

3. Activating PPPoE on the Security Device

1. Turn off the power to the DSL modem, the security device, and the three workstations.
2. Turn on the DSL modem.
3. Turn on the security device.

4. Activating DHCP on the Internal Network

Turn on the workstations.

The workstations automatically receive the IP addresses for the DNS servers. They get an IP address for themselves when they attempt a TCP/IP connection.

Every TCP/IP connection that a host in the Trust zone makes to the Untrust zone automatically goes through the PPPoE encapsulation process.

Configuring PPPoE on Primary and Backup Untrust Interfaces

In the following example, you configure PPPoE for both the primary (ethernet0/3) and backup (ethernet0/2) interfaces to the Untrust zone.

WebUI

1. PPPoE Configuration for ethernet0/3 Interface

Network > PPP > PPPoE Profile > New: Enter the following, then click **OK**:

```
PPPoE instance: eth3-pppoe
Bound to interface: ethernet0/3 (select)
Username: user1
Password: 123456
Authentication: Any (select)
Access Concentrator: ac-11
```

2. PPPoE Configuration for ethernet0/2 Interface

Network > PPP > PPPoE Profile > New: Enter the following, then click **OK**:

```
PPPoE instance: eth2-pppoe
Bound to interface: ethernet0/2 (select)
Username: user2
Password: 654321
Authentication: Any (select)
Access Concentrator: ac-22
```

CLI**1. PPPoE Configuration for ethernet0/3 Interface**

```
set pppoe name eth3-pppoe username user1 password 123456
set pppoe name eth3-pppoe ac ac-11
set pppoe name eth3-pppoe authentication any
set pppoe name eth3-pppoe interface ethernet0/3
```

2. PPPoE Configuration for ethernet0/2 Interface

```
set pppoe name eth2-pppoe username user2 password 654321
set pppoe name eth2-pppoe ac ac-22
set pppoe name eth2-pppoe authentication any
set pppoe name eth2-pppoe interface ethernet0/2
save
```

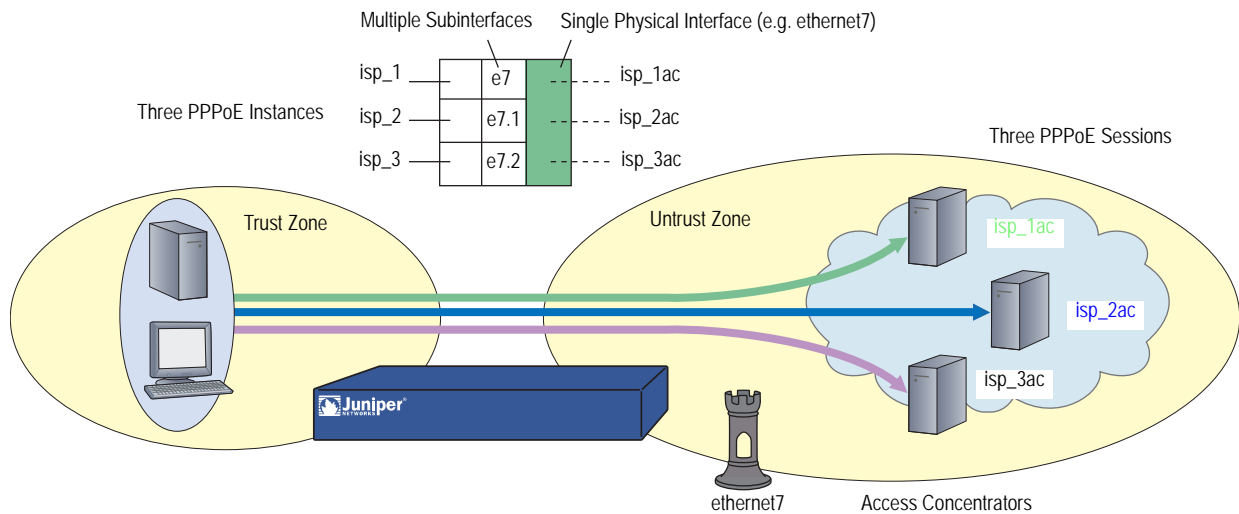
Configuring Multiple PPPoE Sessions over a Single Interface

Some security devices support creation of multiple PPPoE subinterfaces (each with the same MAC address) for a given physical interface. This support allows you to establish a private network connection with one ISP and connect to the Internet through a different ISP using the same physical interface. You can establish these connections using different user or domain names or be connected simultaneously to different ISPs.

The maximum number of concurrent PPPoE sessions on a physical interface is limited only by number of subinterfaces allowed by the device. There is no restriction on how many physical interfaces can support multiple sessions. You can specify username, static-ip, idle-timeout, auto-connect, and other parameters separately for each PPPoE instance or session.

To support a PPPoE session, a subinterface must be untagged. An untagged interface uses `encap` (not a VLAN tag) to identify a VLAN for a subinterface. `Encap` binds the subinterface to PPPoE encapsulation. By hosting multiple subinterfaces, a single physical interface can host multiple PPPoE instances. You can configure each instance to go to a specified Access Concentrator (AC), allowing separate entities such as ISPs to manage the PPPoE sessions through a single interface. For more information about VLANs and VLAN tags, see *Volume 10: Virtual Systems*.

Figure 79: PPPoE with Multiple Sessions



In the following example, you define three PPPoE instances, specify an Access Concentrator (AC) for each, then initiate each instance.

- Instance `isp_1`, username `user1@domain1`, password `swordfish`, bound to interface `ethernet7`. The AC is `isp_1ac`.
- Instance `isp_2`, username `user2@domain2`, password `marlin`, bound to subinterface `ethernet7.1`. The AC is `isp_2ac`.
- Instance `isp_3`, username `user3@domain3`, password `trout`, bound to subinterface `ethernet7.2`. The AC is `isp_3ac`.

WebUI

1. Interface and Subinterfaces

Network > Interfaces > Edit (for ethernet7):

Enter the following, then click **OK**:

Zone Name: Untrust

Network > Interfaces > New (Sub-IF):

Enter the following, then click **OK**:

Interface Name: ethernet7.1
Zone Name: Untrust

Network > Interfaces > New (Sub-IF):

Enter the following, then click **OK**:

Interface Name: ethernet7.2
Zone Name: Untrust

2. PPPoE Instances and AC

Network > PPP > PPPoE Profile > New:

Enter the following, then click **OK**:

```

PPPoE Instance: isp_1
Enable: Enable
Bound to Interface: ethernet7
Username: user1@domain1
Access Concentrator: isp_1ac

```

Network > PPP > PPPoE Profile > New:

Enter the following, then click **OK**:

```

PPPoE Instance: isp_2
Enable: Enable
Bound to Interface: ethernet7.1
Username: user2@domain2
Access Concentrator: isp_2ac

```

Network > PPP > PPPoE Profile > New:

Enter the following, then click **OK**:

```

PPPoE Instance: isp_3
Enable: Enable
Bound to Interface: ethernet7.2
Username: user3@domain3
Access Concentrator: isp_3ac

```

3. PPPoE Initiation

Network > PPP > PPPoE Profile > Connect (for isp_1)

Network > PPP > PPPoE Profile > Connect (for isp_2)

Network > PPP > PPPoE Profile > Connect (for isp_3)

CLI**1. Interface and Subinterfaces**

```

set interface ethernet7 zone untrust
set interface ethernet7.1 encaps pppoe zone untrust
set interface ethernet7.2 encaps pppoe zone untrust

```

2. PPPoE Instances and ACs

```

set pppoe name isp_1 username user1@domain1 password swordfish
set pppoe name isp_1 interface ethernet7
set pppoe name isp_1 ac isp_1ac
set pppoe name isp_2 username user2@domain2 password marlin
set pppoe name isp_2 interface ethernet7.1
set pppoe name isp_2 ac isp_2ac
set pppoe name isp_3 username user3@domain3 password trout
set pppoe name isp_3 interface ethernet7.2
set pppoe name isp_3 ac isp_3ac
save

```

3. PPPoE Initiation

```
exec pppoe name isp_1 connect
exec pppoe name isp_2 connect
exec pppoe name isp_3 connect
```

PPPoE and High Availability

Two security devices that support PPPoE in Active/Passive mode can handle failover of a PPPoE connection. Upon initiation of the connection, the primary device synchronizes its PPPoE state with the backup device. Because the passive device uses the same IP address as the primary device, it does not have to make a new PPPoE connection once it becomes the primary. Therefore, it can maintain communication with the Access Concentrator after failure of the primary. This is necessary when the PPPoE interface supports VPN connections, and these connections must continue, using the same interface IP after failover. For more information about HA configurations, see *Volume 11: High Availability*.

License Keys

The license key feature allows you to expand the capabilities of your Juniper Networks security device without having to upgrade to a different device or system image. You can purchase the following type of keys:

- Advanced
- Capacity
- Extended
- Virtualization
- GTP
- Vsys
- IDP

Each security device ships with a standard set of features enabled and might support the activation of optional features or the increased capacity of existing features. For information regarding which features are currently available for upgrading, refer to the latest marketing literature from Juniper Networks.

The procedure for obtaining and applying a license key is as follows:

1. Gather your authorization code and device serial number.

Authorization Code: A pass key required to generate and activate the license key that you or your company have purchased for your Juniper Networks security device. Note: The Authorization Code is required to generate your license key—it is not the actual license key.

Device Serial Number: A unique 16-character code Juniper Networks uses to identify your particular security device when generating license keys. You can find the device serial number at the bottom or back of the device. You can also find the serial number in the device information section in the GUI or by executing a “get system” command on the CLI.

2. Sign in to the Juniper Networks License Management System (LMS) at www.juniper.net/generate_license/, select the **Firewall/IPSec VPN and Intrusion Prevention** link, then follow the instructions in the system user interface.
3. The Juniper License Management System provides the license key in one of two ways:
 - Download the license key to your computer.
 - Receive an email that contains your license key.
4. Install the license key in one of the following ways:

WebUI

Configuration > Update > ScreenOS/Keys > Select **License Key Update (Features)** > click **Browse** > select the file with the license key, then click **Apply**.

CLI

```
exec license-key key_num
```

Registration and Activation of Subscription Services

Before your Juniper Networks security device can receive regular subscription service for antivirus (AV) patterns, Deep Inspection (DI) signatures, anti-spam, or web filtering, you must do the following:

- Purchase a subscription
- Register the subscription
- Retrieve the subscription key

Retrieving the subscription license key activates your services on the device for the time period purchased. Your specific service-activation process depends upon which services you purchased and the method you used to purchase them.

Trial Service

To allow you to use AV, DI, anti-spam, or web-filtering services, the security device provides a trial period. During this period, the device can obtain services on a temporary basis. To retrieve eligible trial license keys from the entitlement server, use the **exec license-key** update trials CLI command.

- No Juniper Networks security device comes with DI already enabled. To obtain trial DI service, you must start a WebUI session, then click **Retrieve Subscriptions Now** in Configuration > Update > ScreenOS/Keys. This action provides a one-time, one-day DI key.
- If your device has AV service bundled at the time of purchase, then the device has preinstalled trial service. This trial service lasts up to 60 days.
- Anti-spam
- No Juniper Networks security device comes with web filtering already enabled. This feature does not have a trial service.



CAUTION: To avoid service interruption, you must register your Juniper Networks security device as soon as possible after purchasing your subscription. Registration ensures continuation of the subscription.

Updating Subscription Keys

If there is any software with an expiration date installed in the security device, the device periodically connects to the entitlement server to retrieve the subscription keys. The device connects to the entitlement server, License Management System, under all of following conditions:

- Key expires in two months
- Key expires in one month
- Key expires in two weeks
- Key expires
- 30 days after key expires

NOTE: To delete a single license key from the key file, use the **exec license-key delete name_str** CLI command.

Adding Antivirus, Web Filtering, Anti-Spam, and Deep Inspection to an Existing or a New Device

After purchasing AV, web-filtering, anti-spam, or deep inspection (DI) subscriptions to add to your existing Juniper Networks security device, perform the following steps to activate the services:

1. After ordering the subscription, you should receive an authorization code, via email, from Juniper Networks or your authorized VAR. This code is a readable document that contains information you need to register your subscription.
2. Make sure the device is registered. If it is not currently registered, go to the following site:

<http://tools.juniper.net/subreg/>
3. Register the subscription authorization code to the device.
4. Confirm that your device has Internet connectivity.
5. Retrieve the subscription key on the device. You can do this in one of two ways:
 - In the WebUI, click **Retrieve Subscriptions Now** from Configuration > Update > ScreenOS/Keys.
 - Using the CLI, run the following command:

`exec license-key update`
6. You must reset the device after the key has been loaded.

You can now configure the device to automatically or manually retrieve the signature services. For instructions on configuring your security device for these services, see the following sections:

- “Fragment Reassembly” on page 4-58
- “Antivirus Scanning” on page 4-62
- “Web Filtering” on page 4-100

System Clock

It is important that your Juniper Networks security device always be set to the right time. Among other things, the time on your device affects the set up of VPN tunnels and the timing of schedules. First, to ensure that the device always maintains the proper time, you must set the system clock to the current time. Next, you can enable the daylight saving time (DST) option, and you can configure up to three (Network Time Protocol) NTP servers (one primary and two backups) from which the device can regularly update its system clock.

Date and Time

To set the clock to the current date and time, you can use the WebUI or the CLI. Through the WebUI, you do this by synchronizing the system clock with the clock on your computer:

1. Configuration > Date/Time: Click the **Sync Clock with Client** button.

A pop-up message prompts you to specify if you have enabled the DST option on your computer clock.

2. Click **Yes** to synchronize the system clock and adjust it according to DST, or click **No** to synchronize the system clock without adjusting it for DST.

Through the CLI, you set the clock by manually entering the date and time using the following command:

```
set clock mm/dd/yyyy hh:mm:ss
```

Daylight Saving Time

Daylight saving time (DST) is a widely used system of adjusting the official local time forward in summer months in order to save energy and allow more daylight for work and school activities.

DST is observed differently in various countries. Accordingly, you can choose the appropriate DST settings that apply to your country.

You can set the DST adjustment to recur on a weekday schedule (for example, the first Sunday in April) or on a specific date. You also can set DST adjustment not to recur, in which case you can adjust DST settings only for a single year.

Time Zone

You set the time zone by specifying the number of hours by which the local time for the security device is behind or ahead of GMT (Greenwich Mean Time). For example, if the local time zone for the device is Pacific Standard Time, it is 8 hours behind GMT. Therefore, you have to set the clock to **-8**.

If you set the time zone using the WebUI:

```
Configuration > Date/Time > Set Time Zone_hours_minutes from GMT
```

If you set the time zone using the CLI:

```
device -> set clock timezone number (a number from -12 to 12)
```

or

```
device-> set ntp timezone number (a number from -12 to 12)
```

Network Time Protocol

To ensure that the security device always maintains the right time, it can use Network Time Protocol (NTP) to synchronize its system clock with that of an NTP server over the Internet. You can do this manually or configure the device to perform this synchronization automatically at time intervals that you specify.

Configuring Multiple NTP Servers

You can configure up to three NTP servers on a Juniper Networks security device: one primary server and two backup servers. When you configure the security device to synchronize its system clock automatically, it queries each configured NTP server sequentially. The device always queries the primary NTP server first. If the query is not successful, the device then queries the first backup NTP server and so on until it gets a valid reply from one of the NTP servers configured on the device. The device makes four attempts on each NTP server before it terminates the update and logs the failure.

When you manually synchronize the system clock, and you can only do this using the CLI, you can specify a particular NTP server or none at all. If you specify an NTP server, the security device queries that server only. If you do not specify an NTP server, the device queries each NTP server configured on the device sequentially. You can specify an NTP server using its IP address or its domain name.

Configuring a Backup NTP Server

You can specify an individual interface as the source address to direct NTP requests from the device over a VPN tunnel to the primary NTP server or to a backup server as necessary. From among the other interface types available, you can select a loopback interface to perform this function.

The security device sends NTP requests from a source interface and optionally uses an encrypted, preshared key when sending NTP requests to the NTP server. The key provides authentication.

In the following example, you configure the primary NTP server and two backup NTP servers by assigning an IP address to each server. Next, you set each server to send NTP requests from the Trust interface. After that, you set the key-id and preshared key for each server.

WebUI

Configuration > Date/Time: Enter the following, then click **Apply**:

```

Primary Server IP/Name: 1.1.1.1
Primary server Key ID: 10
Source interface: Select Trust from the list.
Preshared Key: !2005abc
Backup Server1 IP/Name: 1.1.1.2
Primary server Key ID: 10
Source interface: Select Trust from the list.
Preshared Key: !2005abc
Backup Server2 IP/Name: 1.1.1.3
Primary server Key ID: 10
Source interface: Select Trust from the list.
Preshared Key: !2005abc

```

CLI

```

set ntp server 1.1.1.1
set ntp server backup1 1.1.1.2
set ntp server backup2 1.1.1.3
set ntp server src-interface trust
set ntp server backup1 src-interface trust
set ntp server backup2 src-interface trust
set ntp server key-id 10 pre-share-key !2005abc
set ntp server backup1 key-id 10 pre-share-key !2005abc
set ntp server backup2 key-id 10 pre-share-key !2005abc
save

```

Device as an NTP Server

A security device can also work as an NTP server serving NTP requests from its subnet peers. To use this feature, you need to enable NTP service on any of the Layer 3 interface with an IP address.

NOTE: Currently, ScreenOS supports only Unicast mode.

WebUI

Network > Interfaces > Edit (for ethernet0/0) > Basic: check the NTP Server check box, then click **Apply**.

CLI

```

set interface interface ntp-server
save

```

Maximum Time Adjustment

For automatic synchronization, you can specify a maximum time adjustment value (in seconds). The maximum time adjustment value represents the acceptable time difference between the security device system clock and the time received from an NTP server. The device only adjusts its clock with the NTP server time if the time difference between its clock and the NTP server time is within the maximum time adjustment value that you set. For example, if the maximum time adjustment value is 3 seconds, and the time on the device system clock is 4:00:00 and the NTP server sends 4:00:02 as the time, the difference in time between the two is acceptable and the device can update its clock. If the time adjustment is greater than the value you set, the device does not synchronize its clock and proceeds to try the first backup NTP server configured on the device. If the device does not receive a valid reply after trying all the configured NTP servers, it generates an error message in the event log. The default value for this feature is 3 seconds and the range is 0 (no limit) to 3600 (one hour).

When you manually synchronize the system clock, and you can only do this using the CLI, the security device does not verify the maximum time adjustment value. Instead, if it receives a valid reply, the device displays a message informing you of which NTP server it reached, what the time adjustment is, and the type of authentication method used. The message also asks you to confirm or cancel the system clock update.

If the security device does not receive a reply, it displays a timeout message. This message appears only after the device unsuccessfully attempted to reach all NTP servers configured on the device.

NOTE: When issuing requests using the CLI, you can cancel the current request by pressing **Ctrl-C** on the keyboard.

NTP and NSRP

NetScreen Redundancy Protocol (NSRP) contains a mechanism for synchronizing the system clock of NSRP cluster members. Although the resolution for synchronization is in seconds, NTP has sub-second resolution. Because the time on each cluster member might differ by a few seconds due to processing delays, Juniper Networks recommends that you disable NSRP time synchronization when NTP is enabled on both cluster members and they can each update their system clock from an NTP server. To disable the NSRP time synchronization function, enter the **set ntp no-ha-sync** CLI command.

Setting a Maximum Time Adjustment Value to an NTP Server

In the following example you configure the security device to update its clock at five-minute intervals from NTP servers at IP addresses 1.1.1.1, 1.1.1.2, and 1.1.1.3. You also set a maximum time adjustment value of 2 seconds.

WebUI

Configuration > Date/Time: Enter the following, then click **Apply**:

```

Automatically synchronize with an Internet Time Server (NTP): (select)
Update system clock every minutes: 5
Maximum time adjustment seconds: 2
Primary Server IP/Name: 1.1.1.1
Backup Server1 IP/Name: 1.1.1.2
Backup Server2 IP/Name: 1.1.1.3

```

CLI

```

set clock ntp
set ntp server 1.1.1.1
set ntp server backup1 1.1.1.2
set ntp server backup2 1.1.1.3
set ntp interval 5
set ntp max-adjustment 2
save

```

Securing NTP Servers

You can secure NTP traffic by using MD5-based checksum to provide authentication of NTP packets. You do not need to create an IPSec tunnel. This type of authentication ensures the integrity of NTP traffic. It does not prevent outside parties from viewing the data, but it prevents anyone from tampering with it.

To enable the authentication of NTP traffic, you must assign a unique key ID and preshared key to each NTP server you configure on a security device. The key ID and preshared key serve to create a checksum with which the security device and the NTP server can authenticate the data.

Table 38 describes the two types of authentication for NTP traffic.

Table 38: NTP Traffic Authentication

Type	Description
Required	When selected, the security device must include the authentication information—key id and checksum—in every packet it sends to an NTP server and must authenticate all NTP packets it receives from an NTP server. Before authentication can occur between a security device and an NTP server, the administrators of the security device and the NTP server must first exchange a key id and a preshared key. They have to exchange these manually and can do so in different ways such as via email or telephone.
Preferred	When selected, the security device must first operate as in Required mode by trying to authenticate all NTP traffic. If all attempts to authenticate fail, the security device then operates as if you selected no authentication. It sends out packets to an NTP server without including a key id and checksum. Essentially, although there is a preference for authentication, if authentication fails, the security device still permits NTP traffic.

Index

- A**
- access policies
 - See policies
 - address books
 - addresses
 - adding..... 104
 - modifying..... 105
 - removing..... 108
 - entries..... 104
 - group entries, editing..... 108
 - groups..... 105
 - See also addresses
 - address groups..... 105, 166
 - creating..... 107
 - editing..... 108
 - entries, removing..... 108
 - options..... 106
 - address negation 186
 - addresses
 - address book entries..... 104 to 108
 - defined 166
 - in policies 166
 - IP, host and network IDs..... 47
 - private..... 47
 - public..... 47
 - aggregate interfaces 37
 - alarms
 - thresholds..... 172
 - ALGs
 - for custom services..... 167
 - MS RPC..... 128
 - RTSP 129
 - Sun RPC..... 127
 - application option, in policies..... 167
 - ARP..... 82
 - ARP, ingress IP address..... 84
 - auth users
 - pre-policy auth..... 171
 - run-time auth process..... 170
 - run-time authentication..... 170
 - WebAuth 171
 - authentication
 - Allow Any..... 171
 - policies..... 170
 - users 170
- B**
- bandwidth 173
 - guaranteed 173, 194, 199
 - managing..... 194
 - maximum..... 173, 194, 199
 - maximum, unlimited..... 194
 - priority 198, 199
- bridge groups
 - logical interface 37
 - unbinding..... 46
- C**
- CLI, set arp always-on-dest..... 74, 77
 - clock, system
 - See system clock
 - custom services..... 121
 - custom services, in root and vsys 122
- D**
- DHCP..... 96, 100, 243
 - client 225
 - HA 231
 - PXE scenario..... 237
 - relay agent 225
 - server 225
 - DiffServ 173, 200, 214
 - See also DS Codepoint Marking
 - DIP..... 98, 140 to 143
 - fix-port..... 142
 - groups..... 153 to 155
 - PAT..... 141, 142
 - pools 169
 - pools, modifying..... 143
 - DNS 217
 - addresses, splitting..... 223
 - lookups 218
 - lookups, domain..... 223
 - servers 244
 - servers, tunneling to 223
 - status table 219
 - Domain Name System
 - See DNS
 - DS Codepoint Marking..... 194, 200, 214
 - DSL..... 239, 244
 - dynamic IP (DIP) pools 143, 169
- F**
- function zone interfaces 38
 - HA 38
 - management..... 38

G	
graphs, historical	172
groups	
addresses	105
services.....	138
H	
HA	
DHCP.....	231
interfaces, virtual HA.....	39
<i>See also</i> NSRP	
historical graphs	172
I	
ICMP services.....	126
message codes	126
message types.....	126
interfaces	
addressing.....	46
aggregate.....	37
binding to zone	44
connections, monitoring.....	63
default	48
DIP	140
down, logically	61
down, physically	61
function zone.....	38
HA function zone	38
interface tables, viewing.....	43
IP tracking (<i>See</i> IP tracking)	
L3 security zones	46
loopback.....	58
MGT	38
modifying.....	48
physical in security zones	36
policy-based NAT tunnel.....	39
redundant	37
secondary IP addresses.....	50
state changes.....	61
tunnel	39, 39 to 42
up, logically.....	61
up, physically.....	61
viewing interface table	43
virtual HA.....	39
VLAN1	81
VSI.....	38
zones, unbinding from.....	45
interfaces, monitoring.....	68 to 73
loops.....	69
security zones.....	73
Internet Service Provider (ISP).....	223
IP addresses	
host IDs	47
interfaces, tracking on.....	63
L3 security zones	46 to 47
Manage.....	95
network IDs	47
ports, defining for each.....	104
private	46
private address ranges	47
public.....	46
secondary.....	50, 51
secondary, routing between.....	51
IP pools	
<i>See</i> DIP pools	
IP tracking	
dynamic option	64
interfaces, shared.....	64
interfaces, supported.....	63
object failure threshold.....	65
rerouting traffic	63 to 78
vsys.....	64
weights	65
IP tracking, failure	
egress interface, on.....	75 to 76
ingress interface, on	76 to 78
tracked IP threshold.....	64
ISP	223
K	
keys, license	250
L	
L2TP policies	168
license keys	250
logging	172
loopback interfaces	58
M	
Manage IP.....	95
MGT interface.....	38
MIP	11
MIP, to zone with interface-based NAT.....	94
modes	
NAT, traffic to Untrust zone.....	79
Transparent.....	80
MS RPC ALG, defined.....	128
N	
NAT mode.....	92 to 97
interface settings.....	95
traffic to Untrust zone	79, 94
NAT-src, Route mode.....	98
negation, address	186
NetInfo	226
netmasks	47, 166
network, bandwidth.....	193
NSRP	

- DHCP 231
 - DIP groups 153 to 155
 - HA session backup 171
 - NTP synchronization 257
 - redundant interfaces 37
 - VSIs 38
 - NTP 255 to 258
 - authentication types 258
 - maximum time adjustment 256
 - multiple servers 255
 - NSRP synchronization 257
 - secure servers 257
 - servers 255
- P**
- packet flow 10 to 12
 - PAT 141
 - physical interface
 - logical interface 36
 - policies 3
 - actions 167
 - address groups 166
 - address negation 186
 - addresses 166
 - addresses in 166
 - alarms 172
 - application, linking service to explicitly 167
 - authentication 170
 - bidirectional VPNs 168
 - changing 189
 - counting 172
 - Deep Inspection (DI) 169
 - deny 167
 - DIP groups 154
 - disabling 189
 - editing 189
 - enabling 189
 - functions of 159
 - global 162, 174, 184
 - HA session backup 171
 - ID 166
 - internal rules 164
 - interzone 161, 174, 175, 178
 - intrazone 161, 174, 182
 - L2TP 168
 - L2TP tunnels 168
 - lookup sequence 163
 - management 174
 - managing bandwidth 194
 - maximum limit 107
 - multiple items per component 185
 - name 168
 - NAT-dst 169
 - NAT-src 169
 - order 190
 - permit 167
 - policy context 185
 - policy set lists 163
 - position at top 169, 190
 - reject 167
 - removing 191
 - reordering 190
 - required elements 160
 - root system 164
 - schedules 172
 - security zones 166
 - service book 109
 - service groups 138
 - services 166
 - services in 109, 166
 - shadowing 189, 190
 - traffic logging 172
 - traffic shaping 173
 - tunnel 167
 - types 161 to 162
 - verifying 189
 - virtual systems 164
 - VPN dialup user groups 166
 - VPNs 168
 - policy-based NAT, tunnel interfaces 39
 - Port Address Translation
 - See PAT
 - priority queuing 198
 - private addresses 47
 - public addresses 47
 - PXE 237
 - PXE server 237
- Q**
- QoS 193
- R**
- RFCs
 - 0792, *Internet Control Message Protocol* 126
 - 1349, *Type of Service in the Internet Protocol Suite* .. 173
 - 1918, *Address Allocation for Private Internets* 47
 - 2132, *DHCP Options and BOOTP Vendor Extensions* 230
 - 2326, *Real Time Streaming Protocol (RTSP)* 130, 134
 - 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers* 173
 - Route mode 98 to 101
 - interface settings 99
 - NAT-src 98
 - RSH ALG 127
 - RTSP ALG

defined	129	ALGs.....	167
request methods	131	in vsys.....	122
server in private domain.....	134	subinterfaces	3
server in public domain	136	creating (root system).....	49
status codes	133	deleting.....	50
rules, derived from policies.....	164	subscriptions	
run-time authentication.....	170	registration and activation	251 to 253
		temporary service	252
S		Sun RPC ALG	
schedules.....	156, 172	call scenarios	127
SCREEN, MGT zone	28	defined	127
ScreenOS		system clock.....	253 to 258
function zones	33	date & time	254
global zone.....	28	sync with client	254
overview.....	1	time zone	254
packet flow	10 to 12	system parameters.....	257
policies	3		
security zones.....	2, 28	T	
security zones, global.....	2	tags, VLANs	3
security zones, predefined.....	2	time zone.....	254
tunnel zones	29	trace-route	85
virtual systems	9	traffic	
zones	25 to 33	counting	172
ScreenOS interfaces		logging.....	172
security zones.....	3	priority.....	173
subinterfaces	3	shaping.....	193
secondary IP addresses	51	traffic shaping	193
security zones	2	automatic	194
determination, destination zone.....	12	service priorities	198
determination, source zone.....	10	Transparent mode	80 to 92
global.....	2	ARP/trace-route	83
predefined.....	2	blocking non-ARP traffic	82
security zones, interfaces	3	blocking non-IP traffic	82
physical	36	broadcast traffic	81
service book		flood.....	83
entries, modifying (CLI).....	123	routes.....	82
entries, removing (CLI).....	123	unicast options	83
service book, services		tunnel interfaces.....	39
adding.....	122	definition.....	39
custom.....	109	policy-based NAT	39
custom (CLI).....	122		
preconfigured	109	U	
service groups.....	138 to 140	unknown unicast options	82 to 87
creating	138	ARP	84 to 87
deleting.....	140	flood.....	83 to 84
modifying.....	139	trace-route.....	85
service groups (WebUI).....	138	URL filtering	
services	109	See web filtering	
defined	166		
drop-down list	109	V	
ICMP	126	VIP.....	11
in policies.....	166	VIP, to zone with interface-based NAT	94
timeout threshold.....	123	virtual HA interfaces.....	39
services, custom	121	virtual routers	

<i>See</i> VRs	
virtual systems	9
VLAN zone	81
VLAN1	
interface	81, 87
zones	81
VLANs, tags	3
VPNs	
policies	168
to zone with interface-based NAT	94
tunnel zones	29
VRs	
forwarding traffic between	4
introduction	4
W	
web filtering	172
WebAuth, pre-policy auth process	171
wireless interface	
logical interface	36
Z	
zones	25 to 33
defining	30
editing	31
function	33
function, MGT interface	38
global	28
global security	2
Layer 2	81
tunnel	29
VLAN	33, 81
zones, ScreenOS	25 to 33
predefined	2
security interfaces	3
zones, security	2, 28
determination, destination zone	12
determination, source zone	10
global	2
interfaces, monitoring	73
interfaces, physical	36

